Alma Mater Studiorum – Università di Bologna

DOTTORATO DI RICERCA IN

Culture Letterarie e Filologiche

Ciclo XXXVI

Settore Concorsuale: Area 11/A4

Settore Scientifico Disciplinare: M-STO/08

DIVING INTO ARTWORKS INTERPRETATIONS THROUGH THE LENSES OF SEMANTIC DATA: AN APPLICATION ON PANOFSKY'S ICONOLOGICAL STUDIES

Presentata da: Sofia Baroncini

Coordinatore Dottorato

Prof. Marco Antonio Bazzocchi

Supervisore

Prof.ssa Francesca Tomasi

Co-supervisori

Dott.ssa Marilena Daquino
Prof. Aldo Gangemi

Esame finale anno 2024

Abstract

Iconology and Iconography are branches of art history that focus on recognizing and interpreting the subject matter represented in visual artworks and the deeper meanings possibly conveyed. Due to its complex and interpretative nature, iconographical and iconological knowledge is usually recorded in catalog entries in free text fields, making information retrieval challenging. However, having such curated data described with the accuracy and flexibility offered by ontologies would allow us not only to explore the art history scholarly discourse through quantitative analysis but also to enable a new way to access cultural heritage objects described in the network of Linked Open Data (LOD) through the narratives of experts' interpretations. The current doctoral thesis aimed to fill this gap by presenting 1) an ontology for describing iconographical and iconological interpretations, 2) an RDF dataset including a selection of iconographical and iconological interpretations by the art historian Erwin Panofsky, chosen as a case study, and 3) a quantitative analysis to verify whether data-driven iconological inquiries can be valuable for the domain. The resulting Iconology dataset, described through the newly created ICON ontology, covers core aspects of the domain according to three levels of interpretation and provides every subject recognition with provenance information. Although limited to the interests of the art historian, results show that the availability of such curated, authoritative data, described with a high level of granularity, allows us to quantitatively address traditional art historical questions on several topics and conduct new historiographical inquiries on the art historian's method. Furthermore, the implementation of exploratory data analysis on an online interactive dashboard shows the potential of domain-specific information retrieval. This contribution opens up the possibility to link artworks at their content level, allowing content-based research questions in art history to cross into the linked open data realm.

Keywords: Iconography, Iconology, Art Interpretation, Digital Art History, Linked Open Data, Semantic Web Data Analysis

. . . . . .

To my grandma,
who desired to study
in a time when it was not possible
to do so.

. . . . . .

Acknowledgements

I am incredibly grateful to have spent my Ph.D. years within the Digital Humanities research group at the University of Bologna. This group taught me many things, primarily how much humanity can be maintained even in professional practice. You have taught me by daily example what collaboration and selfless sharing of knowledge means. I am talking about more than just the open-access science in our academic practice, primarily relational, social, and human science. None of this would have been possible without your help and contribution.

I want to thank my supervisor, Francesca Tomasi, one of the first people who believed in the project's potential and supported it, always being there in times of need. Above all, I thank Marilena Daquino, who has been by my side week by week, accompanying me in this complicated challenge I have chosen. She has supported me in the uncertainties, relapses, and human confidences of a path that was not easy and asked me to put myself on the line in different aspects. Doing research means throwing oneself into the sea with the desire to explore it, continuing to swim on each passing wave enthusiastically, knowing that one never has sufficient means to face it. I want to thank all my colleagues, especially Bruno Sartini, for being a key reference point and for the naturalness with which they spent time and expertise to help me with Python, RDF, SPARQL, and more. Having had the good fortune of meeting you has made me a better person.

During these years, I have had several opportunities to learn and explore Semantic Web technologies applied to art history. First of all, I thank the lecturers of the Digital Humanities Master's course at the University of Bologna, thanks to whom I acquired the first tools to move into the world of digital humanities. I thank the organizers of the Digital Art History Summer School (DAHSS) and the International Semantic Web Summer School (ISWS) for creating international environments of exchange and learning, in which I could exchange ideas and thoughts with both students and experts. I thank Leuphana University in Lüneburg and the Provenance Lab research group for hosting me during my visiting period and sharing their exciting project on Digital Art History applied to the provenance of artworks.

Heartfelt thanks go to the Sony CSL Paris research center and the MUHAI project for hosting and allowing me to actively participate in the life of such an inspiring laboratory of ideas. In particular, I thank Luc Steels, who believed in my potential following a simple chat and conveyed to me how important it is to maintain passion, enthusiasm, and humanity in one's research path. I thank Remi Van Trijp, who supervised and enriched the project during my time at Sony CSL and left me with precious advice. Thanks also go to the members of the VUB Artificial Intelligence Lab in Brussels, who had been a fundamental support for the project's technical issues.

I thank the reviewers for the care with which they read and gave suggestions to improve the work.

Last but not least, I thank my family, friends, and my partner for their love and constant support, which have been fundamental reference points along my path, being by my side during challenging and exciting moments.

Contents

Abstract i

Table of Contents vii

List of Figures xiv

List of Tables xx

List of Abbreviations xxiii

1 Introduction 3

I BACKGROUND AND STATE OF THE ART 7

2 The study of artworks content and meaning 9
- 2.1 What is an interpretation act . . . 10
- 2.2 Iconography and Iconology . . . 12
  - 2.2.1 Aby Warburg and his influence on following scholars . . . 13
    - 2.2.1.1 Aby Warburg's method . . . 14
    - 2.2.1.2 Scholars of the Warburg's school . . . 16
  - 2.2.2 Erwin Panofsky and the theorization of interpretation act . . . 18
  - 2.2.3 Other theoretical approaches proposed formalizing the interpretation act and criticism towards Panofsky . . . 21
- 2.3 The adopted theoretical perspective . . . 25

3 Art description and analysis in the computational era 27
- 3.1 Cataloguing standards . . . 27
- 3.2 Controlled vocabularies . . . 28
- 3.3 Ontologies . . . 30
  - 3.3.1 Art history . . . 31
  - 3.3.2 Interpretations and meaning . . . 32
- 3.4 Computer vision . . . 33
- 3.5 Data analysis for art history . . . 34
- 3.6 Iconography and Iconology statements on the Semantic Web: a survey 35
  - 3.6.1 Knowledge graph evaluation methods . . . 37
  - 3.6.2 Selection of the knowledge graphs . . . 39
  - 3.6.3 Evaluation criteria . . . 41

3.6.3.1 Evaluation Methodology . . . 43
3.6.3.2 Iconographical and Iconological Schema Granularity 44
3.6.3.3 Semantic Validity of Iconographical and Iconological Triples . . . 46
3.6.3.4 Iconographical and Iconological Column Completeness . . . 48
3.6.3.5 Interlinking of artworks . . . 49
3.6.3.6 References to external taxonomies of Art and Culture 50
3.6.3.7 Intralinking potential for subject comparisons . . . 51
3.6.4 Results and discussion . . . 53

II SEMANTIC WEB AND ICONOLOGICAL INTERPRETATIONS 57

4 Approach 59
4.1 Research questions, hypotheses and assumptions . . . 59
4.2 Approach to the research . . . 62
4.2.1 Domain study . . . 62
4.2.1.1 Definition of domain characteristics . . . 63
4.2.1.2 Definition of domain interests . . . 63
4.2.2 Ontological modeling . . . 64
4.2.2.1 Ontology alignment . . . 65
4.2.3 Dataset creation . . . 66
4.2.3.1 Dataset creation workflow . . . 66
4.2.3.2 Database creation . . . 68
4.2.3.3 Dataset creation . . . 68
4.2.3.4 Technical requirements . . . 71
4.2.4 Ontology and data evaluation . . . 71
4.2.5 Analysis . . . 73
4.2.5.1 Iconographical and iconological research questions . 73
4.2.5.2 New inquiries . . . 73

5 Domain study and the ICON Ontology 77
5.1 Domain study . . . 77
5.1.1 Characteristics of the domain . . . 77
5.1.2 Interests of the domain . . . 82
5.1.2.1 Cultural phenomena . . . 82
5.1.2.2 Symbols . . . 84
5.1.2.3 Iconographies . . . 86
5.1.2.4 Evidence . . . 88
5.1.2.5 Visual Citation . . . 88
5.2 Requirements . . . 91
5.3 Modeling . . . 92
5.3.1 First design iteration: Recognitions . . . 92
5.3.2 Second design iteration: Pre-iconographical Recognitions (level 1) . . . 96
5.3.3 Third design iteration: Iconographical Recognitions (level 2) 97
5.3.4 Fourth design iteration: Iconological Recognitions (level 3) . 99

5.3.5 Ontology extension for simple interpretations 100
5.4 Refactoring: reuse and alignment to relevant ontologies and ontology design patterns 103
5.4.0.1 Recognitions as situations 105
5.4.0.2 Interpretations as descriptions 107
5.4.0.3 Describing artwork content 108
5.4.0.4 Interpretation details 108
5.4.0.5 Subjects 110
5.4.0.6 Symbols 110
5.4.0.7 Expression of Style 111
5.4.0.8 Citations, sources, evidences 112

6 The Iconology dataset **113**
6.1 Sample selection and description 113
6.2 Database creation 114
6.2.1 Database modeling 114
6.2.2 Tabular framework 115
6.2.2.1 Table for Cultural Objects 115
6.2.2.2 Tables for Book description 117
6.2.2.3 Recognitions and interpretations 117
6.2.2.4 Table for First Level description 118
6.2.2.5 Table for Visual Citations description 118
6.2.2.6 Table for Second Level description 118
6.2.2.7 Table for Symbols description 121
6.2.2.8 Table for Third Level Description 121
6.2.3 Description rules and database population 121
6.2.4 Database cleansing 122
6.2.4.1 Rules for normalization 122
6.3 RDF data creation 123
6.3.1 Ontological modeling 123
6.3.2 Main objects modeling 123
6.3.3 Modeling challenges 128
6.3.3.1 Artwork specifications 128
6.3.3.2 Uncertainty 132
6.3.3.3 Artwork content 133
6.3.3.4 The concept of evolution 135
6.3.4 Conversion into RDF 136
6.3.4.1 URI creation rules 136
6.3.5 External vocabulary and alignment 141
6.3.5.1 External vocabulary 142
6.3.5.2 Reconciliation 142
6.3.5.3 Hierarchy of aligned objects 144
6.4 A case study description 146
6.4.1 Information extraction and subdivision into levels 147
6.4.2 Registration in the tabular database 148
6.4.2.1 Artwork content interpretation 148
6.4.3 Cultural Object information modeling 153
6.4.4 Pre-iconographical Recognition 154

6.4.4.1 Formal Motif Recognition . . . 155
6.4.5 Level 2 modeling . . . 157
6.4.6 Level 3 modeling . . . 159

7 Analysis **161**
7.1 Quantitative overview . . . 161
7.2 Domain-specific research questions . . . 163
7.2.1 Cultural Phenomena . . . 165
7.2.1.1 SQ1. How many artworks have a cultural phenomenon associated? . . . 165
7.2.1.2 SQ2. What is the frequency of level 1 and 2 subjects occurring with each cultural phenomenon? . . . 166
7.2.1.3 SQ3. What is the frequency with which a piece of evidence supports the recognition of a cultural phenomenon? Of what type is it? (texts, artworks, etc.) . . . 171
7.2.1.4 SQ4. Is there any co-occurrence between cultural phenomena and other types of subjects? . . . 171
7.2.1.5 SQ5. What types of recognition support the recognition of a cultural phenomenon? . . . 174
7.2.1.6 SQ6. What are the artworks that have a cultural phenomenon in common and a different style associated? . . . 177
7.2.1.7 Summary . . . 179
7.2.2 Symbols . . . 180
7.2.2.1 SQ7. Which and how many symbols express the same concept? . . . 180
7.2.2.2 SQ8. What different symbolical meanings can the same symbol have? . . . 180
7.2.2.3 SQ9. How does the usage of symbols change over time and in different contexts? . . . 181
7.2.2.4 SQ10. What are all the symbols motivated by a specific source? Do they vary from the corresponding subjects not citing a piece of evidence? . . . 182
7.2.2.5 SQ11. Do deeper meanings cite as support textual sources? . . . 183
7.2.2.6 Summary . . . 183
7.2.3 Iconographies . . . 184
7.2.3.1 SQ12. What are the representational variations of iconography? I.e., what are the sets of level 1 subjects composing the recognized level 2 subject in artworks? . . . 184
7.2.3.2 SQ13. What are the attributes having a symbolic meaning? . . . 186
7.2.3.3 SQ14. How does the representation of iconography vary over time? Of which level 1 subjects is it composed? . . . 187

7.2.3.4 SQ15. Which attributes allow us to identify representations? What is their frequency? . . . 189
7.2.3.5 SQ16. Among the attributes marked as recognizing, which are the most common and rare? . . . 190
7.2.3.6 SQ17. Which are the most common level 1 subjects not marked as recognizing? . . . 191
7.2.3.7 Summary . . . 192
7.2.4 Evidence . . . 193
7.2.4.1 SQ18. What were the known textual sources to which the artwork refers? Is the artwork involved in a cultural phenomenon? . . . 193
7.2.4.2 Summary . . . 195
7.2.5 Visual Citation . . . 197
7.2.5.1 SQ19. What artworks cite the visual pattern of other ones? . . . 197
7.2.5.2 SQ20. In which cases do those artworks which are involved in a visual citation have a cultural phenomenon? . . . 198
7.2.5.3 Summary . . . 200
7.3 New inquiries into the ICON dataset . . . 201
7.3.1 Quantitative art historiography: characterizing Panofsky's approach . . . 201
7.3.1.1 The usage of textual sources . . . 201
7.3.1.2 Levels of description . . . 202
7.3.1.3 Iconological interpretations . . . 203
7.3.2 Inquiries into the iconological interpretation . . . 204

8 Evaluation **217**
8.1 Dataset evaluation . . . 217
8.1.1 Metrics . . . 218
8.1.1.1 Intrinsic category . . . 218
8.1.1.2 Contextual category . . . 220
8.1.1.3 Representational category . . . 221
8.1.1.4 Accessibility category . . . 222
8.2 Evaluation: domain-specific research questions . . . 227
8.2.1 FAIR evaluation . . . 227
8.3 Ontology evaluation . . . 228
8.3.1 Evaluation dataset . . . 228
8.3.2 Competency questions evaluation . . . 228
8.3.2.1 CQ Level 1 . . . 229
8.3.2.2 CQ Level 2 . . . 231
8.3.2.3 CQ Level 3 . . . 232
8.3.2.4 General CQs . . . 232
8.3.3 Comparison with existing ontologies . . . 234
8.3.4 Automatic Evaluation . . . 240

9 Discussion and conclusions **243**
9.1 Research objectives and contributions . . . 244
9.1.1 Ontological modeling . . . 244
9.1.2 Iconology Dataset . . . 245
9.1.3 Data Analysis and key findings . . . 245
9.1.3.1 Domain-specific Research questions . . . 245
9.1.3.2 Quantitative art historiography: characterizing Panofsky's approach . . . 246
9.1.3.3 Inquiries into the iconological method . . . 247
9.2 Impact of the research . . . 248
9.3 Limitations . . . 249
9.4 Conclusions and future work . . . 250

III IN-PROGRESS APPLICATIONS **253**

10 Tool for art information integration and retrieval **255**
10.1 Narrative-based Art Interpretation . . . 256
10.1.1 Understanding as a question-answering . . . 256
10.1.2 Case study . . . 257
10.2 Integrative Narrative Networks (INN) . . . 258
10.3 Case study for the Lotto painting . . . 259
10.3.1 Queries to knowledge sources . . . 259
10.3.2 Results . . . 260
10.3.2.1 ?Who . . . 261
10.3.2.2 ?What . . . 261
10.3.2.3 ?When . . . 263
10.3.2.4 ?Where . . . 263
10.3.2.5 ?Why . . . 263
10.3.2.6 Second iteration . . . 263
10.3.3 Utilisation of the INN by art historians . . . 264
10.3.4 Discussion . . . 264
10.4 Conclusions and future work . . . 265

Bibliography **267**

List of Figures

2.1 Michelangelo, *The Punishment of Tityus*, 1532, Charcoal drawing, a gift to Tommaso de' Cavalieri, Royal Collection Trust (Wikimedia Commons, public domain) . . . 20

3.1 The gold standard schema created by the application of CQs to the golden example from the literature . . . 46
3.2 Landscape of the knowledge graphs on their iconographical and iconological statements and the structure of the schemas that describe them. . . . 54

4.1 Approach overview . . . 60
4.2 How the research questions are related . . . 63
4.3 Timeline . . . 64
4.4 Workflow adopted for reconciliation . . . 70
4.5 Workflow of ontological modeling and data creation . . . 72

5.1 Definition of the characteristics of the domain: how the entities involved in each level of interpretation relate to each other . . . 82
5.2 ICON ontology classes and properties linked to recognitions . . . 96
5.3 ICON ontology classes and properties describing the pre-iconographical level of interpretation (level 1) . . . 98
5.4 ICON ontology classes and properties describing the iconographical level of interpretation (level 2) . . . 101
5.5 ICON ontology classes and properties describing the iconological level of interpretation (level 3) . . . 101
5.6 Alignment with other ontologies of the artwork description . . . 109

6.1 Modeling of artworks . . . 124
6.2 Modeling of people . . . 125
6.3 Modeling of books . . . 127
6.4 Modeling of books content . . . 127
6.5 Modeling of places . . . 128
6.6 Detail of (feature-of relation) . . . 130
6.7 Feature of relation . . . 130
6.8 Pendant of relation . . . 131
6.9 Modeling of specific location . . . 132
6.10 Caption for LOF . . . 146

6.11 summary of the process of the scholar's text understanding and abstraction and how this content can be described according to the three-level approach, applied to the selected case study 148
6.12 Modeling of the metadata of the Roman relief 153
6.13 Modeling of the pre-iconographical Recognitions of the Roman relief 154
6.14 Modeling of the pre-iconographical Recognitions of the Medieval relief 155
6.15 Modeling of the formal motifs Recognitions of the case studies . . . 156
6.16 Modeling of the Iconographical Recognitions of the Roman relief . . 157
6.17 Modeling of the Iconographical Recognitions of the Medieval relief . 158
6.18 Modeling of the Iconological Recognitions of the Medieval relief . . 159

7.1 a) Frequency of artworks per date; b) frequency of cultural phenomena types; c) the number of recognitions per person responsible; d) term cloud of second level-subjects 162
7.2 Number of artworks having a third-level subject, either a cultural phenomenon, a concept, or both of them 166
7.3 Results of SQ2, showing the co-occurrence of the most frequent 5 level 1 subjects, defined as the number of artworks in which they appear together. 168
7.4 Results of SQ2, showing the co-occurrence of the most frequent 5 level 2 subjects, defined as the number of artworks in which they appear together. 169
7.5 Results of SQ2, showing the percentage of the times the subject appears in artworks having the specified phenomenon type over the total of artworks in which the phenomenon is registered. 170
7.6 a) Results of SQ3, showing the type of cited evidence; b) Results of SQ4, showing the co-occurrence of subject type (support: 0.20, confidence: 0.90) 172
7.7 SQ5 results: Sankey diagram showing the frequency with which recognitions of Cultural Phenomena refer to a specific part of the artwork 175
7.8 SQ6 results: network of cultural phenomena associated with artworks having a specific style 175
7.9 a) SQ6: artworks having both a cultural phenomenon and a style, citing artworks having a style associated, b) SQ6: artworks with a cultural phenomenon and a style, grouped by common style 176
7.10 1) Symbolic meanings having more than one symbol associated (SQ7), and 2) symbols expressing more than one meaning (SQ8) . . 181
7.11 1) Symbols over time and their frequency for each date (SQ9), and 2) symbols having and not having a piece of evidence (SQ10) 182
7.12 a) Interface for having an overview of a single iconography variations (SQ12) and b) for exploring the artworks in which the variation is found 186
7.13 Overview of the iconographies with a level 1 subject which is part of a symbol (SQ13) 188
7.14 Level 1 subject variation over time of the iconography of Saturn (SQ14) 189

7.15 Most rare and frequent recognizing attributes for each iconography (SQ16) . . . 190
7.16 a) Frequency of attributes not marked as recognizing of the iconography of Saturn; b) total frequency of recognizing and non-recognizing attributes of the same iconography . . . 192
7.17 Textual sources grouped according to the artwork's author . . . 195
7.18 Textual sources grouped according to the artwork's date . . . 196
7.19 Interface for comparing the artworks involved in a visual pattern citation. In this example, Durer's *Hercules Killing the Stymphalian Birds* (on the left) visually cites Pollaiolo's *Hercules Killing Nessus* (on the right). All the artwork images are in the public domain . . . 199
7.20 Network of artworks involved in a visual citation (yellow) and their cultural phenomena (red) . . . 199
7.21 a) Percentage of interpretations citing evidence; b) percentages of recognitions citing evidence; c) correlation between recognition and evidence type; d) frequency of piece of evidence per type. . . . 202
7.22 a) Number of artworks having a description of one, two, or three levels; b) levels described in each artwork. . . . 202
7.23 a) Recognitions giving support to other ones divided per type; b) the network of iconological recognitions. . . . 203
7.24 Overview of the level of description of artworks involved in a common cultural phenomenon . . . 206
7.25 Distribution of artworks having only a cultural phenomena description (green) according to the bibliographical source (RQ4.1) . . . 207

8.1 Description of the example using ICON (level 1) . . . 236
8.2 Description of the example using ICON (level 2 and 3) . . . 237
8.3 Description of the example using VIR . . . 238
8.4 Description of the example using Wikidata . . . 238

10.1 Lorenzo Lotto, *Venus and Cupid*, 1520s, Metropolitan Museum of Art, New York . . . 258
10.2 How the network appears (a) with the 5WH questions and sub-questions still unbound, (b) after answering the ?Who question, (c) the ?What question, and (d) at the end of the information retrieval. 262
10.3 This Figure shows an interface for art historians to steer the INN expansion process. The historian can click on certain questions to expand them further and zoom in and out. . . . 265

List of Tables

2.1 Levels of interpretation according to Panofsky (1955) 20
2.2 Example of description of an artwork (Tityus, by Michelangelo) interpretation through three levels 21

3.1 Classes and properties related to the recognition of artworks (sculptures and paintings if available) in selected knowledge graphs 40
3.2 Overview of the artwork subject presence in the selected graphs (as of 01/12/2022) 42
3.3 Evaluation metrics, the first five criteria are adapted from (Färber et al. 2018), the last criterion is newly developed 44
3.4 Inter-annotator agreement scores as measured by quadratically weighted Cohen's Kappa for semantic validity of iconographical and iconological triples and column completeness per knowledge graph 45
3.5 Properties identifying iconographical and iconological content for each selected knowledge graph 47
3.6 Results for each metric over the selected Knowledge Graphs (as of 01/12/2022) 53

5.1 Definition of the entities involved in the pre-iconographical description according to Panofsky's theory (integrated with Imdahl (2012)) 79
5.2 Definition of the entities involved in the iconographical description according to Panofsky's theory 80
5.3 Definition of the entities involved in the iconological synthesis according to Panofsky's theory 81
5.4 Overview of research questions extracted from the literature and the source 90
5.5 Description of requirements 1-5 according to SEEMP methodology . 93
5.6 Description of requirements 6-7 according to SEEMP methodology . 94
5.7 Direct Reuse of Classes and Properties in ICON 105
5.8 Indirect Reuse of Classes and Properties in ICON: icon element - type - external element 106
5.9 Distinction between subjects and visual subjects depending on Panofsky's levels of interpretation 111

6.1 Field described in the table for cultural objects description 116
6.2 Fields described in the table for pre-iconographical recognitions . . 119
6.3 Fields described in the table for iconographical recognitions 120
6.4 Entities for which the type is indicated in the URI 138

6.5 Identifiers creation rules . . . 140
6.6 Identifiers creation rules (Manifested subjects and recognitions) . . 141
6.7 Cultural object relations aligned with CONA associative relations . 143
6.8 Overview of the alignments made and matching strategy 145
6.9 Cultural object Table filled with the case studies metadata 149
6.10 Expression of the recognized first-level objects in tabular format . . 151
6.11 Expression of the visual citation performed by the Medieval relief of the Roman one . . . 151
6.12 Expression of the second level subjects recognition in a tabular format 152
6.13 Expression of the third level subjects recognition in a tabular format 153

7.1 Overview of research questions extracted from the literature, their specification in sub-questions performed in SPARQL queries. 163
7.2 Overview of the subject types having a support score greater than 0.20 . . . 173
7.3 Overview of the recognitions giving support to the recognition of a cultural phenomenon . . . 174
7.4 Results of SQ11 showing the number of third-level concept recognition citing as evidence a textual source . . . 183
7.5 Example of results obtained by performing SQ12 . . . 185
7.6 Overview of the iconographies having a description at the first level and appearing more than once (SQ12) . . . 185
7.7 Frequency of recognizing attributes of Cupid (SQ15) . . . 190
7.8 Example of copying and copied artworks . . . 198
7.9 The most connected cultural phenomena in the iconological network 204
7.10 The most connected cultural phenomena in the iconological network 205
7.11 Total of artworks having at least a cultural phenomenon, distinguished for the degree of description (multiple levels or only third one) . . . 205

8.1 Performed evaluation metrics . . . 224
8.2 Percentage of addressed and partially addressed sub-questions expressing the RQs of the domain . . . 227
8.3 Test dataset overview of a number of triples, subjects, and interpretations . . . 228
8.4 Information about artworks cited in the queries . . . 229
8.5 CQ 1.2.1 results . . . 230
8.6 CQ 1.3.1 results . . . 231
8.7 CQ 2.6.1 results . . . 232
8.8 CQ 0.2.1 results . . . 233
8.9 'Time the destroyer' description according to the three levels of interpretation . . . 235
8.10 Comparison between Wikidata, VIR, and ICON . . . 240

10.1 Knowledge bases included in the case study. . . . 261

List of Abbreviations

C

CH Cultural Heritage

I

icon Iconography and Iconology

K

KB Knowledge Base

KG Knowledge Graph

L

LOD Linked Open Data

P

PREFIX a-cd: <https://w3id.org/arco/ontology/context-description/>

PREFIX a-dd: <https://w3id.org/arco/ontology/denotative-description/>

PREFIX aat: PREFIX aat: <http://vocab.getty.edu/aat/>

PREFIX arco: <https://w3id.org/arco/ontology/arco/>

PREFIX cito: <http://purl.org/spar/cito/>

PREFIX crm: <http://www.cidoc-crm.org/cidoc-crm/>

PREFIX d: <https://w3id.org/icon/data/>

PREFIX dbo: <https://dbpedia.org/ontology/>

PREFIX dbp: <http://dbpedia.org/property/>

PREFIX dc: <http://purl.org/dc/elements/1.1/>

PREFIX dcmitype: <http://purl.org/dc/dcmitype/>

PREFIX dul: <http://www.ontologydesignpatterns.org/ont/dul/DUL.owl#>

PREFIX fabio: <http://purl.org/spar/fabio/>

PREFIX gndo: <https://d-nb.info/standards/elementset/gnd#>

PREFIX icon: <https://w3id.org/icon/ontology/>

PREFIX mmm: <http://ldf.fi/schema/mmm/>

PREFIX nmo: <http://nomisma.org/ontology#>

PREFIX owl: PREFIX owl: <http://www.w3.org/2002/07/owl#>

PREFIX rdf: <http://www.w3.org/1999/02/22-rdf-syntax-ns#>

PREFIX rdfs: <http://www.w3.org/2000/01/rdf-schema#>

PREFIX schema: <https://schema.org/>

PREFIX sim: <https://w3id.org/simulation/ontology/>

PREFIX skos: <http://www.w3.org/2004/02/skos/core#>

PREFIX wd: <http://www.wikidata.org/entity/>

PREFIX wdt: <http://www.wikidata.org/prop/direct/>

CHAPTER 1

Introduction

Iconography investigates the subject matter and meanings of works of art and their evolution over time and place. To this aim, art historians study the stylistic conventions of a certain period, the literary sources that the artist may have known, and the iconographic types used to depict certain subjects or concepts. In some cases, artists propose a visual rendering based on a specific literary text or on an *iconographical program* (i.e., what a work of art should represent, usually written by a humanist in the court of the patron commissioning the painting). Iconology moves from the results obtained by an iconographical description, pushing the research forward to the understanding of deeper meanings expressed, such as the artist's or a nation's attitude or social, religious, and cultural practices.

Such characteristics make the domain complex, in which knowledge has different grades of certainty, and the personal scholars' point of view has a relevant role in knowledge construction. For these reasons, iconographical-iconological analysis is usually stored in a free text field in catalogs, which makes domain-specific information retrieval challenging. Furthermore, as it is intimately related to their authors, this knowledge is connected to the references in which such interpretations are expressed, mainly available in printed format. The traditional knowledge, well known to art historians, is nowadays stored in scholarly literature, with little possibility of online content retrieval. Interpretation includes artworks highly different for what concerns their type (e.g., prints, painting, sculpture, coins, illuminations), period, style, and place of conservation, making the possibility to compare items fragmentary and limited.

However, the increasing application of computational methods to the field of cultural heritage could change the situation. The recently born discipline of Digital Art History (DAH) highlighted the benefits and need for computational methods in art history. Yet, the application field and the limits of the discipline are still under discussion.

Among the technologies adopted in the fields of CH and DAH, semantic web technologies offer means to represent complex domains of knowledge in a structured format. The semantic web was designed by Tim Berners Lee, one of the founders of the World Wide Web, as an extension of the hypertext web (Berners-Lee et al., 2001). Its basic principle is the description and unique identification of data, which allows one to establish relations with the specific *datum*, conceiving ontologies and RDF as the core means to achieve such a structured knowledge representation. The flexibility offered by the graph data structure (corresponding to the web of data and their relations) raised the interest of many domains, including the humanities, which recognized the technology as a suitable tool for expressing the semantic complexity embedded in humanistic knowledge.

Notably, interest in exploiting the computer for tasks related to museum collections maintenance and art history was already expressed in the 1960s. Lindsay (1966) envisioned a rapid adoption in cataloging practices and the creation of collaborative collections through shared standards, but imagined a reticence of art historians to adopt the computer and make information democratically available. To date, the majority of institutions and international associations (e.g., the International Council of Museums (ICOM) and PHAROS, the International Consortium of Photographic Archives) have embraced the potential of computers in cataloging and aggregating data, providing shared standards of description (e.g., CIDOC-CRM, the guidelines *Categories for the Description of Works of Art* and *Cataloguing Cultural Objects*), aggregators (e.g., Europeana), and some museums datasets publicly available.

The use of computers for art history analytical studies was initiated by Prown, who unveiled the potential of computation for addressing domain research questions. To date, several studies have used computational methods to study artworks better and provide insights from another perspective, seeing computer-based and quantitative analyses as a support for traditional, qualitative research. Such examples are more frequent in economic and social studies of art. Among them, the work by Greenwald (2021) expressed a methodological reflection highlighting the importance that data-driven art history can have for advances in research.

According to several Digital Humanities scholars, semantic technologies can even go beyond being a powerful support tool for scholars. By performing pattern discovery and serendipitous knowledge discovery, the system could detect research problems that still need to be addressed (Hyvönen, 2020).

To the author's knowledge, little attention has been paid to the production and analysis of iconography and iconology data.

To fill this gap, the current work addresses how computational methods can be

helpful for iconographical-iconological domain-specific information representation, retrieval, and analysis. We refer to semantic web technologies as they are feasible tools for expressing the complexity embedded in iconographical and iconological interpretations described above. Furthermore, we aim to pursue the research line of quantitative art analysis proper of the field of art economics and social studies to another subdomain of art history, namely iconological studies, expressed in semantic data.

To this end, three research problems are addressed in the thesis:

- **RP1**: the formal representation of art historians' interpretations in the iconographical and iconological domain of knowledge
- **RP2**: the description of real-world iconographical and iconological interpretations in structured data
- **RP3**: the definition of approaches according to which a domain-specific analysis can be performed quantitatively

Each research problem has a corresponding research objective that this thesis aims to achieve:

- **RO1**: providing reusable means for iconography and iconology formal description
- **RO2**: providing a semantic dataset of domain-specific art interpretations
- **RO3**: addressing domain research questions in a quantitative way to prove the usefulness of the approach for iconographical-iconological inquiries

To achieve such objectives, we formulate the following main questions:

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

In the current study, we apply the research questions to the case study of the art historian Ervin Panofsky. For the relevance that his theoretical approach had in the definition of the discipline, he is an authoritative reference for its formalization. Furthermore, by providing a dataset describing his interpretations expressed according to the formalization he proposed, we have the chance to quantitatively

characterize his approach and verify to which extent the theory emerges when practically applied in real research. However, to provide a descriptive model that can represent art interpretations by other scholars, this study will consider other theoretical approaches and a domain study that includes interpretations by different art historians.

The thesis is structured as follows. Part I provides the state of the art concerning the iconographical and iconological theoretical background and its representation in current technologies. In detail, Chapter 2 provides an overview of the iconographical-iconological domain of study, providing information about the most influential scholars and theoretical approaches that have been proposed. Chapter 3 illustrates the current digital tools that can be used for subject description and a survey assessing the presence and quality of iconographical and iconological statements in the currently available Knowledge Graphs. Part II, composed of Chapters 4 to 9, describes the newly conducted research. While Chapter 4 provides an overview of the adopted approach, Chapter 5 describes the domain study and the development of a new purpose-specific ontology. The creation of the Panofsky interpretation dataset is treated in Chapter 6. The quantitative analysis is conducted in Chapter 7, treating both traditional questions quantitatively addressed and the new inquiries that the novel approach allows. Chapter 8 is dedicated to the evaluation of the dataset and the ontology. The results of the current work are discussed in Chapter 9.

A final part is dedicated to in-progress applications. Specifically, it addresses how art historians and general users can effectively retrieve semantic data about specific artworks. We propose an initial study for an interface for information retrieval in which users can build their interpretation of a work of art by retrieving and integrating the knowledge currently available in Knowledge Graphs.

Part of the work presented in this thesis has already been published. In detail, the survey of the extent to which iconographical and iconological statements are represented in current knowledge graphs, presented in Section 3.6, was published in Baroncini, Sartini, et al. (2023). The proposed ontology is presented in two papers, concerning its development (Sartini et al., 2023) and its extension (Sartini & Baroncini, 2023). A presentation of the dataset is proposed in Baroncini, Daquino, et al. (2023b), while part of the results of the data analysis discussed in the chapters 7 and 8 are published in Baroncini, Daquino, et al. (2023a). Finally, the in-progress application presented in chapter 10 refers to Baroncini, Steels, et al. (2023).

Part I

BACKGROUND AND STATE OF THE ART

CHAPTER 2

The study of artworks content and meaning

Although the history of art concerns the study of a common object, i.e., the work of art, it is divided into branches of study focusing on specific aspects. This diversity is reflected in the books that aim to provide an overview of the discipline and define the range of interests (Adams, 2010; Houston, 2013; Preziosi, 2009).

Since an artwork, due to its nature, cannot usually be fully comprehended only from its objective characteristics, most disciplines aiming at understanding a work of art include a grade of the interpreter's subjectivity. In the context of art comprehension, an interpretation is intended as 'any kind of assignment of meaning or significance to artworks' (Robert, 1996, p. 113).

Through interpretations, art historians can claim different types of explanations about art concerning different aspects, e.g., the content of the art, the trend of the art of a century, or the reception of art by the public (Houston, 2013). Among the possible art-critic approaches, some cover different aspects of the artwork. As shown in the detailed overview of art critics provided by Preziosi (2009), scholars who adhere to the branch of the study of iconography and iconology are described in various sections, including style, anthropology, and meaning. This reflects the interdisciplinary nature of the approach, aiming at identifying not only the subject matter of a work of art (iconography) but also a more comprehensive cultural understanding of images. It highlights how images can be read as historical documents of the context in which they were created, including facts and meanings ranging from the artist's personality to sociocultural phenomena. Such analysis includes multiple aspects of the work of art, examining formal and stylistic aspects and its content, and considering external sources, such as archival documents and literary texts.

2.1 What is an interpretation act

To the author's knowledge, most art criticism agrees with the interpretative nature of the discipline itself since works of art are understood as 'rarely obvious' (Barrett, 2003, p. 199). Furthermore, contemporary criticism is slightly influenced by the postmodern approach, according to which oeuvres can potentially have infinite meanings, making 'the concept of interpretation one of the key concepts of our century' (Stecker, 1994). Some approaches even consider the interpretation as part of the identity of the artwork, as 'a work of art cannot be aesthetically evaluated without an interpretation' (Rantala, 1991, p. 273).

Although the relevance of interpretation for art understanding is assessed, there is no unique agreement on the extent to which art analysis is of a descriptive or interpretative nature. Houston (2013) identifies description, interpretation, and judgment as the three main interpretation activities. The description of an artwork constitutes the first action made while interpreting it, including concrete visual traits, such as the colors, shapes, and everyday objects depicted. Nevertheless, even simple descriptions involve subjective choices (Houston, 2013) that depend on the background of an observer and his familiarity with the representational conventions used to represent the objects (Panofsky, 1955). Art critics define what should be meant as description and interpretation (Houston, 2013, pp. 113-115). Among the reported perspectives, we account for the position of Matthews (1977), who underlines that, while the description is verifiable, the interpretation is not. This perspective aligns with Gombrich's take, who points out that even decorative elements may convey deeper meanings (Gombrich, 1948).

Furthermore, different types of criticism can be conducted (Wollheim & Eldridge, 2015, pp. 124-136). When the art critic seeks to reconstruct the native meaning of the work, he/she conducts an *archeological* research, an act of retrieval of predetermined meanings. On the contrary, other critics see their act of viewing the work as an act of revision. According to them, 'critical objectivity is impossible, as every viewer is bound to perceive works of art from a different position' (Houston, 2013, pp. 116-18). Stecker (1994, p. 205) explores such perspectives on interpretation, concluding that as they are 'two different ways in which the desire for understanding can be satisfied', they deserve different names to be distinguished.

According to Panofsky, the art historian aiming at interpreting a work of art has to perform an intuitive, esthetical re-creation of the artist's *intentio* condensed into the work of art. As the humanist deals with human intentions and creations, to understand such objects, he has to 'engage in a mental process of a synthetic and subjective character: he has mentally to re-enact the actions and to re-create the creations', as their 'meaning can only be apprehended by reproducing and, thereby,

quite literally, realizing, the thoughts that are expressed in the books and the artistic conceptions that manifest themselves' in the [artworks]' (Panofsky, 1955, p. 14). In these terms, every interpretation of art is subjective. Nevertheless, suppose that an art historian conducts such an interpretation. In that case, he/she will, at the same time, conduct an *archeological research* to reconstruct the circumstances under which the work of art was created, such as the genre, artistic practices, subject, cultural aspects of the period, etc. This background knowledge will allow the art historian to re-create the artwork's original meaning as intended for the artist and contemporary viewers, allowing him/her to have a more objective view of the artwork itself. In this process, synthetic intuition and archeological research are a recursive process in which the former's results influence the latter's results and vice versa.

Art critics also seek to establish 'what makes a good, or valid, interpretation' (Houston, 2013, p. 118). Since retrievalists aim to retrieve the meaning embedded in a work of art, interpretations can be correct or incorrect. On the other hand, revisionists, welcoming the fact that alternative interpretations may coexist, base the interpretation evaluation on criteria. A convincing interpretation is based on a 'reasonable description of work of art', coherent in itself, and inclusive of all features of an artwork (Houston, 2013, p. 118).

As stated above, defining when a description becomes an interpretation is challenging. Since it is not within the scope of the current work to solve this hermeneutic issue, we will adopt the definition provided by Stecker, according to which an interpretation is intended as 'any kind of assignment of meaning or significance to artworks' (Robert, 1996, p. 113), consequently including in it the more objective descriptions of the visual representation examined. Moreover, we further specify the meaning of an interpretation in the context of this work by referring to the definition provided by Van Ruymbeke et al. (2017) relative to the archeological field, stating that the views produced by scholars are 'the result of an interpretive reasoning that includes the subjectivity of the author', due to the uncertainty and incompleteness that often characterize archeological data (Van Ruymbeke et al., 2017).

This study will focus on a branch of study, namely iconography and iconology, aiming to reconstruct the original meaning of the artwork in its original context. Therefore, the perspective involved is closer to the so-called retrievalist position. Nevertheless, as the scholars' claims made in the literature point out, such archeological interpretations are not strictly correct or wrong, as it is likely that we do not have available all the pieces of evidence needed to state that the interpretation is certain. Consequently, an interpretation is usually a circumstantial

investigation (Warburg, 1999) often based on literary and archival research to reconstruct the context in which the work was created and the meaning that the artist or the patron wanted to express in the work of art. Criteria such as internal coherence of the interpretation, correspondence between visual images and texts cited as evidence (Ginzburg, 1986, p. 69), the knowledge of style, iconographies, and contemporary sociocultural phenomena (Panofsky, 1955) are consequently fundamental to determine whether an interpretation is acceptable. Sometimes, the lack of evidence does not allow us to decide which interpretation is correct, making multiple acceptable interpretations possible. An example is the different coexisting interpretations argued by Warburg and Gombrich about Botticelli's *Primavera* (Gombrich, 1972; Warburg, 1999).

2.2 Iconography and Iconology

As introduced, many disciplines study artworks from different perspectives. Among them, semiotics and iconography are dedicated to meaning, specifically how meaning is related to specific signifiers, but they have different scopes and origins. Whereas iconography regards precisely the recognition and identification of subjects and meanings represented in artworks and was born in Western art history tradition, semiotic art is one of the application fields of the discipline, as it is interested in exploring any relation between signs and meaning, having its origin in the language studies.

These two approaches ask the same two fundamental questions, namely 1) the question of representation (what do images represent and how?) and 2) the question of the hidden meanings of images (what ideas and values do the people, places, and things represented in images stand for? (van Leeuwen, 2004, p. 4)). However, they diverge in the traditional field of application and outcomes. While the semiotics of art are generally applied to contemporary images, iconology is usually applied to artworks, despite many semioticians working to extend the discipline to include them (Lancioni, 2019). Furthermore, whereas the main interest in semiotics is understanding the mechanisms that create meaning, iconography focuses on recognizing the given meaning. For this reason, since we are interested in studying works of art, we refer to iconographic study and its extension into iconology.

In the past, the two terms *iconography iconology* were used indeterminately (e.g., Ripa used the term *iconologia* as a title for a collection of allegories and personifications). However, more recent descriptions define and discern their meaning, starting from the definition provided by Panofsky (1939). The aim of *iconography* is to identify the subject matter, i.e., what an artwork depicts and represents,

through attributes. It also includes the analysis of how iconographic subjects change over time and the meaning they acquire (e.g., identifying the representation of Time as winged and with a sickle in Renaissance and depicted as Kairos in classical iconography (Panofsky, 1972)). Iconology, instead, focuses on interpreting iconographic subjects as documents of sociocultural phenomena and, therefore, explains the reasons for iconographical changes. Whereas iconography is descriptive, iconology is interpretative and highly based on the interpreter's subjective intuition, as it concerns nonobjective and less verifiable cultural meanings. For this interest, iconology is an interdisciplinary approach that includes anthropology and contextual studies.

Currently, iconography and iconology are associated with the studies conducted by Warburg and the scholars following his school. In the following sections, we provide a brief historical and methodological overview of this approach.

2.2.1 Aby Warburg and his influence on following scholars

In the 20th century European art scene, Warburg introduced a new perspective on art studies that focused on the anthropology of art, in addition to the traditional formal, historical, and stylistic approaches (Gombrich, 1972). Specifically, he delved into the endurance of classical antiquity in Renaissance society. Warburg was born in Hamburg in 1866 and received his education at the University of Bonn, where he not only studied art but also attended lectures on philosophy and psychoanthropology. These experiences sparked his fascination with the origins and changes of civilization (Cieri Via, 1994, pp. 28-29). The philologist Usener was also influential in Warburg's research, as he traced ancient mythology's evolution through etymology (Kany, 1985). In 1891, he discussed his dissertation on Botticelli, which revealed his interest in developing a method of reading images as social memory (Rossi Pinelli, 2019, pp. 272-273). His investigation differed from the prevailing aestheticizing currents dominating the art studies of the period, as each image type focused on the interactions and exchanges between cultures and periods. He adopted an anthropological approach to culture, encompassing art, literature, philosophy, science, craft activities, and superstitions (Ginzburg, 1986, p. 35). His primary goal was to comprehend the life of a particular culture as manifested in the works of art that it produced, achieved by consulting diverse types of evidence and interpreting the artwork as a historical document of such phenomena. This multidisciplinary approach (*Kunstwissenshaft*) addressed themes primarily related to astrology, the Renaissance's classicism (*Nachleben der Antike*), and the reappearance of emotionally charged formal motifs in different eras (*Pathosformeln*). In an international conference held in Ferrara in 1912, during which he examined

the astrological aspects of *Salone dei mesi* in Palazzo Schifanoia (Ferrara, Italy), he used for the first time the term *iconology* to refer to this type of image inquiry (Warburg, 1999).

Located in Hamburg, Warburg's library is a testament to the scholar's broad range of interests, encompassing volumes beyond the conventional scope of art history. Even texts on subjects such as social studies, anthropology, astrology, or other areas of research that may be considered peripheral held significant value for his studies and were frequently consulted by the scholars who frequented the library (Rossi Pinelli, 2019, pp. 272-273). The books were organized in an uninterrupted space, with no rigid distinctions between fields, reflecting the thematic continuity that defined Warburg's approach. The library evolved into a full-fledged institute in 1921 and relocated to London in 1934 to avoid the Nazi regime. Eventually, the institute became part of the university in 1944. The migration of scholars to the United States and Britain surprisingly sparked even greater interest in scholarship. As Panofsky noted in his work *The History of Art* (1953), the American cultural climate was conducive to cultivating a more scientific approach to the discipline.

2.2.1.1 Aby Warburg's method

Warburg did not explicitly define his method to catch the deeper meanings of artworks. The absence of an English translation before 1999 of the volume *The Renewal of Pagan Antiquity*[1] its work is not easily accessible (Rampley, 2001). Furthermore, according to some scholars, Gombrich's biography of Warburg (Gombrich, 1970) reinforced the idea of his various fragmentary research lines, shadowing the presence of an implicitly used method (Kany, 1985, p. 1267).

Ginzburg (1986) attempts to reconstruct a 'Warburghian method' from Warburg's quests and the results of scholars in his circle. Whereas Warburg mainly focuses on the role and meaning that classical antiquity had for the Renaissance, this is just a theme expressing the core question he addresses, namely the function of images in the life of a civilization (Ginzburg, 1986, p. 34). To this aim, he recurs to various documentary sources, from wills and letters to tapestries, to reconstruct the relations between sociocultural aspects and images. He uses figurative evidence, considered historical sources[2] (Ginzburg, 1986, p. 30), and documentary sources to understand a situation (Ginzburg, 1986, p. 44). In this sense, his history of art is located in cultural studies (*Kulturwissenschaft*), as no purely esthetical considerations are included (Ginzburg, 1986, p. 43). Style and form are analyzed in

[1] The volume was published in Italian in 1966, much before the English version. This allowed Italian scholars to access the material, such as Ginzburg (1986), who afforded the methodological question of Warburg's method a few years earlier than Rampley.

[2] 'utilizzazione delle testimonianze figurate come fonti storiche'

relation to content to gain a deeper understanding of the artwork. *Pathosformeln*, i.e., the 'forms of pathos' that Renaissance artists used as direct borrowings of emotional expressions from the Antiques, are an example of this approach. Such forms were, as Gertrud Bing reports, 'evidence of moods and feelings become images' (Thomson, 2014). Indeed, stylistic and formal analyses, on the one hand, and iconographical ones, on the other, are fundamental to understanding the powerful influence such images had on Renaissance artists. For Warburg, mere iconographical analysis was insufficient, as iconography was only one of the possible approaches to fully understanding 'life' in pictures (Ginzburg, 1986, p. 52).

Similarly to Ginzburg, other scholars tried to clarify the underlying methodology that Warburg applied.

Kany (1985, p. 1274) defines his iconology as 'the method of synchronic and diachronic comparison between the formular details'[3]. So, he points out that Warburg's method is based on details that emerge from a philological, comparative approach. This aspect is exemplified by Warburg´s studies of the interests of the Florentine society for antiquity, in which he reconstructs that they were interested in expressions of pathos as manifested by the movement of accessory elements, such as hair and clothes. This detail can be noticed in Botticelli's *Birth of Venus*, representing a theme treated in Homer. Nevertheless, the Greek author did not mention such traits. Indeed, they appear in the description of the contemporary writer Poliziano, who referred to the same passage by Homer. In this example, Warburg does philological research, not only finding the textual reference of the painting, but also measuring the difference between the painting, the contemporary source, and the classical prototype. Such details are not full of symbolism, which the author is fully aware of. The origin of such attention to detail is found in perspective, according to which the greatest is contained in the smallest, belonging to Neoplatonic thought and mysticism (Kany, 1985, p. 1278).

Furthermore, Warburg bases his analysis on a 'synchronical and diachronical' comparative method. In other words, he analyzed historical processes in their progressive evolution over time and between different cultures and places. The second approach included juxtaposing opposite poles and analyzing the continuum among them (Kany, 1985, p. 1273). For example, he traces the process of continuous exchanges between the poles of antiquity and contemporaneity in the Renaissance period. This perspective seems to align with the interest in polarities in the Renaissance highlighted by Rampley (2001, p. 309). In this way, he discovers hidden cultural traditions and traces them back to their sources in a perspective in which the past has never fully ended. According to Kany, such a comparative

[3]'il metodo di confronto sincronico e diacronico fra dettagli formulari'

method is needed because of the symbol's meaning in Warburg's view. Moving from the positions of the philologist Usener, who traced the origin of mythological figures in the repetition of an event that becomes an abstract concept, Warburg sees symbols as 'formulae nourished by the memory of an anguished primordial experience, and at the same time charged with history, endowed with a classificatory function' [4](Kany, 1985, p. 1277).

From such criticism, it emerges that, although Warburg did not explicitly state his approach, some recurrent traits of his research method emerge from his practice.

2.2.1.2 Scholars of the Warburg's school

Scholars at Warburg's Institute begin their research from a shared starting point but interpret it in personal ways. They all analyze images as carriers of meaning and draw from contemporary culture, including literature, philosophy, dominant iconography, and symbology, as evidenced by their work (Rossi Pinelli, 2019, pp. 374-75). Due to the diverse research lines and multidisciplinary approach taken by these scholars (Ginzburg, 1986), the question of whether a 'Warburghian method' exists is still open to debate.

Fritz Saxl (1890-1948) played a crucial role in fostering the library's cultural growth in the role of librarian from 1913 and substituting Warburg in running the institute, along with Gertrud Bing, from the '20s. Saxl continues Warburg's research line by investigating the combination of myths, astrology, and pictorial representations, using a method that interpreted variants of images as manifestations of the histories of ideas (Rossi Pinelli, 2019, pp. 373-77). As he received his training at the Vienna school after Dvořák, Wölfflin, and Schlosser, he reemployed the methods of inquiry he learned to relate art history to other branches of history, namely politics, literature, religion, and philosophy (Cieri Via, 1994, p. 52). His research studied figurative evidence as historical sources based on stylistic analysis (Ginzburg, 1986, pp. 45-51). Ginzburg (1986) reports that Saxl is considered the scholar who follows Warburg's method more consistently.

Saxl collaborated with **Erwin Panofsky**, who joined the University of Hamburg in 1920, on several occasions. Their main contribution was the essay *Classical Mythology in Medieval Art*, in which they extended Warburg's research on the survival of antiquity to the still unexplored period of the Middle Ages. According to Garin, Warburg's method that emerged in the collaboration of the two scholars is characterized by concreteness and philological precision, adherence to things, interdisciplinary approach, and rupture of academic or traditional separations (Ginzburg, 1986, p. 39).

[4]'formula nutrita del ricordo di un'angosciosa esperienza primordiale, e insieme carica di storia, dotata di funzione classificatoria'

Fundamental was Panofsky's contribution to defining a formalization of art history and a method of iconological investigation in which the relationships between the form and content of images were established. Hence, he formulated the theory of subject matter and meaning interpretation divided into three levels, which will be discussed in more detail later.

His theory was profoundly influenced by the research of **Ernst Cassirer** on symbolic forms (Panofsky, 1955). While Cassirer was a professor at the University of Hamburg from 1919 to 1933, he got in touch with Warburg and the scholars of his library, with whom he had a proficuous exchange and material for his theory of symbolic forms. Although his philosophical approach differed from the historical one followed by the institute, he recognized the construction of a historical basis for a 'science of the spirit' as a core principle of institute research (Cieri Via, 1994, p. 61). As the Warburg school, Cassirer was interested in studies on the Renaissance and myth (Cieri Via, 1994, pp. 59-64).

Rudolf Wittkower (1901-1971), who was at the Institute between 1934 and 1956, belongs to the second generation of scholars of the Institute, in which he was the co-founder and editor of the Institute´s journal. In his research activity, he was concerned with symbols and their migration between cultures in interdisciplinary research that intertwined the figurative and ethnoanthropological fields. He contributed to the formalization of the discipline through his presence in London and the United States.

Edgar Wind (1900-1971) was the first student of Erwin Panofsky, whose doctoral thesis was also revised by Ernst Cassirer. Cassirer initially inspired Wind's primary approach to methodology through Panofsky, yet his close friendship with Aby M. Warburg brought him closer to Warburg's 'cultural history' approach. Ginzburg (1986, p. 36) underlines that he attempted to include the iconological method in a philosophy of culture deeply influenced by Cassirer.

Gombrich characterized a second phase of the institute's studies, joining it in the 1930s, when he was already in London, and becoming its director from 1959 to 1976. Gombrich studied in Wien with Julius von Schlosser and addressed 'problems of the vision, reception, and representation of works of art' (Rossi Pinelli, 2019, p. 423). His frequentation of Freud and Ernst Kris deepened his psychological interests. In the art study, he applied the interdisciplinary approach of *kunstwissenshaft* that he learned in his training in Wien. According to him, the interpretation of the meaning of a work of art must start with the largest number of documents possible, possibly questioning the initial postulates to avoid coming to erroneous conclusions based on corrupted texts (Cieri Via, 1994, p. 172). The uncertainty of the results that can be obtained from the Warburgian and Panofskian approaches, based on irrational parallelisms among cultures and in which the style is intuitively

interpreted as the expression of the culture of the time, in an anti-historical perspective, led the scholar to focus on problems of style, turning increasingly toward a psychoanalytic approach (Ginzburg, 1986, pp. 58-69).

Among art historians who were influenced by the iconographic and iconological approach, mention should be made of the Dutch **Henri van der Waal** (1910-1972), the founder of Iconclass, a classification system of art subjects (Sciolla, 1997; Waal, 1952)[5]. Influenced by the theories of Ernst Cassirer, Aby Warburg, and Panofsky (Irving Lavin, 1992, p. 39), he was interested in studying images as cultural documents to highlight the connections between art and society. According to him, the study of iconography was the most appropriate way to study these relations (Sciolla, 1997, p. 521). In his view, the term *iconology* must be read in the broad sense of 'theory of images', intended as an investigation that, by equally examining the study of the forms, contents, and functions of artworks, seeks to understand the 'function that the image possesses in a certain society' (Van Den Heuvel, 2024).

In the 1960s, Van Der Waal worked on a second classification system, called *Beeldleer*, less known than Iconclass as it was never published nor finished (Nauta, 1993). The system provides a formalization of art historical topics and of the connections among them, to the extent that the 'entire system might be conceived of as the crystallization of a vision of the connection between (visual)art historical matters' (Nauta, 1993, p. 38). In detail, it provides an interlinked structure of terms, organized into eight categories, thought to be used as a tool to guide 'iconological exploration, both in teaching and in historical art research, to methodologically shape art historical studies (Van Den Heuvel, 2024).

2.2.2 Erwin Panofsky and the theorization of interpretation act

Among the art historians following Warburg's research line, Erwin Panofsky had a relevant role in his attempt to furnish a theorization of Art History as a scientific discipline and in providing a first definition of the interpretation act while interpreting its content and subject matter (Panofsky, 1955, 1972)[6].

Although some critics of his approach have been raised, it is widely recognized that his formalization was crucial for the discipline (Liepe, 2019; Müller, 2014; Rampley, 2001). Aby Warburg's works provide a basis for understanding Panofsky's iconology, as both Panofsky and Warburg illustrate the interconnections among cultures. Nevertheless, Warburg never explicitly expressed his method as a system.

[5]The Iconclass system will be more thoroughly described in Chapter 3.

[6]The current section is based on section 2 of Baroncini, Sartini, et al. (2023), for which S.Baroncini is responsible.

For this reason, Panofsky 'has most often been regarded as completing much of Warburg's work, endowing it with greater philosophical rigor' (Rampley, 2001, p. 305).

According to him, different types of meaning are recognized when interpreting a work of art subdivided into three layers. The depth to which the artwork can be understood depends on the background knowledge of the observer: the more he knows the artist, stylistic conventions, and cultural context of him/her, the more the interpretation at each level is correct, including more profound insights on cultural meanings. Although different types of meaning are recognized at each level, the overall process of interpretation, as it concerns the artwork in its unity, is organic and indivisible.

The first layer, the pre-iconographical description, deals with recognizing factual (e.g., objects, people, actions) and emotional meanings, namely primary or natural subjects. In detail, this description is achieved by recognizing pure forms (i.e., combinations of forms and colors) as carriers of primary subjects. The pure forms recognized are called 'artistic motifs', and their combinations are called 'compositions'. An enumeration of the recognition of artistic motifs constitutes a pre-iconographical interpretation of the artwork (Panofsky, 1955, p. 28). To recognize such subjects, the practical experience of the observer is usually sufficient. Nevertheless, some cases require the knowledge of the stylistic conventions adopted to distinguish, for example, if a person depicted in mid-air is an apparition, as it would be in a Renaissance painting, a person from afar (e.g., Japanese context), or just a depicted person in a medieval representation.

Suppose that the observer is familiar with the literary sources that the artist knows. In that case, the subjects already identified at level 1, viz. the artistic motifs or compositions, can be recognized at the second level by the combination of them with concepts and themes, obtaining, for example, characters (e.g., Venus), personifications (e.g., Virtue), or events (e.g., the Battle of Cascina). Recognized artistic motifs are called images or *Invenzioni*, the term used by ancient theorists to identify stories and allegories. Allegories are defined in opposition to stories as 'combinations of personifications and/or symbols', although there are many intermediate possibilities between them (Panofsky, 1955, p. 29, note 1).

Finally, by knowing and understanding the cultural and social aspects of the artist's time, it is possible to read the artwork and subjects identified at previous levels as symptoms of contemporary society, of the artist's beliefs and personality, or as expressions of voluntarily inserted meanings.

Correct identification of previous levels determines the subjects recognized in the deeper strata: 'a correct iconographical analysis presupposes a correct identification of the motifs' [p. 30].

The scholar emphasizes that the first two levels describe facts and are, respectively, in the domain of forms and iconography. On the contrary, the last level is in the domain of iconology, which is a synthetic intuition rather than a description.

Table 2.1 resumes the synoptic table in Panofsky, 1955, pp. 40-41, integrating it with a further explanation of concepts.

Table 2.1: Levels of interpretation according to Panofsky (1955)

Level	**Type**	**Subject identified**	**Recognized elements**	**Necessary background**
1	Pre-iconographical description	Natural or primary subject, namely factual and expressional meaning	Artistic motifs and their combinations (compositions): pure forms recognized as carriers of primary meanings	practical experience
2	Iconographical description	Secondary or conventional subject	Images and their combinations (invenzioni, i.e., stories and allegories): artistic motifs recognized as carriers of a secondary meaning	literary sources describing themes and concepts familiar to the artist
3	Iconological interpretation	Intrinsic meaning or content	Symbolic values: artistic motifs, images, stories, and allegories are recognized as manifestations of underlying principles of a cultural context	familiarity with cultural phenomena, tendencies, attitudes

To illustrate Panofsky's theory and present an example in which each level of interpretation is covered, we describe Michelangelo's Tityus interpreted in (Panofsky, 1972, p. 216).

Figure 2.1: Michelangelo, *The Punishment of Tityus*, 1532, Charcoal drawing, a gift to Tommaso de' Cavalieri, Royal Collection Trust (Wikimedia Commons, public domain)

The drawing (Figure 2.1) shows a naked, laying man attacked by a vulture (level 1). It represents the story of Tityus (level 2), punished by Apollo for having

assaulted his mother Leto by chaining him to a rock in Hades while two vultures eternally devour his liver, considered the seat of physical passions (symbol, level 2). The story had been commonly interpreted by Michelangelo's contemporaries as an allegory of the tortures caused by immoderate love (allegory, level 2). On this basis, Panofsky claims that the artist depicted this story as a symbol of his personal passion for Tommaso Cavalieri (level 3), to whom he gifted a corpus of drawings pervaded by Neoplatonic meanings (level 3). Table 2.2 shows how this interpretation can be subdivided into levels. For its completeness, this drawing will be considered as an example for the evaluation of the artworks in the currently available knowledge graphs in Section 3.6.3.

Table 2.2: Example of description of an artwork (Tityus, by Michelangelo) interpretation through three levels

Level	**Description**
1	Nude, laying man, attacked by a vulture
2	Tityus; story of Tityius, whose liver is devoured by a vulture; liver as the seat of physical passions; story of Tityus as an allegory of the tortures caused by immoderate love.
3	Agonies of sensual passion, enslaving the soul and debasing it even beneath its normal terrestrial state according to the Neoplatonic theory; Expression of the agonies of sensual Passion that pervaded Michelangelo after he had met Tommaso Cavalieri, for whom he realized the drawing.

2.2.3 Other theoretical approaches proposed formalizing the interpretation act and criticism towards Panofsky

Since scholars declined the iconological method in different, personal ways, alternative theories of the interpretation act were proposed. Some clearly state the limits or pitfalls of Panofsky's one, proposing extensions or new formulations (e.g., Van Straten, Imdahl, Taylor). In this section, we propose a brief overview of such theories.

Wittkower

In his essay titled 'Interpretation of Visual Symbols,' Rudolf Wittkower delves into the topic of how a work of art, with its visual symbolism derived from the artist's personal perception, can be comprehended by an observer attempting

to reconstruct its depth of meaning through interpretation(Wittkower, 1977, pp. 173-216). Wittkower distinguishes between a 'rational interpretation,' consisting of four levels of meaning, and an 'emotional interpretation.' The first level of rational interpretation, the 'literal representational' level, involves recognizing the literal meaning through object experience and knowledge of conventional representation codes. The second level pertains to the iconographic theme ('literal thematic' meaning), requiring familiarity with the artist's conventional code for attributes, allegory combinations, actions by mythological figures, and information from written or oral sources. An observer lacking awareness of these codes may fail to comprehend the work's subject matter (Wittkower, 1977, p. 176). However, more than recognizing the theme is required to fully understand the picture, as it may have an intentional meaning. A third-level meaning, namely, the 'multiple meaning', may be expressed. For example, in Raphael's fresco depicting Pope Leone I and Attila in the Vatican Rooms, a parallelism between the historical event and contemporaneity is voluntarily established, as the pope bears the portrait of Pope Leone X. This evokes the meaning of the power of the Church through centuries by recalling two military successes in which the two popes were involved (the retreat of the Huns in 452 and the win of the Church during the Novara battle in 1513).

Finally, the last level investigates the recognition of 'expressive meaning', namely the traces of the artist's personality. Differently from Panofsky, Wittkower's last level of interpretation focuses on understanding the extent to which an observer can grasp the message that the artist wanted to give to the work rather than an understanding of the society that influenced its production and modes of expression. The more the interpreter reaches a deeper level, the less objectively the interpretation results can be evaluated.

Besides the rational interpretation, an emotional one occurs, which is intuitive understanding depending on theoretical beliefs (e.g., aesthetics).

Van Straten

van Straten (2012) makes Panofsky's model more detailed by distinguishing the presence of intentional meaning expressed by the artist from unintentional meaning that sheds light on the context of production. Accordingly, his iconographic and iconological interpretation model is divided into four levels. He adds an intermediate level between Panofsky's iconographical description and iconological interpretation for identifying intentional or symbolic meanings. The last level concerns the uniquely iconological interpretation of the unconscious meanings projected by the cultural context into the artwork.

Imdahl

Although Imdahl (2012) generally agrees with the levels proposed by Panofsky, he criticizes that the relevant iconic language of visual images is ignored in the proposed theorization. For this reason, he proposes adding a parallel interpretation level concerning the iconic sense given by formal aspects of the representation.

Taylor

From Taylor's perspective, an *icon* interpretation highly relates to textual sources, making it impossible to interpret this in cases where such evidence expressing a society's culture and cultural meanings is unavailable. Hence, his criticism of Panofsky does not concern his consistent reliance on textual sources. Rather, he states that Panofsky's three levels of meaning are insufficient for describing the interpretation of an iconography (Taylor, 2008, pp. 1-14). Instead, he proposes ten levels, answering respectively to the following questions:

1. What does an image resemble?
2. What kind of things does it depict?
3. What specific thing or person, if any, does it depict?
4. What kind of things does it represent?
5. What specific thing or person, if any, does it represent?
6. Which iconographic type or scene does it represent?
7. Which text, if any, does it illustrate?
8. What, if anything, does it symbolize?
9. What was its apparent purpose?
10. What is its actual purpose?

A key aspect is distinguishing between the terms *resemblance, depiction, and representation*. He illustrates why all these levels of understanding are needed by interpreting a complex case, namely the iconography of Saint Kümmernis and its frequent overlapping with the iconography of *volto santo*. He interprets Hans Springinklee's woodcut showing *The fiddler and the golden shoe* (1513). In it, the crucified figure resembles a man (level 1). It is not possible to establish what kind of things it depicts (in this case, a statue; level 2) until we identify the specific object depicted, namely the statue of Christ known as *Volto Santo* located in Lucca Cathedral (level 3). Nevertheless, the woodcut under examination intends to

represent a woman (level 4), specifically Saint Kümmernis (level 5). According to the legend, a beard grew on the face of the beautiful woman to preserve her faith in Christ after her father insisted that she marry a prince. Then, the father had her crucified. As Saint Kümmernis shares attributes of Christ (male appearance and the cross), the two iconographies overlapped, and the legend of Saint Kümmernis started to be represented with the iconography of Volto Santo in Northern Europe. The scene depicted is a story (level 6) written in the saints' legend (level 7), telling that the saint gave a golden shoe to a fiddler praying for her. Even though it is highly probable that such an image had a symbolic meaning, we do not have enough elements to hypothesize what it was (level 8). The picture's apparent purpose may be devotion (level 9), but the actual purpose may be satyrical (level 10).

According to Taylor, he formulated the levels as feeding back one another rather than seeing a linear process, as shown by levels 2 and 3, in which identifying the object depicted (level 3) led to identifying the type (level 2).

Gombrich

Gombrich does not accept the validity of different levels of meaning, as it simplifies the problem and is highly problematic. For example, at the first level of interpretation, it is difficult to clearly define what is considered a subject and what is purely decorative. On the contrary, he analyzes the understanding of the meaning of a work by investigating the criteria leading to a correct interpretation and what and how many meanings the same work can express (Gombrich, 1972). According to him, iconology is the discipline aimed at reconstructing a lost iconographical program of each work of art that defines both the depicted themes and their intended meanings. He affirms that a correct interpretation is made step by step, and the correctness of each step should be verified before proceeding. Applying Hirsch's theory of literary text interpretation to visual art, he states that the first step consists of recognizing the genre of the painting, which will also determine the themes and whether more profound interpretations may exist. The scholar indicates how proper iconological interpretation must be based on the study of 'institutions,' in other words, the context and function of the object, rather than on the 'study of symbols,' and stresses how it is necessary to return to the study of documents following any iconological interpretation, so as not to risk basing 'another even bolder hypothesis' on one's conjectures (Gombrich, 1978, pp. 32-33).

Although it is a central reference for formalizing the discipline, some alternative theories described above show that criticism against Panofsky's theory was raised. Gombrich refuses to divide the interpretation into levels, judging it as simplistic.

Some other scholars underline that defining a 'natural meaning' is problematic. Furthermore, Imdahl proposes adding an iconic level to the model since, according to him, Panofsky's approach focuses on textual sources without giving the proper relevance to the iconic language of artworks. Taylor accepts the subdivision into levels, but he proposes a more thorough framework of ten levels, judging Panofsky's one insufficient to represent the complexity of meaning unveiled by an iconological interpretation. However, he underlines the centrality of documentary sources for conducting an iconographic and iconological analysis. Similarly, Gombrich agrees with a more faithful approach to documentary sources, cautioning about the possible problems raised by interpretations based on conjectures.

Several scholars moved critics opposite to the previous ones, claiming that Warburg's iconological approach was narrowed down to a study of the transformation of artistic motifs through their interaction with texts. One criticism is that, contrary to iconology, which focuses on the unconscious culturally related symbolic meanings of forms, he tends to 'focus on the conscious artistic use of symbols and conventions' (Ginzburg, 1986, p. 57; Rampley, 2001). Finally, Moxey (1986) underlines that his work reflects some biases as he refers to the traditional genre hierarchy, preferring allegories over landscape, still life, and portraiture, and refers to the art tradition as it was established. (Moxey, 1986, p. 270)

It is true that Panofsky's theory was proposed nearly one century ago and that a clear definition of the subjects included in the analysis is complex. Nevertheless, a shareable research method or framework to support art interpretation has not yet been validated because of the high level of subjectivity involved in each claim and the various interpretations of the discipline itself that the scholars gave in their practical research. Despite its limitations, many scholars still recognize his theory as central (Liepe, 2019), providing art interpretation with the first and most robust framework. Moxey, who reported some of the critics described above, underlines that 'Panofsky's 'iconological' method still offers the discipline one of the most sensitive approaches to the understanding of the art', as 'the system of checks and balances that characterizes Panofsky's 'iconological' method has proven to be the door through which it has become possible to essay an interpretation of works of art that does justice to their complex historicity' (Moxey, 1986, pp. 271-72).

2.3 The adopted theoretical perspective

The introductory overview provided in this chapter regarding the art criticism concerning art interpretation, with a specific reference to iconography and iconology, reported on the varied panorama of approaches proposed and the single art historians' take on them. This situation challenges the definition of a common

theoretical framework. However, to the aim of the present study, we make explicit the theories and definitions that we chose as representative of the domain.

In the current study, we intend an **interpretation act** as any assignment of meaning and significance to artworks, as all the authors identifying a descriptive phase included a percentage of subjectivity even in it, also making simple descriptions of what is depicted in a work of art subjective. Furthermore, we refer to a branch of art history that sees criticism as the retrieval of the meaning of a work of art as intended by the artist or by the society in which and for which it was produced. Indeed, we are mainly interested in investigating past artworks from a historical perspective. Nevertheless, in our perspective, multiple contrasting interpretations can coexist. Whereas only one interpretation is ultimately theoretically accurate, the lack of clear evidence that often occurs (e.g., the lack of the artist's explanation or the iconographical program) makes it impossible to judge whether an interpretation is correct. Furthermore, even if pieces of evidence are present, the highly subjective and interpretative nature of iconological understanding determines the impossibility of an objective judgment. Nevertheless, strategies for defining whether an interpretation is likely correct exist, such as the citation of evidence and the consideration of contemporary stylistic and artistic conventions and ideas proposed by Panofsky.

Furthermore, for the high relevance he had and for being the broadest and most thorough theoretical systematization provided, **we refer to the three-layered model proposed by Panofsky**. His theory gains particular importance when dealing with standard description of artworks, as it is the reference for international standard description of subject matter and content[7]. Nevertheless, we consider the fundamental contribution provided by other scholars, such as Warburg's importance of visual citation and style, Imdahl's importance given to the iconic language, and van Straten's subdivision of conscious and unconscious meanings.

[7] e.g. By Getty´s Categories for the Description of Works of Art (CDWA) and Categories for Cultural Objects (CCO). This aspect is discussed in more detail in Section 3.1.

CHAPTER 3

Art description and analysis in the computational era

In recent years, increasing attention has been paid to a structured representation of art-related topics. In the context of the description of artwork subject matter and content, several notable works were conducted, providing domain-specific cataloging standard guidelines, controlled vocabularies, and ontologies. Diverse attempts in applying Computer Vision to artworks to recognize a subject show promising results, too. Furthermore, data analyses of the art market and cultural-related context hint at the kind of understanding that data may provide.

However, little data about iconography and iconology are available in the Knowledge Graphs.[1] Therefore, we conducted a study to assess their presence and correctness among the available KGs.

3.1 Cataloguing standards

To be able to share and relate artwork data on the Web, each institution must annotate its holdings according to shared standards, which can be done at different levels of detail. Cataloging standards offer a set of guidelines for entities that describe a resource. For what concerns the cataloging of works of art, an international reference point is the Getty *Catalogs for the Description of Works of Art* (CDWA)[2], partially incorporated into the guide *Cataloguing Cultural Objects*

[1] For a conceptual definition of a knowledge graph, we refer to (Fensel et al., 2020):

> Knowledge Graphs are massive semantic nets that integrate various and heterogeneous information sources to represent knowledge about specific domains of discourse.

On a technical level, we define KGs as sets of triples (subject, predicate, object) encoded in a serialization of the Resource Description Framework, or RDF (McBride, 2004)

[2] https://www.getty.edu/research/publications/electronic_publications/cdwa/18subject.html

(CCO) (Harpring et al., 2006).

The CDWA (Categories for the Description of Works of Art) is a set of guidelines created by the Getty Vocabulary Program to standardize the descriptions of works of art held in various institutions. It comprises 540 categories and subcategories of information organized around basic categories that can be used to identify a work of art. The *Subject Matter* category allows us to describe the subject and meaning of an artwork. In detail, the *Specific Subject Type* subcategory allows classifying different reading levels, as defined by Panofsky (description, identification, interpretation), to distinguish subject matter from content.

Both CDWA and CCO underline that adopting a simplified description of the approach by Panofsky 'can be helpful in indexing subjects for purposes of retrieval'[3]. Following the alignment first proposed by Shatford (1986), they define the second and third level, that is, the identification of themes, narratives, iconographies, and meanings, as the *aboutness* (i.e. what the work is about). In contrast, the first level and eventually the second corresponding to the *ofness* (viz. what can be seen by a non-expert interpreter (Žumer et al., 2012, pp. 207-208; Klenczon & Rygiel, 2014). If the subject corresponds to the work itself (e.g., the term *architecture* used for describing a cathedral) and does not refer to a subject depicted by the object (e.g., a drawing representing a cathedral), the term *isness* shall be used. The concepts of *ofness*, *aboutness*, and *isness* are a core aspect of knowledge organization initiatives (ISKO, IFLA) and are discussed in more detail in Zeng et al. (2009) and Hjorland (2016).

3.2 Controlled vocabularies

A controlled vocabulary is 'an organized arrangement of words and phrases used to index content and/or to retrieve content through browsing or searching' (Harping, 2013, p. 13). Some vocabularies describing cultural objects offer a taxonomy for describing artwork iconography.

Getty Vocabularies

The Getty Research Institute released five vocabularies,[4] specifically for the description of artworks (Cultural Objects Name Authority), iconography (Iconography Authority), general terms for describing artworks (Art and Architecture Thesaurus), artists (Union List Of Artists Names), and places (Thesaurus of Geographical Names). Each item described in the vocabulary has a unique identifying

[3]For the alignment of the concept of the subject matter to the main cataloging standards we refer to section 16 of *Metadata Standard Crosswalk*, available at https://www.getty.edu/research/publications/electronic_publications/intrometadata/crosswalks.html

[4]https://www.getty.edu/research/tools/vocabularies/

code, related terms, and sources. When possible, alignment with other standards is provided. As vocabularies have a hierarchical structure, each term specifies the hierarchy to which it belongs and the type of hierarchical relationship present. Internal association relationships within the vocabulary are generally provided as well.

The **Iconography Authority** (IA) includes proper names, relationships, and dates for iconographical narratives, characters, events, and literary works in a standard-compliant format. As it has a thesaural structure, it includes equivalence, associative, and hierarchical relationships. The thesaurus is aligned with other ones (e.g., Library of Congress and Iconclass), and it is collaboratively updated and enriched. Although the Getty vocabularies are progressively being published as Linked Open Data, the Iconography Authority has not been published in this format yet.

The **Art and Architecture Thesaurus** (AAT) includes generic concepts, not identifiable by a proper name, needed to catalog and retrieve information about art and architecture. Its development began in the 1970s to respond to the need for art journal indexing services and art libraries, which started to automate the indexing process[5]. His wide scope reflects the aim to meet the needs of various user groups (e.g., historians, architects, librarians, museum personnel, etc.) to allow cross-collection information retrieval. Descriptions of materials, abstract concepts, styles, roles, and techniques are included. AAT is one of the vocabularies exposed as Linked Open Data viewable in JSON-LD, RDF, N3 / Turtle, and N-Triples.

Iconclass

Iconclass[6] is a classification system for iconographical description with a standardized, controlled terminology widely used for research purposes and information retrieval in museums.

Iconclass is the first controlled vocabulary developed even before modern computer science's birth, designed in the 1940s by Henri van der Waal (van Straten, 2012). Based on a system of numbers and letters identifying each entry, it can uniquely identify each resource according to W3C Semantic Web indications. Each identifier is expressed as a URI[7] and is available in LOD.

Iconclass is one of the sources of the Getty Vocabularies. Unlike these, its structure does not have associative relationships between entries other than the

[5] https://www.getty.edu/research/tools/vocabularies/aat/about.html

[6] https://iconclass.org/

[7] A URI (Uniform Resource Identifier), in a Semantic Web context, is the unique identifier for resources. URIs of resources can be semantically linked (using RDF properties) to other resources (and their URIs). For additional information on the concept of URIs, we refer to https://www.w3.org/TR/webarch/

hierarchical relationship. Iconclass has a hierarchical structure based on ten macro-categories, identified by numbers from 0 to 9, and no associative relationships between entries are included. The first category was introduced later to describe contemporary art subjects (0 Abstract, Non-representational Art), including abstract shapes, colors, and concepts. The following sections describe the physical world (2 Nature ; 3 Human Being, Man in General) and aspects concerning society (1 Religion and magic, 4 Society, Civilization, Culture, 5 Abstract Ideas and Concepts, 6 History). The last three categories (7 Bible, 8 Literature, 9 Classical Mythology, and Ancient History) deal with the most famous literary sources of subjects in Western art. The remaining sections describe the world surrounding (2 Nature, 3 Human Being, Man in General) and aspects concerning society (1 Religion and magic, 4 Society, Civilization, Culture, 5 Abstract Ideas and Concepts, 6 History). Each category is further deepened by more specific terms, identifying 450 entries that gradually include more specific items. Specifications may refer to general concepts or individuals (e.g., *the Announcement of Christ's birth*, code 73A5), which can be more thoroughly described (e.g., *the Annunciation with Mary kneeling*, code 73A523).

Although both Getty Vocabularies and Iconclass are excellent solutions to support most iconographic recognition tasks, there are naturally situations in iconographical and iconological studies that cannot be represented through these taxonomies. The study by Baroncini et al. (2021) highlights the different situations in which a complex representation can not be fully represented. For instance, relations among subject entries and possibly shared symbolical meanings they express can not be stated. Moreover, the transmission of iconographical themes over time, treated by a consistent number of iconological studies, and their variation according to the period's style can not be described. Furthermore, Getty Vocabularies allows associating a date with a subject, but the stylistic variants cannot be related to the respective historical period. Likewise, Iconclass and Getty Vocabularies allow specifying the source of the iconography. However, it is impossible to relate a source to a variant of a subject (e.g., Durer's Rape of Europa depicted following the description provided by the Renaissance writer Poliziano). Therefore, to fully express the complexity embedded in the branch of iconological studies, other more flexible means are needed.

3.3 Ontologies

In the field of computer science, the term ontology identifies a formalization of a domain that consists of the definition of concepts and relations among them. As

a result, an ontology is a taxonomy describing the relevant term in a hierarchical structure (Staab & Studer, 2010). The Web Ontology Language (OWL)[8] is the standard used to define such concepts and relations.

3.3.1 Art history

In the context of art interpretations, several attempts have been made to create models that cover some specific elements related to interpretations (i.e., symbolic meanings) or the whole act of interpretation of a cultural heritage object[9]. CIDOC-CRM (Bekiari et al., 2021) is a widely used ontology in the context of cultural heritage. It has an event-based structure and covers fundamental aspects of the life cycle of a cultural heritage object. Carboni et al. extended it with the Visual Representation ontology (VIR) (Carboni & de Luca, 2019). VIR ontology explores the concept of visual representations in artworks. It associates the portion (called *iconographical atom*) of the cultural heritage object to the recognized subject by representing an interpretation act. Thanks to such structure, further specifications of the assertion can be added, such as the person expressing the interpretation and the sources on which it is based. Although founded on a thorough consideration of interpretation level theories described in Chapter 2, VIR focuses only on subjects of level 2, considering iconographies and their attributes, consequently lacking a clear distinction between levels.

The preliminary study in Baroncini et al. (2021) further extends VIR by adding an iconological interpretation class linking the concepts of the artwork and external cultural phenomena. It is evaluated on 11 real case studies taken from the literature in iconology, which illustrate a wide variety of aspects included in an iconological analysis. In addition, the work is based on a careful theoretical comparison of the main iconological and iconographical interpretation theories. For its comprehensive overview of iconographical and iconological theories, along with the real-base evaluation, it is used as a source for the ontology developed as part of the current Ph.D. project, which has to be seen as its development and refinement. We deepen this study by developing aspects not already considered, such as a more detailed description of level 1 and level 2 subjects and integrating multiple interpretations by different art historians (see section 5.3).

Gartner (2020) proposes an ontology to facilitate and automate the identification of subjects (level 2) in works of art through logical inferences. No alignments were possible to this ontology because it has not been released.

The ARCO ontology (V. A. Carriero et al., 2019) was developed to model

[8] https://www.w3.org/2001/sw/wiki/OWL

[9] Section 3.3 is based on the version published in (Sartini et al., 2023)

artifacts from Italian cultural heritage by converting the information contained in traditional catalog sheets into Linked Open Data. Among the possible aspects modeled for an artwork, some classes were designed to describe its iconographical apparatus. Apart from this class, the schema does not mention any distinctions between different levels of interpretations. Most of the information about the iconographical and iconological interpretations in ArCo are provided through natural language descriptions with the property `dc:description` or `core:description`[10], not exploiting the full potential of the Semantic Web (Sartini & Gangemi, 2021). Sartini et al. Sartini et al. (2021) modeled symbolic meanings in cultural heritage in the Simulation Ontology. This ontology does not consider the hermeneutic act of interpretation of the association of symbolic meanings with artworks. However, its conceptualization of symbolic meanings using n-ary relationship classes that link a symbol, its symbolic meaning, and the cultural context in which the symbol-symbolic meaning relationship proposes a thorough description of symbols. Therefore, we reuse the Simulation ontology classes and properties to express the symbolic meanings in our work, inserting them in the context of an interpretation of an artwork.

3.3.2 Interpretations and meaning

In the context of knowledge organization, several ontologies addressed the concept of interpretation. CIDOC-CRM models assertions with the class `E13_Attribute_Assignment`, which relates the assertion made by one agent to the object considered. Since each assertion reflects the agent's opinion, multiple, contradictory assertions may be represented. The concept of interpretation is applied broadly, including measurements and other types of scientific observations. Similarly, the class `Interpretation` of Arco is intended to describe every piece of information asserted by an agent about an object on the basis of stated sources.[11] The CIDOC-CRM extension CRMinf deepens the concept expressed by `E13_Attribute_Assignment` distinguishing the type of argumentation and if the belief resulting from the argumentation holds true or not. The concept is further explored by the CRMsci, another CIDOC-CRM extension, which integrates CRMinf by formalising the shared scientific process adopted across different domains and the scientific activities involved. In detail, of great interest is the class `crmsci:S4_Observation`, subclass of `crmsci:I1_Argumentation` and of `crm:E13`, expressing the scientific observation of physical events or reality which is done

[10] See https://dati.beniculturali.it/lodview-arco/resource/HistoricOrArtisticProperty/0500653281.html dc:description value.

[11] https://dati.beniculturali.it/lodview-arco-onto/ontology/context-description/Interpretation.html

directly or through measurements. It represents the transition between reality and propositions‘´ (Doerr et al., 2014). Furthermore, the VIR ontology adds a domain-definition of `crmsci:S4 Observation` by declaring its subclass `vir:IC12 Representation`, which represents an assignment of a solely iconographical status to a physical object.

The same topic is also addressed by the history domain to represent the frequent case of disagreeing historians' interpretations of the same events. As reported by Frank (2019), several ontologies afford the theme by modeling different views of the same observed events, such as the SEM ontology[12] MIDM (Van Ruymbeke et al., 2017), the ODP Event-Model-F[13], expanding DOLCE+DnS Ultralite (DUL), in which a distinction between facts and interpretations has already been made (Gangemi & Mika, 2003). The STAR model developed in the Releven project (Andrews, 2023) reuses CIDOC-CRM to describe the provenance of the assertion and its authority. HiCO (Daquino & Tomasi, 2015)[14] goes further by adding contextual information to the interpretation, such as interpretation type and criterion.

3.4 Computer vision

Besides the current work that focuses on representing the content of the artwork as descriptive metadata, it is crucial to assess what information can be currently extracted from the artworks using alternative techniques. As the main object of study, i.e., the artwork, is of a visual nature, we provide a brief overview of the knowledge that can currently be extracted from artworks' pictures through computer vision (CV) techniques. Computer vision is a branch of Artificial Intelligence that focuses on processing visual images to interpret and understand them.

Since most CV algorithms are trained on photographs, the task is challenging when applied to art pieces. Nevertheless, the literature shows an increasing interest in applying it to works of art, offering solutions beyond this limitation. Several attempts have been made in Computer Vision (CV) to detect objects, subjects, and visual patterns in artworks. Some studies concern the detection of iconographies with a focus on Saint's identification in the context of Western Christian Art. Among them, Stork et al. (2021) exploit the art practice of using attributes to identify a saint (i.e., the keys held by a man indicate that the figure is Saint Peter), performing good results in identifying ten saints. Nevertheless, the ambiguity that attributes have (i.e., saints identified by the same object) prevents

[12] https://semanticweb.cs.vu.nl/2009/11/sem/
[13] http://ontologydesignpatterns.org/wiki/Ontology:Event_Model_F
[14] http://purl.org/emmedi/hico

them from extending the methodology to all the Christian saints. Similarly, Milani and Fraternali (2021) trained a Convolutional Neural Network (CNN) to automatically classify a corpus of artworks according to 19 iconographical classes, mainly identifying Christian characters. Madhu et al. (2019) focused on the facial and body recognition of Mary and Gabriel in the *Annunciation* scene by applying a gender detection algorithm. Another approach, taken by Marinescu et al. (2020) to improve the precision of object detection in paintings, introduces semantic information about the time in which objects were used to correct CV output. Further studies have addressed the central theme of visual pattern detection in artworks. Ufer et al. (2021) developed style-independent research, which detects visual patterns in the overall image or a portion of it, offering the opportunity to search for a combination of patterns. The retrieval of images having similar compositional structure is addressed by Madhu et al. (2023), based on the iconic language of the images afforded by Max Imdhal. Impett (2020) trained an algorithm to detect recurring poses in artworks based on Warburg's attention to gestures and body position in art. The results are shown in the *Gesture Atlas* project[15], in which figures from artworks are clustered according to gesture similarity. Brandhorst (2022) suggests the use of mixed Iconclass-based retrieval and Computer Vision search to retrieve similar subjects that have the same visual arrangement.

The progress made in identifying icons and subjects through CV is promising. Nevertheless, some aspects could be more challenging, such as resolving ambiguities in recognizing saints and historically feasible object detection. Moreover, to the author's knowledge, no CV algorithm included the detection of canonical deeper meanings, e.g., symbolical meanings or moral messages. Furthermore, it seems highly challenging to detect with such an approach the complexity of cultural meanings unveiled by iconological studies.

Therefore, the availability of structured data on iconographies, historical objects, and deeper meanings is essential to advance the state-of-the-art of CV performances for the recognition of iconographical and iconological aspects in works of art.

3.5 Data analysis for art history

Museums are increasingly publishing open source data on their collections. Nevertheless, to date, only a few studies have exploited computational methods for art analysis. The major areas in which computational or quantitative approaches are cultural economics and cultural and social art history. Indeed, being interdisciplinary fields, they combine quantitative analyses over large datasets of sociology

[15] https://biblhertz.github.io/atlas/

and economy with art topics. The artworks-related inquiries of cultural economics usually involve art market studies, generally focusing on a specific period. Among them, we cite the recent study by Greenwald (2021) in which the art circulated in France, England, and the United States during the XIX century is quantitatively analyzed. The study provides a robust methodological approach in which domain-specific research questions drive quantitative inquiries.

For its familiarity with data analysis, cultural economics is also the domain in which many studies exploit semantic web technologies and artificial intelligence. Schich et al. (2017, 2014) perform network analysis of the movement of notable historical figures, quantitatively seeing some characteristics of the art market, such as the importance of certain centers (e.g., Paris) (Kienle, 2017). Filipiak et al. (2016) and Rother et. al. (2022, 2023) use the Natural Language Processing technique to extract semantic information from provenance datasets and represent them in RDF format.

Other types of analysis were proposed based on image analytics. Klinke underlines the potential that a distant viewing has for seeing similarities and patterns in artwork images. Similarly, Manovich proposes cultural analytics applied to images (Manovich, 2015, 2020). A similar interest is witnessed by Kim et al. (2014), who quantitatively compare artworks of different periods on some parameters, such as brightness and use of colors.

To the best of the authors' knowledge, no quantitative analysis of iconographical and iconological topics is currently available.

3.6 Iconography and Iconology statements on the Semantic Web: a survey

Recent years have witnessed a growing interest in linked open data describing Cultural Heritage (Davis & Heravi, 2021).[16]. Despite many Cultural Institutions releasing their data only in a simple tabular form, several Knowledge Graphs (KG) address the description of artworks in a more structured, logical form. Some of them, e.g. Wikidata (Vrandečić & Krötzsch, 2014), have a general scope and are created collaboratively, while others, (e.g. ArCo (V. A. Carriero et al., 2019), Zeri & Lode (Daquino et al., 2017) are generated by the conversion of authoritative data from cultural institutions. Other art-related projects exposed RDF datasets to provide

[16]This section is based on the article (Baroncini, Sartini, et al., 2023) In particular, section 3.6.2 refers to article Section 4, for which S. Baroncini is responsible. Sections 3.6.1 and 3.6.3 refer to the article's Sections 2 and 5, for which B. Sartini is responsible. Section 3.6.4 is based on article's section 6, for which B. Sartini and S. Baroncini are responsible. Both authors contributed equally to the research.

data about the socio-cultural context in which artworks were created. An example is the Golden Agent project[17], which provides a set of LOD datasets containing information about art creation and market of the Duch Golden Age, a particularly prosperous period for art making. To this aim, they extract resources concerning the production and consumption of creative goods in that period from biographical records (Brouwer & Nijboer, 2017), offering extremely valuable research for domain studies.

In this diversified setting, it is crucial to assess the coverage, accuracy, and reliability of the available data to allow their reuse for domain-specific purposes. While many studies addressed the problem of KG evaluation methods, to the authors' knowledge, a survey on Art History information stored in Knowledge Graphs, comprehensive of an assessment of the data quality, is still missing. Therefore, the work in this section aims to evaluate the coverage of the content represented in visual works over existing KGs, focusing on iconographical and iconological aspects (i.e., artistic subjects and their symbolic and cultural meanings). We survey KG evaluation methodologies and adapt some of their metrics to the considered domain of knowledge. Furthermore, theories concerning the iconographical and iconological (*icon*) domain are reviewed to assess how much KGs cover information about visual items' subject and content description.

Semantic web technologies offer an opportunity to express semantically complex information formally. For this reason, they are a suitable means to express fields of study as complex as iconography and iconology at the required granularity.

Artwork contents should be analyzed both isolated, i.e., by identifying relevant features and associating them to features of other artworks (e.g., the study of patterns recurring in different subjects (Warburg, 1999; Wittkower, 1977)). Therefore, the knowledge emerging from an analytic approach is mostly missed when an artwork's content is described just by a general subject term.

The traditional sources of knowledge are natural language descriptions of artworks found in texts. However, texts need knowledge extraction methods to enable further analysis and interlinking, limiting the computational reuse of that knowledge (Sartini & Gangemi, 2021).

Another problem is the lack of advanced ontologies that provide a detailed semantic form for artwork description data (see Section 3.3.

In addition, since iconographical-iconological analysis can potentially involve very different types of cultural objects, often stored by different institutions, the major benefits of storing information about this domain in knowledge graphs include at least:

[17] https://www.goldenagents.org/

- The opportunity to answer domain-specific questions through quantitative analysis (e.g., which attributes and meanings were related to the mythological character of Mercury across the centuries?);
- Accessing and querying interlinked information about worldwide objects that could not otherwise be experienced together (e.g., all artworks with political implications stored in different museums worldwide);
- Formally expressing the semantic complexity of the topic (e.g., the levels of meanings of an artwork and its relations to external resources, such as other artworks, texts, etc.).

In this section, we assess the available data accuracy, reliability, and interoperability in the iconographical and iconological domain of knowledge. Therefore, the major benefit is to provide domain experts with a clear state of the quality of semantic, domain-specific data available online. Other benefits include improving current data reuse following LOD principles and fostering the creation of a shared semantic description framework for iconology and iconography. With this analysis, we show the reusability potential of the existing KGs based on defined *icon* requirements. Finally, the main findings of this survey are shown in a landscape (Figure 3.2) in which KGs are positioned according to their performance in the chosen metrics.

3.6.1 Knowledge graph evaluation methods

The reference standard for data quality is ISO 25012[18], in which 15 data quality characteristics are identified and classified as inherent and system-dependent data quality. Even if only a few data quality assessment methodologies explicitly refer to the standard (Radulovic et al., 2018), its metrics are generally included in all the proposed methods. Although the standard constitutes a reference point, knowledge graphs require specific evaluation methods and metrics. They show differences from traditional relational databases in their structure (graph versus table), the reasoning possibilities that can be applied to them, and their interoperability and interconnections (Janev et al., 2020). Ji et al. (2022) surveys evaluation metrics and methodologies for representation learning, knowledge acquisition, and completion, with additional analyses over temporal KGs and applications developed from them. Paulheim (2017) provides a series of refinements methods to increase the quality of knowledge graphs. Pellegrino et al. (2023) evaluates Cultural Heritage knowledge graphs regarding their suitability for question-answering tasks. Zaveri et al. (2016) proposes a conceptual framework for quantitative and qualitative

[18] https://iso25000.com/index.php/en/iso-25000-standards/iso-25012

metrics in evaluating knowledge graphs from a study of more than 100 scholarly publications. Various general metrics for knowledge graph quality evaluation and applications thereof are provided in (Färber et al., 2018). We reuse parts of these metrics, adapting some to focus on the fields of iconography and iconology (see Section 3.6.3). Behkamal et al. (2014) presents a similar study but uses the Goal-question metric paradigm to assess the quality of knowledge graphs. Ringler and Paulheim (2017) also compares several general domain knowledge graphs in their content coverage. It contains interesting reflections, particularly regarding coverage of artistic fields, in which YAGO and DBpedia seem to be the most detailed. Heist et al. (2020) uses coverage and a metric for evaluation, although this work does not mention cultural heritage-related findings. Shenoy et al. (2022) evaluates Wikidata on schema violations and deprecated entities, looking at its history of updates. Freire and Isaac (2020) also evaluates Wikidata's completeness in the description of data related to cultural heritage. To do so, the information contained in it is compared with the information available on Europeana (Isaac & Haslhofer, 2013), which is used as a gold standard for completeness. This study does not mention specific aspects related to iconography and iconology. Issa et al. (2021) offers a thorough study on the completeness metric when evaluating knowledge graphs. Finally, Ruan et al. (2016) introduces the concepts of queriability to knowledge graphs, developing a framework for evaluating *quality in use*, applying it to DBpedia and YAGO. Queriability is a very interesting concept when extracting relatively complex sets of information from knowledge graphs. Such complex relationships might be present in knowledge graphs that describe artworks with high granularity. However, to verify the queriability of the *icon* content, a first assessment of what is currently included in a knowledge graph is needed.

In summary, previous work evaluates knowledge graph suitability for some automatic tasks or their content in terms of various metrics that go from completeness to accuracy to quality in use. Some of them focus on specific fields (such as cultural heritage). There is no study yet that evaluates specific aspects related to iconology and iconography in knowledge graphs, which would require a specific evaluation due to the complexity of the information expressed in this domain of knowledge (Baroncini et al., 2021; Sartini et al., 2021). Therefore, the contribution of the current survey is to adapt a selection of general metrics from the literature to the domain-specific needs, adding a newly created metric. As a result, this contribution attempts to provide a domain-specific overview of the quality of available data according to the domain focus of interest and research questions.

3.6.2 Selection of the knowledge graphs

To collect the most representative RDF data about the description of the artwork, we need to consider which kind of cultural objects can represent a visual subject and can have a cultural meaning. Potentially, every image representing a subject that can be invested with a cultural meaning can be considered by an iconographical-iconological interpretation. To narrow down the research in the art history field, we focus our selection on paintings, sculptures, frescoes, visual subjects on coins (numismatics), and illuminations. Therefore, in this survey, we considered graphs containing data on Cultural Heritage, Museums, Libraries (manuscripts' drawings and decorations), and numismatics. In addition, we included general-purpose knowledge graphs likely containing information about artworks such as Wikidata, DBpedia (Auer et al., 2007) and YAGO (Rebele et al., 2016).

We used the following methodology. We first define our object of interest, namely artworks and information about their subject and meaning. Then, we collect the KGs through 1) the analysis of literature concerning a survey or evaluation of CH KGs (Bikakis et al., 2021; Pellegrino et al., 2023; Savnik et al., 2021) and 2) direct search on the web, through a manual keyword search on Google Database Index[19] and other main databases search engines.[20][21][22] This led to 56 graphs. These graphs were further pruned according to the criteria of their online availability through a SPARQL endpoint[23]. We considered these criteria fundamental to assessing data that follows the principle of availability and reusability of the Semantic Web (Wilkinson et al., 2016), according to its shared standards[24].

Only 27 of 56 graphs were active online, 18 with a SPARQL endpoint. The KGs for which the SPARQL endpoint was not responsive and those without information about subjects were discarded. Consequently, we obtained nine graphs. Table 3.2 gives an overview of the number of artworks having a subject, distinguishing between URIs and literals.[25][26] This analysis was conducted through SPARQL

[19] https://datasetsearch.research.google.com/
[20] https://datahub.io/
[21] https://triplydb.com/
[22] https://www.kaggle.com/
[23] SPARQL is the query language used to retrieve information from RDF data. SPARQL endpoints are online services linked to specific knowledge graphs that let users query knowledge graphs through SPARQL queries. For additional information about SPARQL and SPARQL endpoints, we refer to https://www.w3.org/TR/rdf-sparql-query/
[24] https://www.w3.org/standards/semanticweb/query
[25] Literals represent basic data types, such as strings, boolean values, and integers. They are not assigned a URI, and therefore, they can only be referred to as the object of a triple and never as the subject. Literals contain unstructured information (such as natural language descriptions) that might require additional processing before being machine-readable. For additional information about literals, we refer to https://www.w3.org/TR/rdf11-concepts/
[26] For an overview of the relations considered to identify subjects and other *icon* information, see Section 3.6.3.2

queries and by consulting the KGs' documentation. The selection process of our analysis highlights how information about cultural heritage is very scarce when considering data that follows Semantic Web principles, as few domain-specific knowledge graphs are available under those conditions. This makes the inclusion of general domain knowledge graphs essential to assess how *icon* aspects are described in the Semantic Web, as the majority of *icon* data is stored in them. From a structural perspective, we would expect the ontological schemas[27] of domain-specific knowledge graphs to describe *icon* information with a higher degree of granularity compared to general ones. This assumption is proved wrong by our results (section 3.6.4), as Wikidata performs better than domain-specific KGs.

Table 3.1: Classes and properties related to the recognition of artworks (sculptures and paintings if available) in selected knowledge graphs

Name (Abbreviation)	**Artwork (paintings and sculptures if possible)**
ArCo	<artwork> a arco:HistoricOrArtisticProperty
Zeri&Lode (Zeri)	<artwork> a fabio:ArtisticWork
Nomisma	<artwork> nmo:hasObverse <something> <artwork> nmo:hasReverse <something>
Wikidata	<artwork> wdt:P31 wd:Q3305213 (Painting) <artwork> wdt:P31 wd:Q860861 (Sculpture)
RDS Platform (SARI)	<artwork> a gndo:Work <artwork> a gndo:formOfWorkAndExpression
Europeana	<artwork> a http://vocab.getty.edu/aat/300033618 <artwork> a http://vocab.getty.edu/aat/300047090
National Digital Data Archive of Hungary (ND_Hungary)	<artwork> a dcmitype:Image
DBpedia	<artwork> a dbo:Artwork
YAGO	<artwork> a schema:Painting <artwork> a schema:Sculpture

One critical aspect we encountered while doing this analysis is the proper identification of what is a work of art. While some graphs use a specific class or property to express it (e.g., `fabio:ArtisticWork` in Zeri & Lode), others do not have a unique way to identify it. In some cases, e.g., Wikidata, many specific classes are used, subclasses of a general 'visual work'. In others, e.g., SARI's RDS platform, the class 'Work' corresponds to many different types of cultural objects specified by a controlled vocabulary. Although this granularity in the artwork

[27]We consider the ontological schema as the portions of those ontologies that knowledge graphs actually use as a data model. There exist several general domain schemas, such as Dublin Core https://www.dublincore.org/specifications/dublin-core/, Simple Knowledge Organization System (SKOS) https://www.w3.org/TR/skos-reference/, or Friend of a Friend (FOAF) http://xmlns.com/foaf/0.1/, that are reused in many different knowledge graphs. Domain-specific knowledge graph schema might include specifically developed ontologies in their schema. See the ArCo Ontology (https://w3id.org/arco/ontology/arco) for the ArCo knowledge graph, or the Nomisma ontology (https://nomisma.org/ontology) for the Nomisma knowledge graph

description is appreciable, it may generate a few issues when approaching data quantitatively. First, the selection of what is considered an artwork is left to the user, whom subjective decisions may influence in this definition. Second, the high number of entities to be included in a SPARQL query can influence the server response.

In the context of this study, we selected which classes could be considered artworks from the analysis of the documentation or from data retrieval. We decided to focus our attention on paintings and sculptures when available (if the information present in the knowledge graphs made them distinguishable from other artworks), as they are universally considered as artworks with at least a subject. When paintings and sculptures were unavailable in the studied knowledge graph, we shifted our attention to the most prominent class in the schema that could represent an artwork (as the numismatic items in Nomisma). On the other hand, when the total number of sculptures and paintings was too little for conducting an evaluation (e.g., in SARI's RDS platform), we included in the analysis broader terms, such as prints, illustrations, and graphics. Table 3.1 summarizes classes that define artworks from the selected KGs, along with properties that link information relevant to iconography and iconology.

3.6.3 Evaluation criteria

Following the approach presented in (Wang & Strong, 1996), we define metrics that go beyond accuracy, as we are interested in i) the coverage of the KGs schemas and their data, ii) the references and interlinking with existing taxonomies that identify subjects in art (Iconclass, Getty), iii) alignments, and iv) linking to external knowledge graphs to foster poly-vocality in art interpretations. These general metrics were adapted to evaluate the specific knowledge domain to obtain a specific quality assessment of domain data. In addition, these metrics acquire a particular relevance for the domain studies, which analyze the relations between cultural objects, their sources, and multiple interpretations. Following the Panofsky theory explained in Section 2.2.2, we are interested in analyzing whether the current knowledge graphs distinguish between elements that belong to the first, second, and third levels of interpretation. We are therefore looking for clear distinctions when it comes to the description of natural elements depicted in a painting, the recognition of subjects and symbols, and the reflections of the influence of the cultural period in which the artwork was created on the artwork itself and vice versa.

Taking this into consideration, we applied parts of the framework formulated in (Färber et al., 2018) in the evaluation of the chosen KGs. This study proposes the

Table 3.2: Overview of the artwork subject presence in the selected graphs (as of 01/12/2022)

Short Name	Artwork	Percentage of artworks having a subject (URI)	Average of subjects (URI) defined per artwork	Percentage of artworks having a subject (literal)	Average of subject (literal) defined per artwork
ArCo	2111726	45.86%	1.01	100	1.22
Fondazione Zeri	20082	99.99%	1.19	0	0
Nomisma	566732	21.16%	1.1	0	0
Wikidata	669857	26.76%	3.37	0	0
SARI	339	72.57%	1.11	0	0
Europeana	13861	9.32%	2.38	33.8	1.64
ND_Hungary	11655	0%	0	55.97	6.04
DBpedia	12250	93.93%	5.7	4.8	1.09
YAGO	29324	12.75%	1.02	0	0

possibility of a weighting system applied to each metric according to the importance of the task in the context of the evaluation. In our case, we give more weight to the evaluation criteria referring to the elements addressed the most in the literature of *icon* studies. Specifically, we assign the maximum weight (1) to those criteria that we consider wholly related to iconography and iconology evaluation, 0.8 to those that we consider closely related, and 0.6 to those criteria that we consider partially related. All other criteria are excluded; considering their weight would be 0, they were not computed. Therefore, of all the categories described by (Färber et al., 2018), we focus only on column completeness, schema completeness, semantic validity, reference to external vocabularies, and interlinking via owl:sameAs.[28] We adapted all metrics cited above to address the specific tasks of evaluating the *icon* content. As a result of the adaptation, we decided to rename them to address their new specific purpose. Column completeness was changed into Iconographical and Iconological column completeness (IICC), semantic validity became Semantic validity of Iconographical and Iconological triples (SVIIT) schema completeness became Iconographical and Iconological schema granularity (IISG), reference to external vocabularies became References to external taxonomies of art and culture (RETAC) and Interlinking via owl:sameAs became Interlinking of artworks (IA). The differences and specific changes applied to these metrics will be explained in the sub-paragraphs of this section. Finally, we added a new metric to measure intralinking potential for subject comparisons (IPSC).

Table3.3 summarizes (i) the reused metrics plus the newly created one, (ii) their adaptation to the *icon* field, and (iii) the weight assigned to the metric. We applied these measurements to the knowledge graphs listed in Section 3.6.2. We then grouped these metrics into 2 macro-categories, namely i) structure of the knowledge graphs, which includes IISG, IA, RETAC, IPSC, and ii) content of the knowledge graphs, which includes SVIIT and IICC. The analysis results and the formulas used to calculate the overall score will be discussed in Section 3.6.4.

3.6.3.1 Evaluation Methodology

Of the chosen metrics, three (interlinking of artworks, references to external taxonomies of art and culture, and intralinking potential for subject comparisons) could be processed automatically by analyzing the data, one through an analysis of the schemas of the various KGs (iconographical and iconological schema granularity), and two required qualitative evaluations (semantic validity of iconographical and iconological triples and iconographical and iconological column completeness). For all automatic evaluations, a series of SPARQL queries were launched on the

[28]The cited categories will be thoroughly explained in the following part of this section

Table 3.3: Evaluation metrics, the first five criteria are adapted from (Färber et al. 2018), the last criterion is newly developed

Area	Criterium	Adaptation	Weight [0-1]
Content	Semantic Validity	Semantic Validity of Iconographicaland Iconological Triples (SVIIT)	1
Content	Column Completeness	Iconographical and Iconological Column Completeness (IICC)	1
Structure	Schema Completeness	Iconographical and IconologicalSchema Granularity (IISG)	1
Structure	Using External Vocabulary	References to External Taxonomiesof Art and Culture (RETAC)	0.8
Structure	Interlinking via owl:sameAs	Interlinking of Artworks (via variousproperties) (IA)	0.6
Structure	Intralinking Potential for Subject Comparisons (IPSC)		0.6

analyzed graph, and some will be listed as examples in the following subsections. For the metrics that required a qualitative evaluation of the content, we extracted random representative samples of the knowledge graphs and evaluated the graphs manually on those samples through annotations.

All annotations were performed by two annotators.[29] In the annotation process, they could express their inability to evaluate the veracity of some of the triples if the information contained in the knowledge graph was unreachable (broken links) or too scarce to fully assess its quality. We used Cohen's Kappa (using quadratic weights) (Cohen, 1960) to measure the agreement score between the annotators. The triples considered invalid by annotators were mutually excluded when computing these agreement metrics.[30] Given the general agreements of the two annotators for all the different samples annotated, as shown in Table 3.4, we decided to average the evaluation scores of the two annotators for both the qualitative categories.

In the following paragraphs, the metrics and our computations to obtain them are described in natural language and their mathematical formulas.

3.6.3.2 Iconographical and Iconological Schema Granularity

This metric is a re-elaboration of the 'Schema completeness' metric in (Färber et al., 2018). Schema granularity aims to verify to what extent the ontologies and vocabularies, and corresponding classes and properties instantiated in the knowledge graphs, cover the domain of interest. In this work, we verify to what extent the schema of the knowledge graph is suited for the complete description of *icon* elements. Based on the comparison of theories of art interpretation discussed

[29] One of the annotators was an Art History domain expert.
[30] Only 3,3% of total evaluated triples was considered invalid

Table 3.4: Inter-annotator agreement scores as measured by quadratically weighted Cohen's Kappa for semantic validity of iconographical and iconological triples and column completeness per knowledge graph

Knowledge Graph	Semantic Validity	Column Completeness
Yago	1.00	0.65
Nd hungary	0.82	0.62
ArCo	0.77	0.77
Zeri	0.66	0.78
Nomisma	1.00	1.00
Sari	0.78	0.68
Europeana	0.82	0.79
DBpedia	0.89	0.66
Wikidata	0.76	0.90

in Section 2.2.2, we formulate the following competency questions, which should be possible to apply to a work of art or a part of it (e.g., a detail or a scene of a pictorial cycle) (Uschold & Grüninger, 1996):

- What are the pre-iconographical elements that appear in a work of art?
- Which actions are depicted in a work of art?
- What are the subjects of a work of art?
- What are the represented symbols in a work of art?
- What are the represented stories in a work of art?
- What are the represented allegories in a work of art?
- What are the intrinsic meanings associated with a work of art?
- Which cultural phenomena are reflected in a work of art?
- What are the corresponding external taxonomies for the identified iconographical terms?

We then created a gold standard interpretation on the example from Michelangelo's work, able to answer those competency questions, as shown in Figure 3.1. We first aligned the properties used in each KG to our example and computed schema granularity as the division between the number of properties of the example that have been aligned and the total number of properties in the example. Given N as the number of properties of the gold standard and N_{akg} as the number of properties of the same gold standard aligned to the properties of the schema of the knowledge

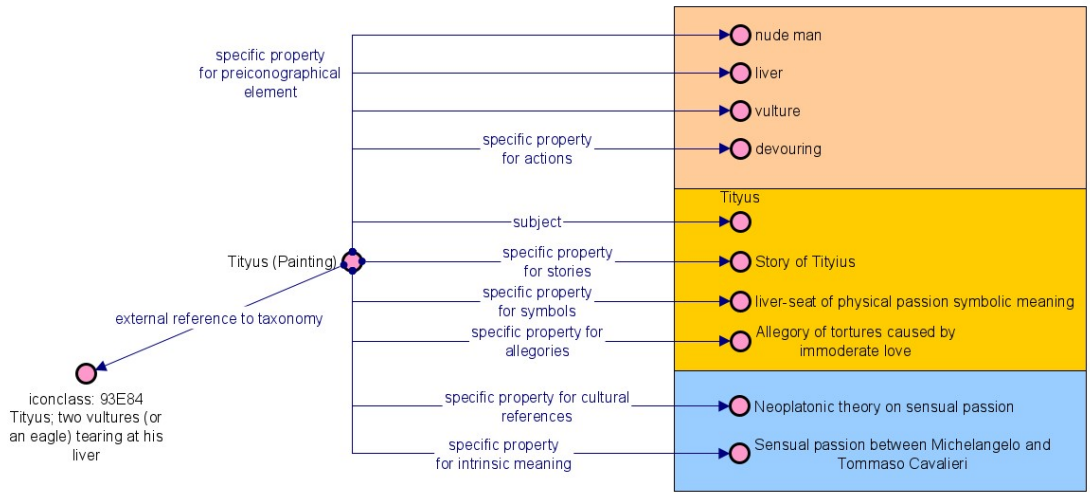

Figure 3.1: The gold standard schema created by the application of CQs to the golden example from the literature

graph, we measure the Iconographical and Iconological Schema Granularity ($IISG$) of a knowledge graph as

$$IISG(kg) = \frac{N_{akg}}{N}$$

Table 3.5 shows those properties that were recognized as expressing *icon* content and were aligned to the gold standard.

We weigh this metric as 1 because a schema that permits to express *icon* statements, respecting the required granularity given by the complexity of their field, is essential to correctly and completely store information on this matter.

3.6.3.3 Semantic Validity of Iconographical and Iconological Triples

This metric was modified from the 'Semantic Validity' of (Färber et al., 2018), whose purpose is to define whether all the statements of triples in knowledge graphs hold true or not. In our study, we consider the semantic validity of *icon* triples only: we evaluate whether triples that refer to a subject, depicted element, or symbol associated with a painting hold true. To evaluate this, we take a subset of the *icon* statements in each KG. Those statements link the artwork to one of the elements relative to the three layers of interpretation explained in Section 2.2.2, agnostic to the property used. We compute this metric by taking a random sample of 100 iconographical/iconological triples from each knowledge graph, evaluating whether the triple is correct (1), partially correct (0.5), or wrong (0). Given S_{ictkg} as the random set of iconographical triples extracted from a knowledge graph, and $S_{evictkg}$ as the evaluation scores set given for each triple $\{sc_1, sc_2...sc_x\}$, and x as the sample size[31] to be extracted from the knowledge graph, the Semantic Validity

[31] In our case set as 100

Table 3.5: Properties identifying iconographical and iconological content for each selected knowledge graph

Name	Iconographic and Iconologic Properties
ARCO	arco-cd:hasSubject arco-dd:hasIconographicOrDecorativeApparatus arco-cd:iconclassCode arco-cd:subject dc:subject
Zeri	fabio:hasSubjectTerm
Nomisma	nmo:hasPortrait nmo:hasIconography nmo:hasControlMark
Wikidata	wdt:P180 (depicts) wdt:P921 (main subject) wdt:P1257 (depicts iconclass notation) wdt:P4878 (symbolizes) (qualifier of wdt:P180) wdt:P6022 (expression, gesture or body pose) (qualifier of wdt:P180)
SARI	gndo:topic gndo:gndSubjectCategory
Europeana	dc:subject
ND Hungary	dc:subject
DBpedia	dc:subject dbp:subject dbp:symbol dbp:symbols
YAGO	schema:about

of Iconographical and Iconological Triples ($SVIIT$) is measured as follows

$$SVIIT(kg) = \frac{\sum_{i \in S_{evictkg}} i}{x}$$

As, this metric offers key insights on the quality of the *icon* content of knowledge graphs, which is fundamental for their reuse, we give it a weight of 1.

3.6.3.4 Iconographical and Iconological Column Completeness

This metric, in (Färber et al., 2018), considers the general column completeness of knowledge graphs. In our work, we focus only on the column completeness of *icon* statements. Considering the potentiality expressed in a knowledge graph through the iconographical and iconological schema granularity, we evaluate the column completeness as the schema in use. We extract subgraphs from the analyzed KGs that contain all the *icon* triples associated with 100 randomly selected artworks per KG. This evaluation considers two aspects:

1. the expected number of layers of an artwork. Generally, a landscape only contains elements belonging to the first layer, a portrait contains the first layer and then the identification of the subject (second layer), and more complex artworks that represent cultural and religious themes can also be analyzed at a third, iconological level. Despite the potential for every visual image to have a deeper level of interpretation (van Straten, 2012), we decided to expect a third layer only in artworks presenting an explicit cultural subject. This is meant to not affect the artworks' evaluation with the bias of over-interpretation, criticized by some scholars (Gombrich, 1948)

2. the number of layers covered by the current description in the knowledge graph.

We then divide the covered layers by the expected layers for each artwork in the subset. Having a maximum of three layers, the possible scores for each artwork can be 0 (0 covered layers out of 3 expected, 0/2, 0/1), 0.33 (1/3), 0.5 (1/2), 0.66 (2/3), 1 (1/1, 2/2, 3/3). We do not expect artworks to be described meticulously by indicating every single element of level 1, every single recognizable subject, allegory, and symbol of level 2, and every single intrinsic meaning and culturally related meaning of level 3[32]; for this evaluation, having at least one element for every expected level was considered enough. Given A as the set of the randomly

[32]Especially considering that in the field of iconography and iconology, there could be potentially endless different interpretations of a painting, and it is not possible list them all

sampled artworks in the knowledge graph of size x[33] $\{a_1...a_x\}$, EL as the array of expected layers (a number from one to three) for each artwork

$$EL = \begin{bmatrix} el_1 & el_2 & el_x \end{bmatrix}$$

in A, and CL as the array of covered layers for each artwork

$$CL = \begin{bmatrix} cl_1 & cl_2 & cl_x \end{bmatrix}$$

we create the array SL that contains the divisions between covered and expected layers

$$SL = \begin{bmatrix} \frac{cl_1}{el_1} & \frac{cl_2}{el_2} & \frac{cl_x}{el_3} \end{bmatrix}$$

and then we measure the Iconographical and Iconological Column Completeness (IICC) of a knowledge graph as follows

$$IICC(kg) = \frac{\sum_{i \in SL} i}{x}$$

We consider this metric as important as having a schema that permits a certain degree of granularity in artwork descriptions; therefore, we give it a weighing of 1.

3.6.3.5 Interlinking of artworks

We adapted the metric 'Interlinking via owl:sameAs' described by (Färber et al., 2018) to only apply to artworks. 'Interlinking' is considered as the connection between entities belonging to different knowledge graphs. Although less central than the other used metrics (weight = 0.6), we decided to include it because aligning artworks across different knowledge graphs fosters poly-vocality in art interpretation, especially if these knowledge graphs have been manually curated.[34] We measure this metric by dividing the number of artworks in a knowledge graph that are connected to their corresponding versions in external knowledge graphs by the total number of artworks present in a knowledge graph. The main property used to align artwork across different KGs is `owl:sameAs`, but we also looked at other possible alignments from the analyzed KGs.[35]

[33] In our case set as 100

[34] We acknowledge that poly-vocality can be achieved also by giving iconographical and iconological assertions a provenance (even in the same knowledge graphs), although for this work we only focus on statements agnostic to the provenance of the interpretation, which would require another specific study

[35] The link to external artworks is expressed 1) in Europeana through the relations `dc:relation` or `edm:relatedTo`, 2) in Wikidata through different `wikibase:identifier`, 3) in ARCO and Zeri&Lode through `rdfs:seeAlso`, beyond `owl:sameAs`

Given KG as the set of triples $\{t_1...t_n\}$ in a knowledge graph (a triple being a sequence of subject, predicate, object $\{s_i, p_j, o_k\}$), A as the set of artworks $\{a_1...a_m\}$ denoted by s_i or o_k, and R_a as the set of relationships $\{r_1...r_z\}$ that are used to align an artwork in a knowledge graph to the same artwork in other knowledge graphs, we consider $A_a = \{a_1...a_w\}$ as a subset of A if

$$\forall a_i \in A_a : a_i \in A \wedge (\exists p_j \exists o_k : (a_i, p_j, o_k) \in KG \wedge p_j \in R_a)$$

and we measure Interlinking of Artworks (IA) as

$$IA(kg) = \frac{n(A_a)}{n(A)}$$

Two example queries launched on DBpedia to count the number of artworks and the number of artworks aligned to different KGs can be seen in listings 3.1 and 3.2 respectively.

Listing 3.1: SPARQL query launched on DBpedia to count the number of artworks

```
SELECT (COUNT(DISTINCT ?artwork) as ?tot)
WHERE { ?artwork rdf:type/rdfs:subClassOf* dbo:Artwork }
```

Listing 3.2: SPARQL query launched on DBpedia to count the number of artworks aligned to external knowledge graphs

```
SELECT (COUNT(DISTINCT ?artwork) as ?tot)
WHERE { ?artwork a dbo:Artwork; owl:sameAs ?x }
```

3.6.3.6 References to external taxonomies of Art and Culture

This metric is a re-elaboration of the 'Using external vocabulary' metric of (Färber et al., 2018). In our work, we focus on the use of vocabulary that belongs to taxonomies of art and culture, which play an important role in artwork descriptions as they provide permanent URIs for specific subjects, scenes, and other *icon* elements represented in artworks. Moreover, they are curated by domain experts, and referring to them gives more authoritativeness to the interpretations. For this analysis, we selected four core taxonomies: Iconclass [36], the Getty Art & Architecture thesaurus,[37] the Getty Iconography Authority vocabulary,[38] and the

[36] http://www.iconclass.org/help/outline

[37] https://www.getty.edu/research/tools/vocabularies/

[38] https://www.getty.edu/research/tools/vocabularies/guidelines/cona_3_6_3_subject_authority.html

Getty Cultural Object Name Authority vocabulary.[39] We measure the references to external taxonomies of art and culture by dividing the number of artworks in a knowledge graph that are associated with at least one of them by the total number of artworks present. Given A as the set of artworks and KG as the set of triples $\{t_1, ..., t_k\}$ in a knowledge graph (a triple being a sequence of subject, predicate, object $\{s_i, p_j, o_k\}$) and T as the set of nodes in a knowledge graph representing a particular subject expressed using a taxonomy of art and culture, we consider an artwork part of the subset A_t that contains artworks with a taxonomy reference if

$$\forall a_i \in A_t : a_i \in A \wedge (\exists p_j \exists o_k : (a_i, p_j, o_k) \in KG \wedge o_k \in T)$$

and we measure the References to external taxonomies of Art and Culture (RETAC) of a knowledge graph as

$$RETAC(kg) = \frac{n(A_t)}{n(A)}$$

The list of taxonomies of art and cultures used for this analysis contains only those that are referenced at least in one of the analyzed knowledge graphs. Increasing the number of taxonomies referenced would not change the evaluation methodology (and its formula). We welcome potential changes to this list to address *icon* aspects of more specific artworks, such as the reference to the Chinese Iconography Thesaurus[40] for a potential analysis on Chinese *icon* statements in the Semantic Web. References to external taxonomies are strictly related to iconography and iconology but are not essential to give a complete artwork description. For this reason, we weigh this metric 0.8.

The query shown in listing3.3 was used to count all the artworks in ArCo, referring to a taxonomy of art and culture (Iconclass).

Listing 3.3: SPARQL query launched on ArCo to count the artworks that have a reference to a taxonomy of art and culture (iconclass)

```
SELECT (COUNT (DISTINCT ?s) as ?tot)
WHERE {
?s a arco:HistoricOrArtisticProperty ; arco−cd:iconclassCode ?res}
```

3.6.3.7 Intralinking potential for subject comparisons

We introduce this metric to highlight the importance of intralinking subjects in the same knowledge graph. We consider 'intralinking' as the connection between entities belonging to the same knowledge graph. Having a URI as a subject of an artwork allows grouping artworks per subject and comparing them in respect to

[39] https://www.getty.edu/research/tools/vocabularies/cona/
[40] https://chineseiconography.org/

having a subject as a literal. Moreover, the same subject can then be aligned to other subjects in different knowledge graphs to foster interlinking in the digital art history LOD field. We measure the Intralinking potential for subject comparison by dividing the number of subjects that are linked to more than one artwork by the number of total subjects. Given S as the artistic subjects (expressed as URIs) in a knowledge graph and S_2 as the artistic subjects that are linked to more than two artworks, we measure the Intralinking potential for subject comparison (IPSC) of a knowledge graph as

$$IPSC(kg) = \frac{n(S_2)}{n(S)}$$

As this aspect is relevant but not fundamental for iconographical content representation, we weight it 0.6. Two example queries that count the number of subjects (URIs) in Europeana, and the number of subjects that are linked to more than one artwork can be seen respectively in listing3.4 and3.5.

Listing 3.4: SPARQL query launched on Europeana to count all the subjects that are URIs

```
SELECT (count(distinct ?sub) as ?tot)
WHERE {?s a skos:Concept ; skos:broader* aat:300033618 .
       ?s2 a skos:Concept ; skos:broader* aat:300047090 .
    {?CHO dc:type ?s ; dc:subject ?sub .
     FILTER (isURI(?sub))}
    UNION
    {?CHO dc:type ?s2 ; dc:subject ?sub .
     FILTER(isURI(?sub))}
}
```

Listing 3.5: SPARQL query launched on Europeana to count all the subjects that are linked to more than one artwork

```
SELECT (COUNT(?sub) as ?tot) WHERE { FILTER (?tot > 1){
SELECT ?sub (COUNT(DISTINCT ?CHO) as ?tot) WHERE {
        ?s a skos:Concept ; skos:broader* aat:300033618 .
        ?s2 a skos:Concept ; skos:broader* aat:300047090 .

{?CHO dc:type ?s; dc:subject ?sub .}
UNION
{?CHO dc:type ?s2; dc:subject ?sub .}
}
GROUP BY ?sub }
}
```

Table 3.6: Results for each metric over the selected Knowledge Graphs (as of 01/12/2022)

Short Name	SVIIT (weight 1)	IICC (weight 1)	IISG (weight 1)	IA (weight 0.6)	IPSC (weight 0.6)	RETAC (weight 0.8)
ArCo	0.8278	0.74	0.3333	0.0026	0.172	0.1238
Fondazione Zeri	0.9925	0.5117	0.1111	0.0005	0.266	0.5449
Nomisma	0.995	0.5	0.2222	0	0.749	0.0001
Wikidata	0.9768	0.74	0.6667	0.699	0.367	0.157
SARI	0.849	0.3783	0.1111	0.997	0.5	0
Europeana	0.4688	0.236	0.1111	0.0073	0.6122	1
ND_Hungary	0.13	0.5392	0.1111	0	0	0
DBpedia	0.655	0.7242	0.2222	0.994	0.41	0
Yago	0.99	0.4825	0.1111	1	0.1675	0

3.6.4 Results and discussion

Results obtained from applying the metrics over the KGs are summarized in Table 3.6 and visualized in Figure 3.2. To give a better overview of the results of the metric evaluation, they were then used to place the knowledge graphs inside a two-dimensional landscape. The landscape coordinates are determined by the two macro-aspects, namely content and structure, described in section 3.6.3. We averaged the metrics relative to these two macro-categories to obtain a score for content and structure. These averages are computed taking into consideration the weights of each metric. Given M_s and M_c as the sets of scores of a knowledge graph relative to its structure and content respectively {IISG, IA, RETAC, IPSC} and {SVIIT, IICC}, WM_s and WM_c as the sets of weights given to M_s and M_c respectively $\{w_{iisg}, w_{ia}, w_{retac}, w_{ipsc}\}$ and $\{w_{sviit}, w_{iicc}\}$ we computed the structure score (SS) of a knowledge graph as follows

$$SS(kg) = \frac{IISG \cdot w_{iisg} + IA \cdot w_{ia} + RETAC \cdot w_{retac} + IPSC \cdot w_{ipsc}}{\sum_{i \in WM_s} i}$$

and the content score (CS) of a knowledge graph as follows

$$CS(kg) = \frac{SVIIT \cdot w_{sviit} + IICC \cdot w_{iicc}}{\sum_{i \in WM_c} i}$$

We divided the graphs into four categories that represent the four quadrants of the landscape, according to their averaged scores, namely: high in content and structure (both scores ≥ 0.5), low in content and high in structure (content $<$

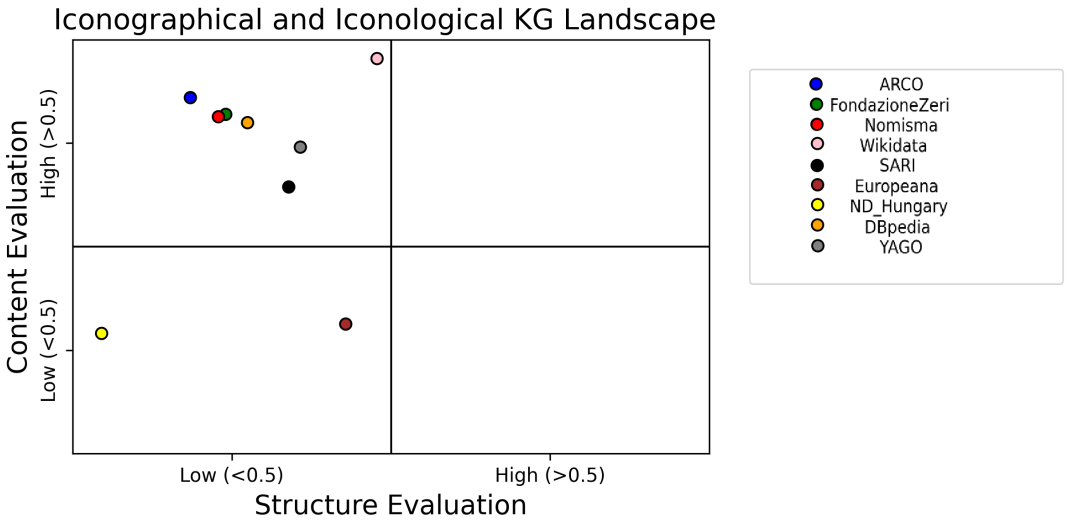

Figure 3.2: Landscape of the knowledge graphs on their iconographical and iconological statements and the structure of the schemas that describe them.

0.5 and structure ≥ 0.5), high in content and low in structure (content ≥ 0.5 and structure < 0.5), low in content and in structure (both scores < 0.5).

Figure 3.2 shows a clear scenario: the data content is generally correct but not thoroughly described. In fact, none of the graphs has acceptable results in the structure quadrants, and most (7 out of 9) present high scores in content. Nevertheless, this result is given by higher rates in semantic validity (six KGs score more than 0.8) rather than in column completeness (only 3 KGs score more than 0.7). Among them, despite being a general-purpose graph, Wikidata performs the best results. It has the best schema granularity, as several properties can be aligned to the prototype schema of Figure 3.1. In addition, its column completeness scores are higher than some Art History graphs. This is because, in contrast with the approach adopted in the other graphs, the first level of interpretation is often described even when a second or third-level subject is identified.

The granularity in the levels' description may influence the intralinking metric since the description of simpler and more generalizable elements of the first level of description can positively affect the capability of comparing artworks that share them. This assumption is evidenced by the fact that graphs such as SARI's platform,[41] where the subjects considered are broad concepts (e.g. 'persons related to art'), perform better results in intralinking. However, it is important to underline that the general purpose of the graph and the restricted number of subjects described can affect this evaluation. For example, Nomisma,[42] having as subjects only deities, personifications, or Roman emperors, scored the highest in this metric.

[41] https://rds.swissartresearch.net/resource/rdsPages:Start
[42] http://nomisma.org/

Other relevant qualitative observations can be made over the results obtained. Firstly, we envision that Art History KGs such as Zeri&Lode, which precisely identify second-level subjects with an acceptable percentage of interlinking to vocabularies, could foster subject retrieval and semantic computational capabilities by adding information on more levels of interpretation. Additionally, ArCo, created by automatic conversion of cultural heritage catalogs, despite having a high result in column completeness, has low rates in subject intralinking (0.172) and in relation to external taxonomies (0.123). This may be due to the highly automatic process through which the knowledge graph was created (V. A. Carriero et al., 2019). The automatic creation of URIs for subjects from strings extracted from catalog data could be improved to avoid duplicates of URIs referring to the same entities, therefore increasing the intralinking potential of the KG. For what concerns references to external taxonomies, Europeana shows the best results. In fact, it is possible to retrieve different types of artworks according to the Getty vocabulary category, allowing feasible reusability and retrieval of information for people knowledgeable about them. Moreover, by defining artwork types in this way, it is also possible to retrieve information without having to know specific classes for types of artworks, shifting from the necessity to know the specific schema of the knowledge graphs to the knowledge of general taxonomies applicable to different linked open data datasets. It is interesting to note that, despite having a perfect score in references to taxonomies of art and culture, Europeana does not have any specific property that links an artwork to a taxonomy (it uses `dc:subject`), which decreased the score obtained in the schema granularity metric. Finally, the National Data Archive of Hungary (Fülöp et al., 2005) scores worst in the general categories, given the absence of subjects expressed as URIs, the only use of `dc:subject` to describe *icon* statements, and the complete absence of references to taxonomies.

To exploit the capabilities of interlinking, inference, and analysis of the semantic technologies applied to *icon* study of artworks, reliable, complete, and well-structured data are required. We assess the data quality of current CH KGs that are openly available, online queryable, and have data on artwork subject descriptions. Our results indicate that only a few KGs describe the artwork's iconography and iconology (Section 3.6.2). To assess their content according to different aspects, we adapt five metrics from prior KG evaluation methodologies (Section 3.6.3), and add a new metric. This set of metrics is used to evaluate the content and the structure of sub-graphs describing artworks' *icon* characteristics. We observe that all KGs perform poorly in the *schema structure* as resulting from a combination of metrics, but the major part of them have high or acceptable scores for the *content evaluation* combined metric (Section 3.6.4). This survey gives a

critical overview of the complexity involved in the correct and exhaustive creation of domain-specific data. Since the artwork *icon* descriptions are generally correct, the current data can be reliable for data reuse and analysis. Nevertheless, a deeper, more accurate description and a better schema are required to enhance all the expressivity that may lay in them. Whereas *icon* descriptions exist, they are not sufficiently interlinked, searchable, and exhaustively described. As a consequence, we recommend i) a more extended reuse of existing domain-specific controlled vocabularies; ii) the development of domain-specific ontologies that thoroughly cover iconography and iconology; and as a result of this, iii) either the creation of new domain data, formally expressed at a finer granularity or the re-engineering of current data following newly developed ontologies. This recommendation is extended to current studies in enhancing iconographical cultural metadata, such as (Bobasheva et al., 2021), which focus on adding new knowledge to artistic linked open data. As shown in this study, the quantity and correctness of the data cover only one side of the coin. It is also important to express the newly generated knowledge with the correct schema that respects the granularity and complexity of iconography and iconology. Finally, from the general perspective of data quality assessment in a specific domain of knowledge, this evaluation can be considered a case-study, which can be generalized for spotting semantic representation issues in other domains.

Part II

SEMANTIC WEB AND ICONOLOGICAL INTERPRETATIONS

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. As a result of the survey on the availability of iconographical and iconological descriptions in current KGs, 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 iconographical and iconological interpretations represent the domain features to foster the access, analysis,

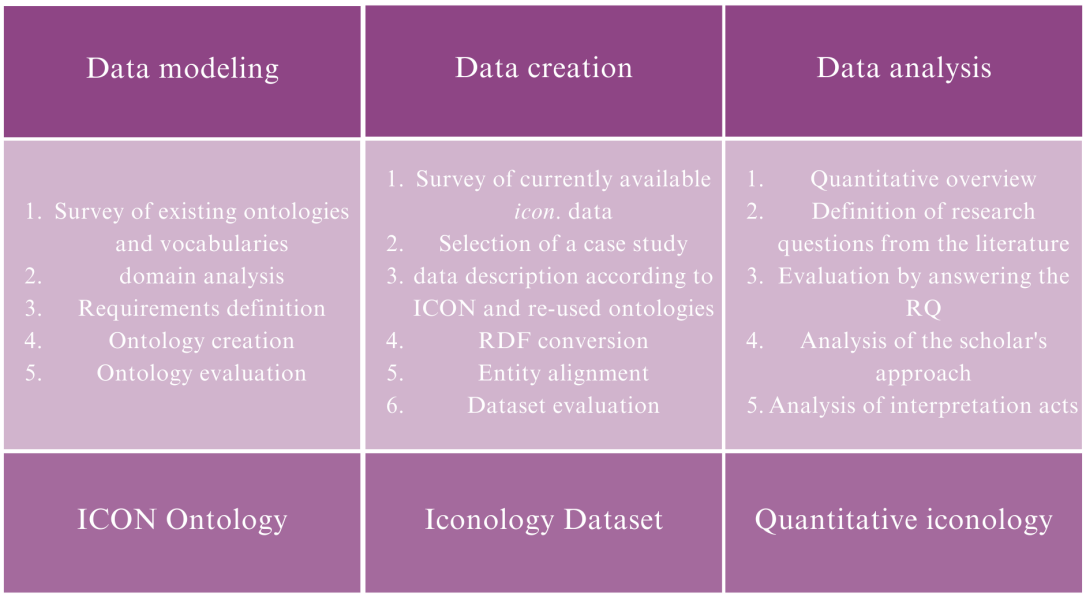

Figure 4.1: Approach overview

and retrieval of iconographical and iconological content?

- MQ2. What would be the advantages of quantitatively analyzing a semantic network of iconographical and iconological interpretations?

The first research question (MQ1) concerns the study of the domain (RQ1) and the ontological modeling (RQ 2). 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 data set creation, and the analysis (see Figure 4.2).

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:

- RQ 2.1: What theoretical models are available in the domain-specific literature for describing the discipline? Among them, which one is feasible for ontological modeling?
- RQ 2.2: 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 (RQ 3) and analysis (RQ 4). In both cases, sub-questions related to the adopted theoretical approach and research questions are defined:

- RQ3. How to create data sources allowing quantitatively analyzing iconographical and iconological 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

4.2 Approach to the research

To answer the research questions, we first conducted a study of the typological characteristics of the iconological field of study, then developed feasible ontological modeling, and manually collected and described data according to it. Next, quantitative analyses were conducted.

The approach is structured into three main parts: 1) domain study and ontological modeling, 2) data creation, and 3) data analysis.

4.2.1 Domain study

The domain study was conducted to understand the characteristics and interests of an iconological approach by identifying 1) typological characteristics of the field, and 2) domain-specific research questions addressed by the discipline. The former is based on a top-down approach, selecting the most feasible theoretical model available. Both definitions will be fundamental for shaping the remaining parts of the research. In particular, whereas the definition of the characteristics of the domain will be fundamental for ontological modeling (Chapter 5), the definition of domain interests is crucial for designing the analysis performed (Chapter 7).

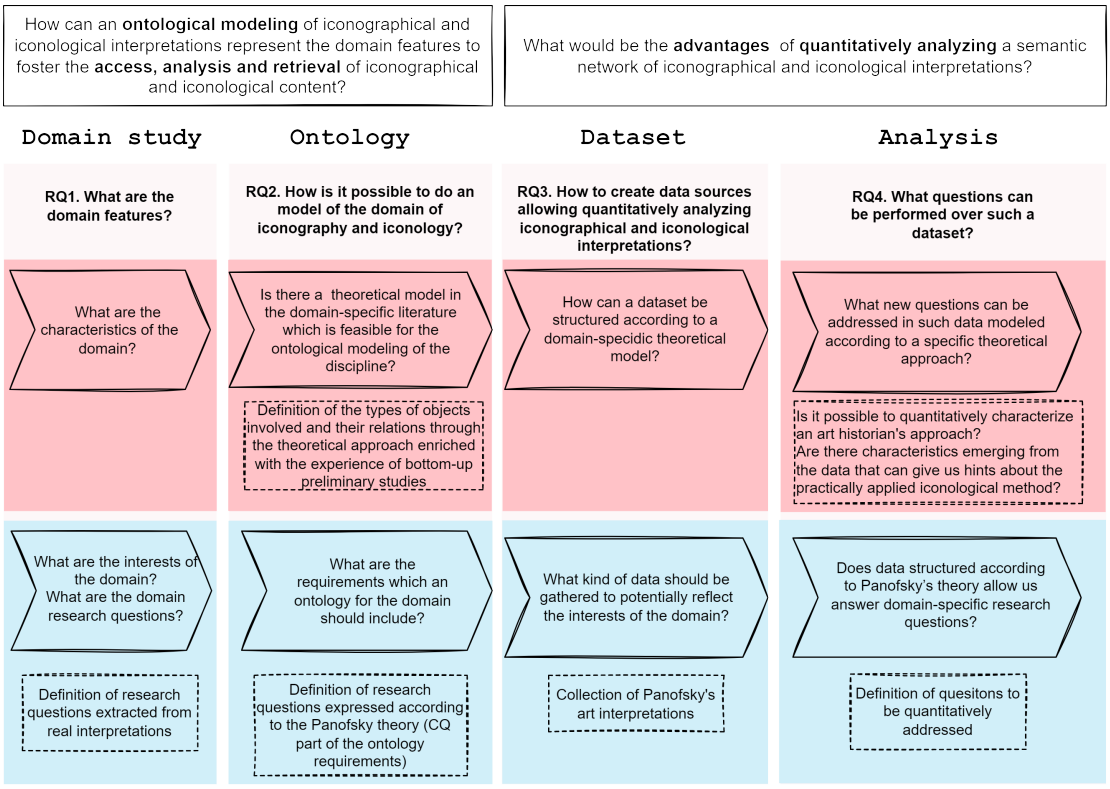

Figure 4.2: How the research questions are related

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 while doing an *icon* interpretation. To this aim, we extract the research questions that researchers are implicitly answering while performing an interpretation with a bottom-up approach from a selection of studies by various scholars.

The approach of specifying domain requirements through research questions is presented by Szabo (2012), one of the first DAH studies in which databases were

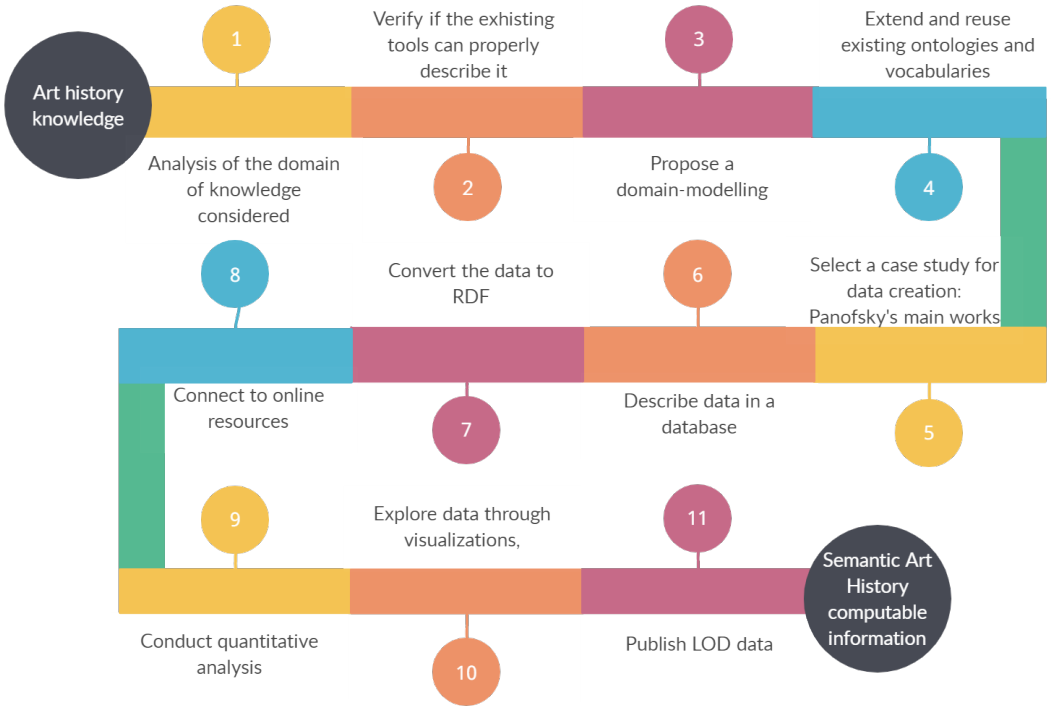

Figure 4.3: Timeline

created for domain-specific data-driven inquiries into artworks study.

The extracted research questions were thematically grouped. As they express complex issues addressed by scholars, they are also used as part of the dataset evaluation. They will be taken into account in the following steps of the approach. During the ontology development, they will be expressed according to the Panofsky theoretical approach to define the Competency Questions. Next, they will be taken into account in the database creation phase. Finally, they will be expressed in subquestions to be quantitatively addressed as part of the data analysis and the dataset evaluation.

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)[1]. 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 method-

[1]The current section is based on the version published in Sartini et al. (2023)
[2]The ontology is available at https://w3id.org/icon/ontology/

ology 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 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). In fact, the methodology tries to solve the 'local problems' included in the so-called 'local space', or the modeling issues related to the specific ontology that is being developed, with the re-use of modeling patterns that come from the 'solution space', such as the ODP. We specifically adopted this methodology when dealing with the re-use of ontology design patterns that were specialized in the context of our domain.

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. We present our alignments and reuse following guidelines proposed by the state of the art (V. Carriero et al., 2020; Osman et al., 2021). Our ontology selection for reuse and alignment 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. With this term, we refer to a data source modeled according to a domain ontology 'established to capture the domain of interest' that defines them as 'semantically interrelated entities and relations' (Abu-Salih, 2021).

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 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 fact allows us not only to perform quantitative analysis on authoritative domain-specific data but also to verify if such a 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 data sets is based on the Best Practices for the Publishing of Open Linked Data published by the W3C[3].

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 the best of the author's knowledge, the currently available graph and semantic data storage platforms (e.g., Omeka-S[4], Neo4j[5]) are complex to customize for storing data with such a level of

[3] https://www.w3.org/TR/2014/NOTE-ld-bp-20140109/#MODEL
[4] https://omeka.org/
[5] https://neo4j.com/

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 4.5 and 4.3, 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 re-defined and enriched 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).

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 The database creation was conducted as follows. Firstly, we created a tabular data structure reflecting the ontology structure, including information on each level of interpretation. The entities to be described reflect the ICON ontology classes. Further relevant aspects (e.g. metadata about artworks, people, books, places) were included in the database design. The artworks described in the preliminary study (Baroncini et al., 2021) were used as a test for database design and improvement.

Database population Secondly, we interpreted Panofsky's claims and described them according to the data structure. The quality of the data was assessed through value validation through 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 a 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 through a Python script.

4.2.3.3 Dataset creation

Ontological modeling The 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 data set is available on a GitHub repository. [6] 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. 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, so as to include new types of entities encountered.

Reused ontologies and vocabularies Following the good practice of ontology reuse (V. Carriero et al., 2020), 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[7] (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. Direct reuse or alignment was also performed at the vocabulary level. Regarding types of artifacts, style, and period terms, Getty vocabulary AAT URIs were directly reused. An internal URI was created and aligned with one or multiple reconciled terms of controlled vocabularies for all the other cases. Places and people names were aligned respectively to TGN or ULAN and VIAF. The artworks were partially aligned with Wikidata. Regarding the depicted subjects, alignments were made 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[8], 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.[9] The results were manually checked and included in the main database to be included through a new RDF conversion.

Figure 4.4 illustrates the general workflow adopted for alignment. The vo-

[6] https://github.com/SofiBar/IconologyDataset/

[7] http://www.sparontologies.net/ontologies/pro

[8] https://www.getty.edu/research/tools/vocabularies/guidelines/cona_3_5_associative_rels-copy.html

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

cabularies for the reconciliation were identified, and an alignment strategy (i.e., either a direct reuse of the URI or an alignment to the external source through `owl:sameAs`) was chosen. When feasible, a script for automatic reconciliation was used and results were manually checked.

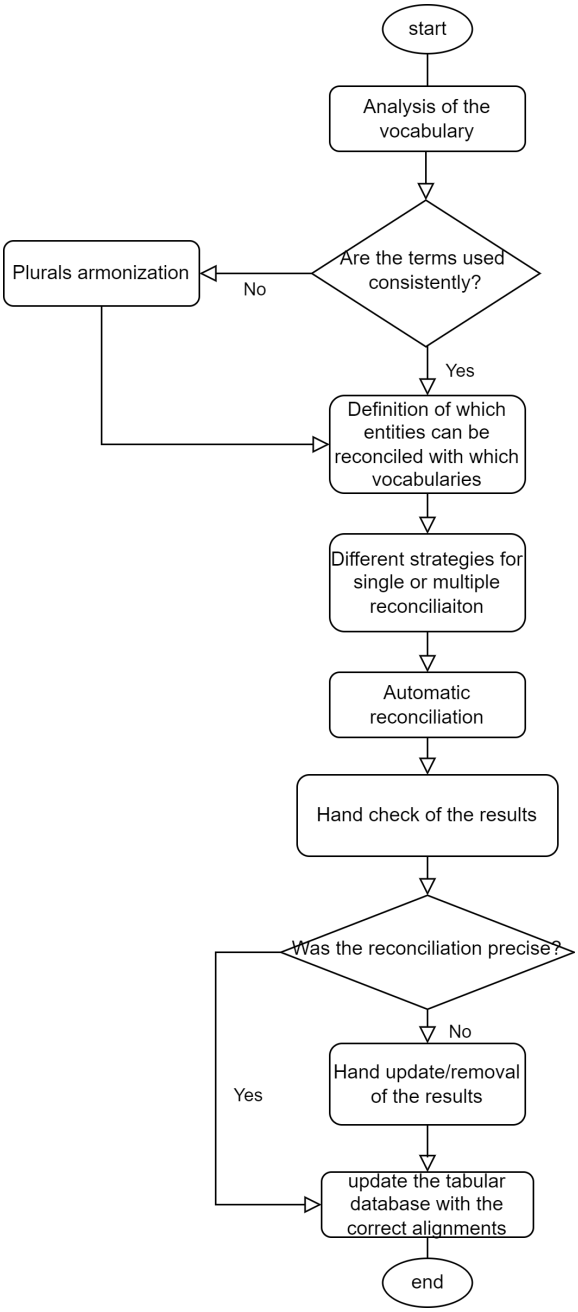

Figure 4.4: Workflow adopted for reconciliation

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.

Data conversion Data were converted to RDF format with a Python script using RDFlib[10], a library for creating and querying RDF data. In the script, the database columns were mapped to the previously defined 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 NC SA) is included in the graph. Further metadata are provided, including author, source, language, and versioning. We reused Dublin Core Terms[11] and the PROV-O ontology[12] for this purpose.

Machine access to data Data are exposed through a SPARQL endpoint[13] and provided with a unique identifier[14].

4.2.4 Ontology and data evaluation

The dataset was evaluated 1) 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 according to SHACL, which is, with ShEx, one of the languages aiming 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[15].

To verify if the dataset could address the domain-specific questions defined in the domain study, we defined subquestions that would be directly expressed in SPARQL queries and we performed them over the dataset. We calculate

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

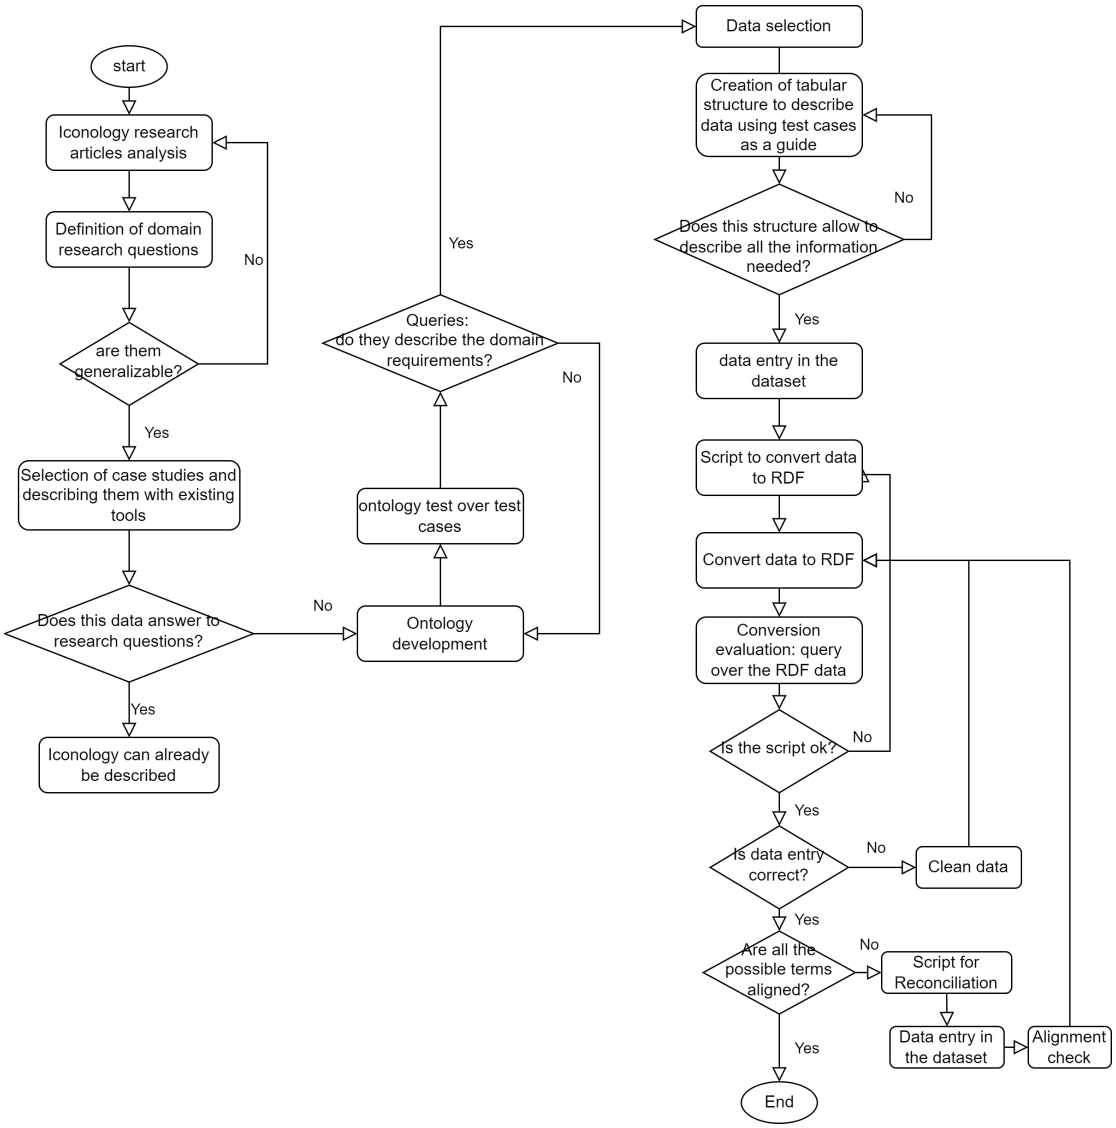

Figure 4.5: Workflow of ontological modeling and data creation

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 data set 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.

The ontology was validated according to multiple dimensions. Its ability to represent and answer aspects of interest of the iconography and iconology domains was assessed by performing the competency questions expressed in the ontology requirements over a subset of the dataset. Then, a comparison with existing

ontologies was conducted through the modeling of the same case study to assess the introduced novelties. Finally, an automatic evaluation 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 was conducted over the above-described RDF dataset of iconological interpretations manually extracted from a selection of Panofsky's books and modeled according to ICON.

The analysis is structured in two main parts. We assess whether iconographical and iconological 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) which further 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 domain-specific 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 (i.e., the percentage of successfully addressed questions) is presented in chapter 8, chapter 7 thoroughly illustrates the performed queries and their results.

4.2.5.2 New inquiries

In this section, we expressed hypotheses about what more questions could be addressed in the resulting data, which could not be answered with the traditional analysis. These research questions are addressed in the analysis part in Section 7.3.

As the data set 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 formalisation 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(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 iconographical and iconological analysis is underlined by Taylor(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)

RQ4.1a2 aims at representing 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)

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.

CHAPTER 5

Domain study and the ICON Ontology

5.1 Domain study

From the preliminary study in (Baroncini et al., 2021), it is clearly evident that a simple description of the subject matter of the artwork is not enough to cover the complexity afforded by iconological studies. The aspects of complexity involved range from deeper levels of comprehension in the artwork itself to intricate relations with textual sources or comparisons with artworks depicting the same subject matter or similar features. In fact, the understanding of deeper psychological traits of the artist or the contemporaneous socio-cultural context lies in the analysis of small details and differences, as Warburg's approach points out.

For a better understanding of the domain, the following sections provide the definition of its characteristics (Section 5.1.1) and interests (Section 5.1.2).

5.1.1 Characteristics of the domain

To the characteristics of the domain, we adopt a top-down approach from the theory of Erwin Panofsky. His systematic definition of the interpretation act is of great value for the attempt at formalization we are conducting. Nevertheless, we take into account the contributions of other art historians for important aspects of the discipline that are not included in Panofsky's theory. The initial research question (Q1.1.) can therefore be rephrased in this way: *According to Panofsky's theory, what are the items involved in an interpretation act? How do they relate to each other? How can such a theory be integrated with the relevant aspects addressed by other art historians?*

To answer this question, we examine Panofsky's theory and extract the types of

entities involved along with their relations. Entities are listed with their definition in Tables 5.1 5.2 5.3, while relations are represented in Figure 5.1.

The content of the essay on which we will focus was already treated in Chapter 2, as it corresponds to the essay introducing Iconography and Iconology, first published in *Studies in Iconology* (1939), and then included in *Meaning in the Visual Arts* (1955). In this section, we examine in detail the entities that the art historian describes with particular reference to the terminology he uses to identify them and the description he gives.

According to the definition provided by Panofsky in this treatise, Iconology is a branch of Art History focusing on the study of **subject matter or meaning** of visual arts. In Panofsky's theory, this concept is central to such an extent that he brings all the art understanding process to the identification of meaning, divided into three forms. Hence, he recognizes the existence of two opposed types of meanings, namely *phenomenical* and *essential*. The former is the type to which the *primary or natural meaning* and the *secondary or conventional* belong, whereas the latter is the type of the *intrinsic meaning or content*.

The three-strata nature of the interpretation act introduced in section 2.2.2 addresses the understanding of such meanings in artworks, each according to the level. In the following paragraphs, we report on the definition of such meanings for each level and on the items needed to interpret them.

The **primary or natural meaning** includes the *factual* and *expressional* meanings and is apprehended during the pre-iconographical description of a work of art (level 1; see Table 5.1). Such meaning is defined as of 'elementary and easily understandable nature' and is understood by relating 'certain visible forms with certain objects known [...] from practical experience' (Panofsky, 1955, p. 26). Whereas factual meanings are acquired by simply identifying visual forms with known objects, the comprehension of expressional meanings (i.e., the emotions that the recognized factual meanings may convey) takes place through empathy.

When interpreting a work of art, the visible forms that will be interpreted are the configurations of lines and colors visible on the artwork surface. Once a visual form is identified with the primary meaning and, therefore, recognized as a carrier of meaning, Panofsky calls it an *artistic motif*. A preiconographical description is consequently defined as an enumeration of artistic motifs[1].

To consider the interpretation correct, some background knowledge is needed. Although Panofsky initially affirms that the sole common practical experience is required, he underlines that, to correctly interpret the forms visible in a work of art, the knowledge of stylistic conventions is required, namely 'the manner

[1]'An enumeration of these motifs would be a pre-iconographical description of the work of art' (Panofsky, 1955, p. 28)

in which objects and events are expressed by forms under varying historical conditions'(Panofsky, 1955, p. 35).

We further introduce another aspect not addressed by Panofsky in his theory, viz., the visual arrangement of subjects, that we consider relevant to the description of this level, as suggested by Imdahl (2012), who underlines the relevance of the iconic language of the image (see Section 2.2.3). Although Panofsky did not explicitly refer to it, he considers visual aspects core features of artworks understanding in some iconological interpretations (e.g., his essay on perspective as a symbolic form of different cultures (Panofsky, 1955)).

Table 5.1: Definition of the entities involved in the pre-iconographical description according to Panofsky's theory (integrated with Imdahl (2012))

Entity	**Description**
Primary or natural meanings	Meanings of 'elementary and easily understandable nature'. They are recognized by relating the knowledge of objects and emotions acquired through practical experience to certain visual forms (Panofsky, 1955, p. 26). Natural meanings can be *factual* or *expressional*
Factual meanings	Primary or natural meanings that are apprehended by identifying visible forms with objects known from practical experience
Expressional meanings	Primary or natural meanings apprehended by empathy
Pure forms	Certain configurations of line and color, or certain peculiarly shaped lumps of bronze or stone
Artistic Motifs	Pure forms recognized as carriers of primary or natural meanings
Pre-iconographical description	An enumeration of the artistic motifs
Compositions	Combination of artistic motifs
Practical experience	The interpreter's 'everyday familiarity with objects and events' (Panofsky, 1955, p. 27)
History of style	'The manner in which objects and events are expressed by forms under varying historical conditions' (Panofsky, 1955, p. 35)
Visual arrangement	Visual arrangements that may exist between the depicted subjects (e.g., pyramidal structure)

A first attempt to formalize the recognition process of natural meanings is provided in Figure 5.1(a). During this act, the interpreter relates the pure forms with the natural meaning, obtaining artistic motifs. He/she makes the interpretation on the basis of his/her practical experience and the history of style.

The second level of meaning understanding concerns the second type of *phe-*

nomenical meaning, namely the **secondary or conventional** meaning (see Table 5.2). Whereas the natural meaning was perceived through the senses, this meaning is intelligible, as it is expressed on the basis of human conventions (e.g., the Western world's conventional meaning of greeting someone by removing the hat). In addition, the secondary meaning is intentionally conveyed by the person performing it (e.g., an artist depicting a Virtue intentionally depicts it according to certain visual conventions to make it recognizable), contrary to the natural meaning, which can be non-intentional. Therefore, secondary meaning consists of specific themes or concepts that are voluntarily conveyed in visual works.

While interpreting the artwork, the observer will recognize the artistic motifs apprehended at the previous level as carriers of such secondary meanings; e.g., a visual form representing a man is recognized as Saint Bartholomew. Such recognized motifs or compositions of motifs are defined by Panofsky *images*. When images carry abstract concepts, they are called symbols or personifications. Combinations of images are *invenzioni*, which is the term used by ancient theorists to define stories and allegories. Usually, combinations of personifications and symbols can be called allegories.

Table 5.2: Definition of the entities involved in the iconographical description according to Panofsky's theory

Entity	**Description**
Secondary or conventional subject matter	Themes or concepts
Images	Motifs as recognized as carriers of secondary meanings
Personifications or symbols	Images conveying the idea of abstract and general notions
Invenzioni (Stories and Allegories)	combination of images
Allegories	combinations of personifications and/or symbols
History of types	'The manner in which, under varying historical conditions, specific themes or concepts were expressed by objects and events'

As the secondary meaning recognized in this level of interpretation depends on human traditions, the prerequisite for such an understanding is 'a familiarity with specific themes or concepts as transmitted through literary sources' (Panofsky, 1955, p. 35). Going further, the correctness of recognition is given by the knowledge of how such themes and concepts were visually conveyed through the depicted objects and events (Panofsky, 1955, p. 37). Knowing how themes and concepts expressed in literary sources were usually depicted (e.g., with what attributes) helps in the disambiguation of uncertain subjects that may be potentially referred

to by multiple textual sources.

The relations among the entities participating in the second level of interpretation have a similar structure to that provided for level 1 (Figure 5.1(b)). The interpreter relates the artistic motifs or compositions recognized during the pre-iconographical description with conventional meanings, obtaining images that may be grouped into *invenzioni*. The interpretation is conducted on the basis of the interpreter's knowledge of concepts and themes from literary sources and of the history of iconographical types.

Table 5.3: Definition of the entities involved in the iconological synthesis according to Panofsky's theory

Entity	**Description**
Intrinsic meaning or content	'a unifying principle which underlies and explains both the visible event and its intelligible significance, and which determines even the form in which the visible event takes shape'(Panofsky, 1955, p. 28)
History of cultural symptoms	'the manner in which, under varying historical conditions, the general and essential tendencies of the human mind were expressed by specific themes and concepts' (Panofsky, 1955, p. 39)
Deeper meanings	meanings intentionally embedded by the artist (e.g., a moral message)

The last type of meaning corresponds to the essential one, recognized during the iconological interpretation (third level). As reported in Table 5.3, Panofsky defines the **intrinsic meaning or content** as the underlying principles manifested through artistic motifs, compositions, and images. Some examples of such principles are the attitude of a nation, a period, a class, or a religious or philosophical persuasion. In other words, the artwork is interpreted as a document of cultural and societal phenomena, as it is a symbol of such phenomena, which manifest themselves in the compositional and iconographical aspects of the artwork. We further partially include van Straten's position, which defines deeper meanings as those meanings voluntarily embedded by the artist in the artwork (van Straten, 2012).[2]

The iconological interpretation occurs through a *synthetic intuition*, an interpretative process 'based on synthesis rather than analysis.' To be correct, the interpreter must know how such socio-cultural phenomena were expressed by themes and concepts, namely the iconographical subjects identified during the iconographical description level.

Similarly to the previous levels, the relationships established by the entities that

[2]His definition is only partially included as, in contrast with Panofsky, van Straten recognizes the conceptual meanings of personifications and symbols as deeper meanings.

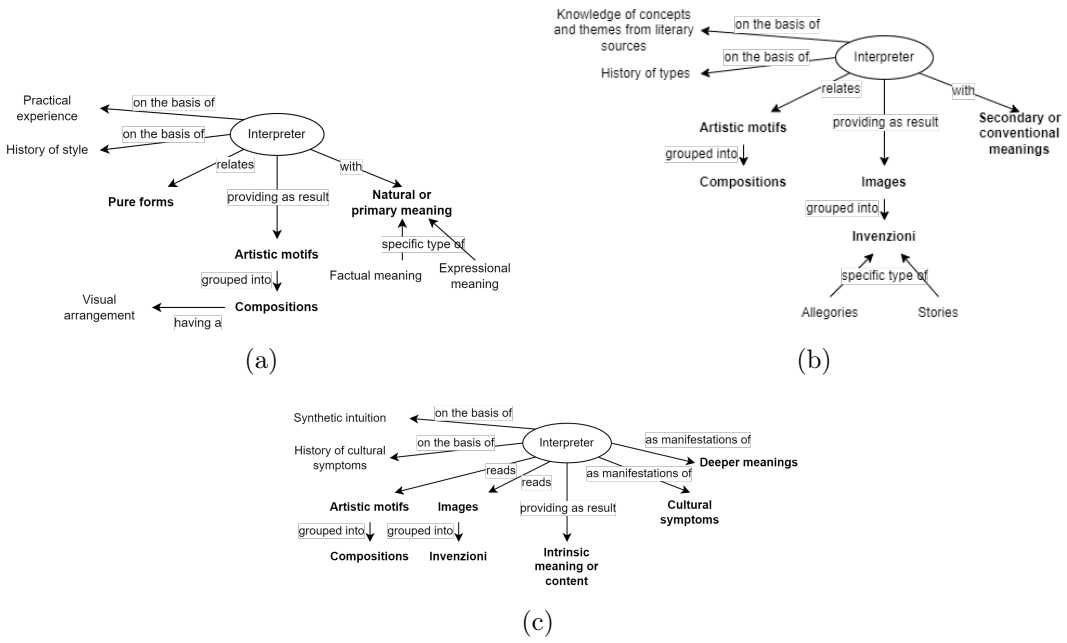

Figure 5.1: Definition of the characteristics of the domain: how the entities involved in each level of interpretation relate to each other

occur in the iconological synthesis can be described as follows (Figure 5.1(c)). The interpreter reads artistic motifs and images as manifestations of cultural symptoms, providing, as a result of the interpretation, the intrinsic meaning. Interpretation is carried out on the basis of synthetic intuition and knowledge of the history of cultural symptoms.

5.1.2 Interests of the domain

To better address possible points of interest, we extracted the questions the scholars implicitly answered in the typological selection of the literature collected in this study. The research question driving this methodology was to identify what interested the iconography and iconology domain (RQ1.2). As a result, we identified 12 questions.

The following paragraphs briefly describe the studies from which the question was extracted. Since the same example can implicitly answer several identified questions, some of them are repeated. As some questions are central to the discipline and are found in many texts, we describe only one or a few cases used as examples. The questions will be conducted on the dataset created during this PhD project as part of the data analysis in Chapter 7.

5.1.2.1 Cultural phenomena

Q1. Which cultural phenomena are witnessed by artworks?

Question 1 constitutes the central question of iconological studies. As representatives of it, we cite Warburg's investigation examining the role of portraits in Renaissance Florence.

In the study *The Art of Portraiture and the Florentine Bourgeoisie* (Warburg, 1999), Warburg delves into the question of the 'effect of *milieu* on the artist', i.e., how the contemporary background influences the artworks (Warburg, 1999, p. 187). He noticed that while during the Middle Ages contemporary human beings were not depicted in the same illusionistic space of a sacred representation, during the Renaissance, there was an increasing number of representations of terrestrial patrons included in such sacred space. In particular, he analyzes the painting in the Sassetti Chapel in Santa Trinita in Florence, realized by the painter Domenico Ghirlandaio. The fresco representing the scene of the *Confirmation of the rule of Saint Francis* portrays not only the Patrons and their family members but also Lorenzo de' Medici, the most influential man in Florence of the time. To explain the massive presence of portraits not limited to the ones of the sole donors and included in the sacred space, he reconstructed the social relations of the family with the Medici and the contemporary function of portraits as votive pictures. Through the aid of 'archival and literary searches'(Warburg, 1999, p. 187), he retraces the dense network of relations between the two families. By analyzing the social value that the portrait had for the well-documented Medici family, who used to place real-size portraits of family members in the churches for votary purposes, he concludes that the massive presence of patrons' portraits in a religious art piece is not, as may appear, a violation of the sacred scene, but the proof of the evolution of the votive function of patrons' portraits.

Q2. Which sources and visual aspects motivate their presence?

It seems frequent that recognizing a cultural phenomenon is based on visual or textual sources that lead to a deeper understanding of the context.

Considering Warburg's study of Ghirlandaio's painting described above, he based the interpretation on several sources. In addition to the visual evidence (i.e., Lorenzo's portrait), the letters demonstrated Sassetti's narrow connection with the Medici family. On this basis, Warburg deepened the study of the role of votary portraits for the Medici family, finding further documentary sources about Lorenzo's real-size portraits, which were displayed in three Florentine churches. This textual evidence led to an understanding of the cultural practice of the portrait as a votary function described above. Finally, such background knowledge allowed a deeper understanding of the phenomenon witnessed in the Ghirlandaio fresco. From this example, it emerges that both visual and textual pieces of evidence are

fundamental for reconstructing the sociocultural context in which the artwork was realized to understand its deeper meanings.

Another relevant example of the use of visual citations can be found in the study by Cardini of the Palazzo Vecchio's Medici Chapel. The iconography of the Magi is historically narrowly linked to the concept of sacralization of the ruler's power (Cardini, 2001, pp. 22-23). The identification of the Magi as the Byzantine Emperor John VIII Palaiologos (adult Mago), of Sigismund of Luxemburg, Holy Roman Emperor from 1433 to 1437 (elder Mago), accompanied by a young Lorenzo de' Medici (young Mago), led to the overall interpretation of the chapel as a political manifesto of the future magnificence of the Medici family. The recognition of the portrait of Sigismund of Luxemburg, which allows this interpretation, is based on a comparison with the portrait on a medal (Cardini, 2001, p.31). In fact, the young descendant of the family, depicted as the young Mago among the rulers of the Western and East Roman empires, is symbolically represented as the future ruler of the known world.

Q3. How does the representational evolution of subjects witness the emergence of cultural changes?

In studying the reception of classical antiquity in the Middle Ages, Panofsky and Saxl (1933) reconstruct the survival of classical themes and their classical visual appearance. They observed that classical subjects were depicted in an utterly non-classical manner. According to them, the reasons are twofold. On the one hand, the classical themes were transmitted through a textual tradition so that the subjects were provided without the classical appearance they had in classical art. On the other hand, the Gothic style reached a full maturity that prevented the classical themes from being depicted with a style far from the sensibility of the contemporaries. Therefore, how themes are transmitted and reinterpreted expresses the cultures that assimilate them.

5.1.2.2 Symbols

Q4. How do symbols evolve?

Wittkower provides a vast reconstruction of how the symbol of the fight between an eagle and a serpent, first found in Babylonian culture, spreads over time and place, passing through a wide range of cultures and eras, reconstructing visual and meaning variations (Wittkower, 1977). We summarize some salient passages to provide an idea of the complex and detailed reconstruction that Wittkower performs of the migration of such a symbol. According to him, the earlier evidence of a bird and a snake was found in Babylon, from which the iconography of a bird

used as a solar symbol spread, reaching the East. In the Indian literature, it can be traced to the bird Garuda fighting with the snake Naga. Evidence can be found in visual representations and myths in Polinesia, Siberia, and German mythology. Furthermore, the hypothesis of migration in American cultures is affirmed since a predatory bird is present in Siberia and Alaska, analogously to the Arabian-Persian Roc, and a bird fighting with a snake is present in Mexico and Peru. Wittkower proceeds the inquiry by tracing the evolution of the symbol during the Roman, Medieval, Renaissance, and Baroque periods.

Another complex example is found in the study of Titian's *Allegory of Prudence* by Panofsky (1955), in which he traces the evolution and resemantization of the three-headed animal figures, which found their origin in the companion of the Egyptian deity Serapis. The creature, represented as three animal heads and a twisted serpent, was interpreted by the Roman writer Macrobius as a symbol of time, as the serpent was a traditional symbol of time. In the following tradition, the creature is associated with Apollo (Petrarch, *Africa*), depicted as a monster with the body of a snake, to finally become an independent symbol of time or prudence with an appearance closer to its original traits.

Q5. Is the symbolic meaning motivated by a specific source?

It is common for traditional attributes to originate in known textual sources. For example, the major part of catholic saints' attributes is the object of their martyrdom (e.g., Saint Sebastian is usually depicted as a man pierced by arrows). Consequently, it is even more reasonable that an art historian will check the literature presumably known by the artist when faced with unusual iconography. Lorenzo Lotto's painting *Venus and Cupid* (Figure 10.1), showing the two deities in a clearly symbolic scene, requires a similar investigation. To interpret it, the art historian Keith Christiansen provides citations from the classical literature explaining the possible symbolical meanings of the depicted objects. For example, the red cloth in the background of the painting is read as the proper decoration of a nuptial chamber (Claudianus, *The Magnate*). The incense had a similar function, according to Sidonius (*Epithalamium*).

Q6. Is the deeper conceptual meaning motivated by sources?

In his well-known essay on Botticelli's *Primavera*, Gombrich extensively uses textual sources to interpret the painting through contemporary Neoplatonic ideas (Gombrich, 1972). According to Vasari, Botticelli's painting was located in the villa at Castello that was built for the adolescent Lorenzo di Pierfrancesco de' Medici, cousin of Lorenzo il Magnifico. Gombrich examines some letters written by the Neoplatonic philosopher Marsilio Ficino to *Laurentius Minor*. In one of the letters,

Marsilio illustrates Venus as a symbol of *Humanitas* with a moral and educational aim toward the boy. In Gombrich's hypothesis, the letter is addressed to Lorenzo di Pierfrancesco, and he supposes that such a concept inspired the realization of *Primavera* as a decoration for his villa and that Ficino and the tutors of the young Medici drew up the iconographical program for the painting. Therefore, Gombrich's hypothesis is that the painting expresses a moral and educational meaning conveyed by an iconographical program imbued with Neoplatonic concepts and ideas.

5.1.2.3 Iconographies

Q7. How do iconographies, their meaning, and attributes evolve over time?

Panofsky reconstructs the genesis of the iconography of Father Time, which, during Baroque and Renaissance, was depicted as an old, winged man with the possible attributes of scythe, hourglass, crutches, or the snake biting its tail. However, the iconographies of the two forms of time in classical culture, namely *aion* and *chairos*, are never depicted with a similar appearance. In fact, the Renaissance iconography of time is a case of *pseudomorphosis*, according to which figures with a classicizing appearance are 'invested with a meaning that, for all their classicizing appearance, had not been present in their classical prototypes' (Panofsky, 1972, pp. 70-71). The iconography of Time derives the core traits of its aspect from Saturn, the god of agriculture, from which the old age and the sickle come. The identification of time with Saturn was favored by the similarity of their Greek names (the god Kronos and *chronos*, the word indicating time), supported by Plutarcus' indication that Kronos means time and by the Neoplatonic correlation of Kronos with Nous, the cosmic mind. Being invested with the meaning of time, Saturn acquires attributes indicating it, such as the serpent biting its tail, and the mythological episode of Saturn devouring his children is allegorically interpreted as time devouring everything. In the late Middle Ages, the sinister character of Saturn prevails, especially in his role as planetary deity. The Renaissance figure of Father-Time is the fusion of such a negative representation of Saturn with a scholastic representation of time (*Temps*), which appeared in the illustrations of Petrarch's *Triumphs* for the first time.

This type of analysis investigates the roots from which the iconography originated. In doing so, not only are its visual representations considered, but also textual sources expressing the meaning that the concept of time had in different cultures and all the concepts, iconographies, and representations possibly related to it.

Q8. What are the attributes that allow us to recognize a subject?

This question is usually raised when an observer tries to identify the subject

he/she is observing, corresponding to the main practice of iconography. To this end, the attributes or details that reveal the subject matter must be identified. van Straten (2012) exemplifies the practice by reading Vermeer's *Woman holding a balance*. At first glance, a genre scene seems to be depicted, showing a woman in her daily life while weighing with a balance. However, a close examination of the balance reveals that it is empty. Parallelism with the painting representing the Last Judgement in the background is, according to Van Straten, a clue for the figure's interpretation. Since the woman weighs something invisible, the action must be symbolic. Christ, during the Last Judgement, symbolically weighs the souls of the faithful. This association allows us to read the figure as a personification of Divine Justice.

Another example in which such a question is addressed and explained methodologically is provided by Panofsky (1955) while interpreting an artwork by Francesco Maffei that presents attributes that may refer to two distinct characters, namely Salomé or Judith. The artwork shows a woman holding a sword and a basin containing a human head. As a first step, Panofsky refers to the literature, as the Bible narrates about two women who beheaded a man, Salomé, who ordered to behead John the Baptist, and Judith, who beheaded Holofernes. According to the text, the head of John was carried on a basin, whereas the one of Holofernes was put in a bag. Although this detail would suggest that the figure is Salomé, she did not kill John herself, whereas Judith did it using a sword. As the sole literary source does not further help identify the figure, the art historian checks whether the iconography with this variation could be found in other previous examples. Whereas there are no known examples of Salomé with a sword, cases in which Judith is represented with the head of Holophernes with a basin are present. Therefore, the comparison with the iconographical tradition and the presence of iconographical types helps to identify the subject depicted in the artwork by Maffei, that is, Judith with the head of Holophernes. In this case, the art historian, to identify the depicted iconography, first looked for clues in the literature and then made a comparison with the iconographical tradition.

Q9. Does the representation of iconography vary?

This question emerges from iconographic and iconological studies in which relevance is given to the evolution of details of the iconography rather than its recognizing attributes. Our example is treated in Panofsky and Saxl (1933, p. 257). When reconstructing the origin of medieval mythography from northern France and England, they notice that, in an illustration of a manuscript from 1420 [3], the three graces are represented, but the one with her back turned is depicted

[3]Illustration of Venus, cod. Vat. lat. 1290, folio 2r. available at https://digi.vatlib.it/

on the left of the group, instead of being in the center, as shown in the classical iconographical type. The fact that medieval mythography was based on a textual tradition and no visual examples were available motivates this position change. Furtermore, the iconographical type was allegorically explained as the given favor (grace with her back turned) returned twice, which clarifies why such detail is kept in the medieval visual representation. Hence, the apparently meaningless variation of the traditional position of the graces reveals that the medieval period relied on a textual tradition of mythographical content and that a purely esthetical trait of the prototype was allegorically explained.

5.1.2.4 Evidence

Q10. What were the known textual sources to which the artwork refers, and what does this knowledge tell us about the thinking of the time?

Although many examples can be provided to highlight the relevant role that textual evidence plays in understanding deeper meanings, we will focus on Warburg's essay on Botticelli's *Nascita di Venere* (Warburg, 1999). Warburg traces the dependence on the contemporary writer Angelo Poliziano in Botticelli's masterpieces. Despite both the painting and the textual source referring to the *Homeric Hymn to Aphrodite*, the painting reflects some details of the description provided by Poliziano in his *Giostra*. In fact, among Poliziano's changes to the classical text, Warburg notes that both include the presence of two *zefiri* rather than one and an accurate description of clothes and hair flying in the wind. Such detail witnesses that the painting aligns with contemporary sources in manifesting the interest in looking for the pathos in antique formulae of movement introduced in descriptive details.

5.1.2.5 Visual Citation

Q11. How do visual shapes migrate and reappear across cultures?

In the first chapter of his book, Wittkower affirms that certain artistic forms usually identified with a specific style and period can reappear in societies far in time and space (Wittkower, 1977) which have been spreading through historically documented migration routes. His analysis of the symbol of the fight between the eagle and the snake, previously described, registers how such symbols are documented from Babylonians onward, and evidence is found in Asia, Japan, Indonesia, Germany, Siberia, and America.

Another type of visual citation that he briefly explores is the influence of a visual formula that spreads in other contexts. He noticed that the visual composition Bernini adopted for his *Rape of Proserpina* (1621-22) spread in Europe, being used

for other subjects such as the *Abduction of Cybele by Saturn*, *Time carrying off the Truth* (Poussin and Rubens), *Time and Opportunity* (David Le Marchand), in which the Opportunity, which is lost for the action of time passing, kills herself.

Question 11 is also the basis of a second type of study conducted by the founder of the iconological approach, namely Warburg's forms of pathos (*Pathosformeln*) elaborated in the classical period, which started to reappear in Renaissance Art to express emotional meanings. For example, he notes that a series of drawings and paintings of the Italian Quattrocento reproduce the scene of the death of Orpheus, making the hypothesis of the presence of a classical prototype copied by such artists as Dürer, Pollaiolo, and Mantegna (Warburg, 1999, pp. 553-59). They maintained the genuine classical appearance and formal composition of the prototype, which expressed the force of emotion to such an extent that the visual appearance of a classical maenad was used, in some cases, to express the despair felt by Magdalene under the crucified Christ, as can be seen in a bas-relief by Bertoldo di Giovanni representing the Crucifixion (1485-1490, Firenze, Museo Nazionale del Bargello), included in table 42 of *Bilderatlas*[4] (Warburg et al., 2020).

A further fundamental *Pathosformel* that Warburg notices is the movement of accessory elements, such as hair or clothes, agitated by air. This feature can be seen in the already introduced Botticelli *Birth of Venus*, and in paintings in which a classicizing figure with fluttering clothes is introduced in a sacred scene, as happens in Ghirlandaio *Birth of Saint John the Baptist* (1485-90), usually described as a 'nymph' from a letter written by André Jolles to Warburg (Warburg, 1999, p. 21).

Q12. Is a visual citation the piece of evidence that documents a cultural phenomenon?

The re-emergence of the formal type of *The Death of Orpheus* previously described witness the 'emotive, rhetorical current within the reawakening of antiquity' (Warburg, 1999, p. 553) that took place in the Early Renaissance, in drawings by Pollaiolo, Durer, and Mantegna. Rather than a stylistic borrowing, the death of Orpheus was a felt theme by contemporaries, as Poliziano's poem Orpheus witnesses. In this sense, the emotion of jealousy expressed by Orpheus beaten by the maenads is embedded in the style and visual form. The fact that the theme represented by its antique visual arrangement and style re-emerged during the early Renaissance witnesses that Italians sought in classic art 'extremes of gestural and physiognomic expression, stylized in tragic sublimity' (Warburg, 1999, p. 558).

An additional example is provided by the *phenomenon of disjunciton* registered by Panofsky and Saxl (1933): whenever the classical shape survives during the

[4]The table is visible at https://www.engramma.it/eOS/core/frontend/eos_atlas_index.php?id_tavola=1042. The topic is deepened by Antal (1937), Wind (1937), Didi-Huberman (2017)

Middle Ages, it is deprived of its classical meaning. Cases can be observed in Italy and southern France, where classical artifacts were available. One example is the Roman relief placed on the external wall of St Mark's Basilica depicting a *Hercules carrying the Herymanthian boar* which is the visual prototype of an *allegory of Christ saving the Christian souls* of the XIII century (see Figure 6.10).

Table 5.4 resumes the research questions that were extracted from the literature.

Table 5.4: Overview of research questions extracted from the literature and the source

Source	**Research Question**
Cultural phenomena	
(Warburg, 1999)	Q1. Which cultural phenomena are witnessed by artworks?
(Cardini, 2001; Warburg, 1999)	Q2. Which sources and visual aspects motivate their presence?
(Panofsky & Saxl, 1933)	Q3. How does the representational evolution of subjects witness the emergence of cultural changes?
Symbols	
(Panofsky, 1955; Wittkower, 1977)	Q4. How do symbols evolve?
(Christiansen, 1986)	Q5. Is the symbolic meaning motivated by a specific source?
(Gombrich, 1972)	Q6. Is the deeper conceptual meaning motivated by sources?
Iconographies	
(Panofsky, 1955)	Q7. How do iconographies, their meaning and attributes evolve over time?
(Panofsky, 1955; van Straten, 2012)	Q8. What are the attributes allowing us to recognize a subject?
(Panofsky, 1972)	Q9. How does the representation of iconography vary?
Evidence	
(Warburg, 1999)	Q10. What were the known textual sources to which the artwork refers, and what does this knowledge tell us about the thinking of the time?
Visual citations	
(Warburg et al., 2020; Wittkower, 1977)	Q11. How do visual shapes migrate and re-appear across cultures?
(Warburg, 1999)	Q12. Is a visual citation the evidence that documents a cultural phenomenon?

5.2 Requirements

Based on the analysis of the iconographical and iconological literature described in Section 2 and on the domain study in the previous section, we formulate the requirements using the SEEMP framework (Suárez-Figueroa et al., 2009)[5]. The terminology was selected mainly from Panofsky's theory (Panofsky, 1955). The output document is described in Tables 5.5 and 5.6.

The purpose of the ICON ontology is to formally represent the domain of knowledge of iconology and iconography with a high granularity level to allow specific quantitative analysis that can be interesting for domain experts. It is intended to be used by i) cultural institutions willing to publish their data about artwork content in linked data, ii) art historians interested in answering iconographical and iconological research questions in a quantitative way, and iii) developers who plan to use computer vision to associate recognized elements to portions of artworks. Therefore, the ontology aims to be implemented in different contexts, meeting the needs of different types of users. We use the OWL2 format to make the ontology available and reusable.

Therefore, the main non-functional requirement[6] is the reuse and alignment to the standards shared across the community to allow reusability. Furthermore, the CQs formulated for functional requirements aim to express the various aspects of the iconographical-iconological approach described in Chapter 2 and Sections 5.1.1 and 5.1.2. We summarize the main themes that can be extracted from the requirements listed in Table 5.6 as follows:

1. The identification of subjects at each level of interpretation needs to be included.

2. The variations of iconographical subjects (e.g., Cupid represented with a bandage and griffon talons, rather than only with traditional attributes, viz. wings and arrows (Panofsky & Saxl, 1933)) must be described.

3. The symbolic and cultural meanings attributed to each subject must be included.

In addition, the relevant characteristics of the approach are considered, namely:

4. The attribution must be subjective.

5. The sources used by the scholar to state its claim must be present.

[5]The ICON ontology development was published in Sartini et al. (2023). Sections 5.2 and 5.3 refer to Sections 4 and 5 of the article. Whereas S Baroncini was responsible for Section 4, the authors B. Sartini, S. Baroncini, and A. Gangemi contributed to Section 5.

[6]see slot 6a of the SEEMP ORSD in (Suárez-Figueroa et al., 2009)

6. The clear distinction between the subjects described at a general level (i.e., the background knowledge necessary for iconographical descriptions cited in Table 2.1 and in section 5.1.1, found in standard vocabularies, describing, e.g., Cupid as a 'child with wings and arrows') and their specific manifestation in a single artwork (e.g., Cupid with griffon talons) needs to be done to allow us to describe variations.

7. The ontology must allow the integration of one claim within the agreeing claims quoted by the art historian as a source of shared and accepted knowledge.

8. The ontology must allow gathering sets of agreeing recognitions made in a coherent situation (e.g., a scholar making an interpretation in a specific paper expanding on other scholars' interpretations, therefore including their claims in his own), that may gather the interdependent recognition made at different levels (e.g., a scholar recognizes the level 2 subject 'Cupid', since he recognized at level 1 the subjects 'child', 'arrows', 'wings').

9. The description of the iconic language of the visual artwork needs to be included, e.g., the relative position of objects and the structure in which they are organized.

10. At least a description of the style should be included.

As Panofsky's theory is considered a representative formalization of the iconological approach, we take most of the ontology terms from his theory. Therefore, we decided to populate the pre-glossary of terms (i.e., the relevant terms extracted by the CQs and their answers) contained in Table 5.6, point 7, by extracting the terms which are answering to CQs directly from the definition of his theory, on the basis of the domain characteristics defined in section 5.1.1. The number following each word indicates the word frequency in the selected article[7], in which Panofsky's theory is fully illustrated.

5.3 Modeling

5.3.1 First design iteration: Recognitions

As explained in Chapter 2 and in Section 5.1.1, works of art can be analyzed through different layers of interpretations that depend on recognitions. A recog-

[7]For this analysis, we referred to the article "Iconography and Iconology: an introduction to the study of Renaissance art' published in (Panofsky, 1955), since it is the last revised version published.

Table 5.5: Description of requirements 1-5 according to SEEMP methodology

	SEEMP Reference Ontology Requirements Specification Document (requirements 1-5)
1	Purpose
	The ontology purpose is to formally represent the domain of knowledge of iconology and iconography with a high granularity level, to provide art historians and cultural institutions a way for expressing complex art subjects and meanings, claims about their interpretations, and interlinking among them.
2	Scope
	The ontology focuses only on the iconographical and iconological interpretations that can be made about the content and meaning of visual artworks. The ontology has a high level of granularity, to correctly represent i) specific data important for domain experts and ii) the subjectivity of each claim.
3	Implementation language
	The ontology has to be implemented in OWL2 language.
4	Intended End-Users
	User 1. Cultural institutions that have a detailed bibliography about artworks looking for a formal language to express it User 2. Art history scholars with complex research questions only answerable with quantitative methods or wanting to express the data they collected in a formal language User 3. Developers using computer vision to associate recognized elements to portions of artworks
5	Intended Uses
	Use 1. Publish structured data about artwork interpretations online and integrate them with existing data to enhance the query potentiality of the cultural institutions' data Use 2. Conduct a specific and detailed quantitative analysis to answer research questions in the domain research field Use 3. Provide a semantic structure for knowledge extraction

nition, in the context of this ontology, is an interpretation act made by an agent (or interpreter, which can be a biological or electronic being) that links works of art to something related to their content. From a conceptual perspective, it is a mental entity reflecting the agent's subjective point of view. From a technical viewpoint, it is an N-ary predicate that cannot be modeled using OWL due to expressivity limitations; therefore, it was turned into an N-ary relationship class.[8] Coherent recognitions on the same artwork are collected and documented by interpretation descriptions (requirement 8, section 5.2).[9] In this iteration, we conceptualize the elements that revolve around recognitions. From the n-ary relationship class `icon:Recognition`, several properties were designed (or reused from existing ontologies) to link it to its interpreter(s) (or agents), the artwork that is being interpreted, supporting sources for the recognitions. In particular, the `aboutWorkOfArt` property links the recognition to the artwork (`Artwork` class). Then the `dul:includesAgent` property (from DOLCE (Gangemi et al., 2002)) links the recognition to the agent who performed it (requirement 4, section 5.2).

[8]More observation on the matter can be found in subsection 5.4.

[9]The distinction between the mental entity of the recognition and the document entity of the description is necessary not only because a description can contain multiple recognitions, but also as a way of separating through coherent criteria different recognitions made on the same artwork (even by the same interpreter). For example, a cultural institution such as a museum might decide to describe an artwork by collecting only some recognitions made by one interpreter and adding more recognitions made by different interpreters to finalize their descriptions.

Table 5.6: Description of requirements 6-7 according to SEEMP methodology

	SEEMP Reference Ontology Requirements Specification Document (requirements 6-7)
6	**Ontology Requirements**
6.a	Non-functional Requirements
	NFR1. The ontology must be based on international standards and, when possible, directly reuse them
6.b	Functional Requirements: Groups of Competency Questions
	CQ level 1. CQ 1.1 What level 1 objects are represented in the artwork? CQ 1.2 What objects are natural elements, expressive characteristics, or actions? CQ 1.3 What level 1 subjects are formally derived or copied from other artworks' level 1 subjects? CQ 1.4 In what compositional structure are the objects organized (e.g., pyramidal arrangement)? CQ level 2. CQ 2.1 What level 2 subjects are identified in each artwork? CQ 2.2 Retrieve, respectively, all characters, events, personifications, named objects, and places recognized at level 2. CQ 2.3 In which story or allegory are the depicted subjects involved? CQ 2.4 Do the level 2 subjects have a symbolic meaning? CQ 2.5 Which is the object that allows character recognition at level 2, i.e., the character's attribute? CQ 2.6 What are the representative variations at level 1 of the same level 2 subject in different artworks? CQ 2.7 What are the level 1 variations of the same level 2 subject involved in different stories or allegories? CQ 2.8 What are the level 1 subjects having multiple interpretations at level 2? Which of them are made in the same descriptive situation? CQ level 3. CQ 3.1 What meanings are expressed by the artworks? CQ 3.2 What cultural phenomena are identified? CQ. 3.3 Who identified the cultural phenomena and on which basis? CQ 3.4 What are the artworks involved in the same cultural phenomenon? CQ 3.5 To which specific subjects at level 1 and 2 does the level 3 recognition refer? CQ 3.6 What are the artworks having both a common cultural phenomenon and a common level 2 subject? General CQ. CQ 0.1 What are the sources that support each subject recognition at each level? CQ 0.2 What is the person responsible for every recognition at each level? CQ 0.3 What are the artworks that are only interpreted on a pre-iconographical level? CQ 0.4 What works of art are interpreted on an iconological level but not on an iconographic one? CQ 0.5 What are the subject recognitions motivating another recognition? Of what type are they? CQ 0.6 What artworks or parts of it have a style associated?
7	Pre-Glossary of Terms (Term, Frequency in studied documents)
	Motif(s) 44; Story(ies), 37; Image, 26; Interpretation, 22; Natural, 16; Iconography, 15; Iconographical, 14; Allegory(ies), 11; Intrinsic meaning, 9; Preiconographical Description 9; Iconographical Analysis, 8; Composition, 7; Expressional, 7; Artistic motifs, 5; Factual, 5; Iconological Interpretation, 4; Iconology, 4; *Invenzioni*, 1

The class `InterpretationDescription` is linked to (one or many) `Recogntion` class(es) that comply with it through several properties according to the type of the recognition, namely: `isCompliantWithPreiconographicalRecognition` for pre-iconographical recognitions and formal motif recognitions,[10] `isCompliantWith IconographicalRecognition` for iconographical recognitions, `isCompliantWith IconologicalRecognition` for iconological recognitions. The CiTO (Shotton, 2010) properties `cito:citesForInformation` and `cito:citeAsEvidence` can be linked to a `icon:Recognition` class to provide sources or other information that support a recognition (requirement 5, Section 5.2). Finally, a recognition can also be used to support further recognitions made on the same artwork or another one. For example, Panofsky recognizes that the figure of Chastity sculpted by Giovanni Pisano on the Pulpit of the Pisa Cathedral is represented with the same appearance of the classical nude iconography of Venus Pudica (formal motif recognition)[11]. This interpretation supports the third-level recognition of the characteristics of the Proto-Renaissance movement in the cultural context of Medieval Tuscany (Panofsky, 1972, p 157). To express this using our ontology, the property `cito:givesSupportTo` can link the supporting recognition to another one (requirement 7, Section 5.2). These elements are also the object of interest of the general competency questions (see Table 5.6, Q0.1 to Q0.5).

Depending on the level of interpretation presented in Table 2.1, four `Recognition` subclasses have been defined:

- `PreiconographicalRecognition` (level 1)
- `FormalMotifRecognition` (level 1)
- `IconographicalRecognition` (level 2)
- `IconologicalRecognition` (level 3)

Recognitions at each level of interpretation may be based on the results of the recognition at one of the previous levels. Therefore, they can be linked together but ultimately are modeled as independent of one another. This choice is made since i) the describer may not have available the descriptions of the lower level(s), ii) the corresponding subjects in the other levels may not be relevant for the recognition, iii) it may be possible that a level 3 recognition (i.e., an `IconologicalRecognition`) is linked to level 1 subjects rather than level 2 ones (e.g., iconological interpretations of a landscape painting, which may not have level 2 subjects (Panofsky, 1955)).

[10] Both a pre-iconographical recognition and a formal motif recognition are described in the first level of pre-iconographical interpretations of Panofsky, so we use the same property to link them to the `InterpretationDescription` class

[11] figure available at https://www.Italianways.com/wp-content/uploads/2014/07/Giovanni-Pisano-Pulpito-Duomo-Pisa-06.jpg

These classes and their specific usage will be described further in the following subsections. Figure 5.2 shows a rendering of the classes and properties of this iteration.

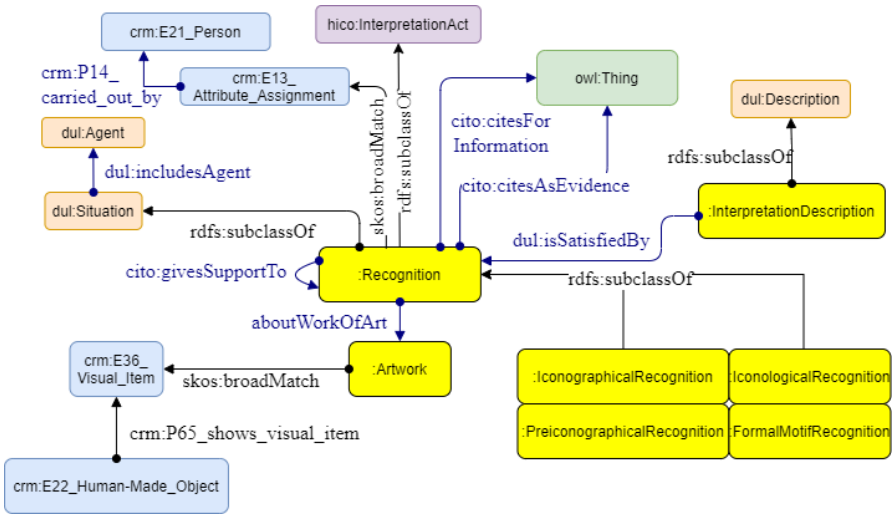

Figure 5.2: ICON ontology classes and properties linked to recognitions

5.3.2 Second design iteration: Pre-iconographical Recognitions (level 1)

In this iteration, we model the recognitions that happen on a Pre-iconographical level. In this level, an interpreter recognizes artistic motifs present in the artwork, and associates to them i) natural objects (a tree, a man, a sword) without identifying specific individuals from those classes which are recognized in level 2 (e.g., tree of life, Saint Joseph, Excalibur; see Section 5.3.3), ii) in the form of expressional meanings [12] (emotions of the depicted elements), iii) qualities about these elements (size, colour, positions), iv) performed actions (see Table 2.1 in chapter 2). Assuming that the agent doing the interpretation act might also be a computer, as in the case of the results of object detection through computer vision, we give the possibility to express coordinates of the portion of the image of the artworks where these elements are detected. Furthermore, these coordinates can be expressed using IIIF URIs (Snydman et al., 2015) that point to a specific portion of the work of art. A series of artistic motifs can be grouped together in a composition that can have a compositional structure[13] (e.g., pyramidal). Additionally, an interpreter might recognize similarities between artistic motifs present in a work of art with other artistic motifs of another work of art, recognizing a prototypical artistic motif or

[12] according to Panofsky, the expressional meanings are the subjects that can be interpreted at the first level of recognition through empathy (Panofsky, 1955, p. 27)

[13] The compositional structure conceptualization is derived from Imdahl's theory (Imdahl, 2012)

composition that is reused in another artwork. For example, the level 1 description of Pisano's figure of Chastity cited above is linked through a formal motif recognition to the level 1 description of Venus Pudica, from which its appearance is derived (i.e., a nude woman covering herself with her arms). Artistic Motifs and compositions are linked to the class `PreiconographicalRecognition` respectively through the properties `recognizedArtisticMotif` and `recognizedComposition`. Only one artistic motif or composition can be linked to a recognition. Compositions are linked to the artistic motifs that take part in them through the `hasPart` property. If the artistic motif refers to a natural object or action with a factual meaning, it is linked to the classes `NaturalElement` or `Action` through the property `hasFactualMeaning`. Otherwise, if what is recognized in the artistic motif is an expressional meaning, the property that links it to expressional meanings is `hasExpressionalMeaning`. If actions, expressional meanings, or natural elements have some specific quality that needs to be highlighted, from the artistic motif the qualities are expressed with the DOLCE `hasQuality` property. When the pre-iconographical recognition is performed by a computer with an object detection algorithm, or when a IIIF URI is provided, it is possible to associate not only the detected objects, but also the coordinates of the image in which they are found. Coordinates of the detected object can be expressed through the data property `hasRegionDescription` that has the `ArtisticMotif` or `Composition` classes as the domain. As mentioned above, the use of IIIF URIs for the format of this data property is also welcomed. The `FormalMotifRecognition` class links the prototypical motif to the copied motif, respectively, using the `hasPrototypicalMotif` and `hasCopiedMotif` properties. Finally, all the coherent formal motif recognitions and pre-iconographical recognitions that take part in an interpretation about a work of art, can be linked to an `InterpretationDescription` class, through the property `preiconographicallyCompliesWith`. Figure 5.3 shows a graphical rendering of the classes and properties used in this interpretation level.

5.3.3 Third design iteration: Iconographical Recognitions (level 2)

In this third iteration, we focus on Panofsky's second level of art interpretation: the Iconographical interpretation. In this level, the interpreter recognizes images and *invenzioni*[14] in an artwork. An image represents the subject depicted as a manifestation in the specific artwork taken into account. It is then linked to second

[14]Invenzione is an Italian word used by Panofsky as an umbrella term for allegories and stories (Panofsky, 1955, p. 59)

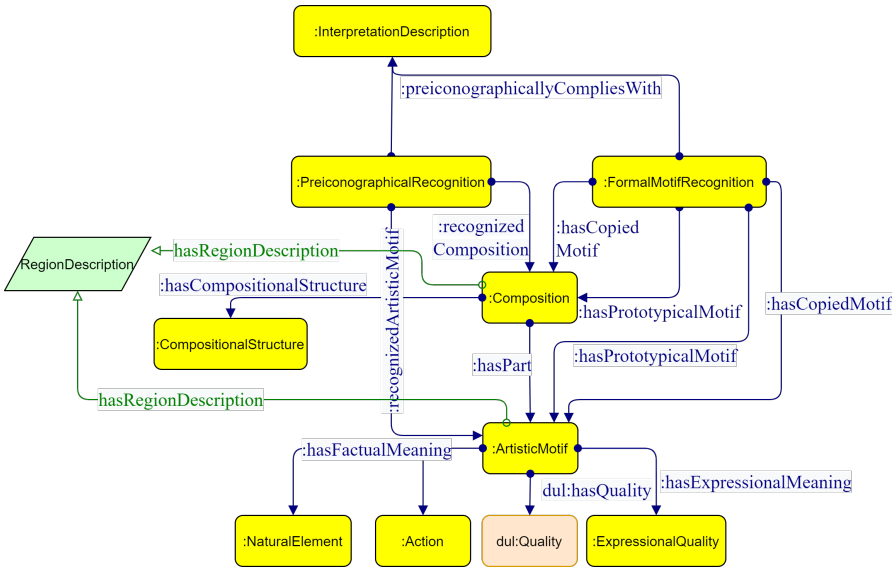

Figure 5.3: ICON ontology classes and properties describing the pre-iconographical level of interpretation (level 1)

level subjects, which are characters, places, events, named objects[15], symbols, personifications, identifying iconographies from an abstract and general point of view. This distinction between the general subject level (i.e., characters, symbols) and the artwork-specific one (image) is functional to identify the variants of a subject in relation to the specific context (i.e., Thor as represented in a specific painting may differ from its common one). An *invenzione*, instead, is the subject matter represented by the combination of general subjects linked to the single images recognized.[16] For example, a viewer might recognize three images in an artwork: the first refers to the general subject of Mary, the second refers to the general subject of Angel Gabriel, and the third refers to the general subject of the Holy Dove. The combination between the general subject of Mary, Angel Gabriel and the Holy Dove is the Annunciation, which, in our ontology terms, would be considered the *invenzione*. The same *invenzione* could be present in multiple artworks, but each artwork maintains its uniqueness by having different images. The classes `Story` and `Allegory` are subclasses of the class `Invenzione`. Stories are more likely to contain characters, named objects, places and events, whereas allegories are more likely to contain symbols and personifications. We give

[15] A named object is a non-living unique element that is often used as an attribute for the recognition of specific characters (Thor's hammer.)

[16] This definition slightly differs from Panofsky's: while he describes an *invenzione* as a form expressing the subject represented by the combination of the single images recognized, we consider it as an individual belonging to the 'sphere of secondary or conventional subject matter, viz., the world of specific themes or concepts manifested in images' (Panofsky, 1955). This decision is motivated by the fact that, the description of real case studies in the modeling phase, emerged that it would be redundant to general stories and allegories both at the conventional level and their manifestation in the specific artwork. Their variations are already clear, considering which subjects are part of them in each particular case.

the possibility to express symbols as just symbolic meanings recognized, or, for a more thorough description, as Simulations (see section 5.4). The classes `Image` and `Invenzione` are linked to the class `IconographicalRecognition` through the respective properties `recognizedImage` and `recognizedInvenzione` (one image or *invenzione* per recognition). The artistic motif belonging to a pre-iconographical level that refers to the recognition of an image can be linked to it with the property `refersToArtisticMotif` (e.g., the recognition of the image that represents Mary Magdalene can be linked to the artistic motif that has the factual meaning of woman). This link is important to ensure that the connection between pre-iconographical elements and the respective iconographical subjects is preserved. If the artistic motif is the principal element that enabled a recognition of an image, then it can be linked to that image through the property `hasRecAttribute` (i.e., the recognition identifying Cupid has recognizing attributes the artistic motifs linked respectively to "wings' and "arrows'). Images are linked to the general subject portrayed through specific properties according to the subject class. The property `hasCharacter` links an image to the class `Character`, likewise: `hasEvent` refers to the class `Event`, `hasPlace` refers to the class `Place`, `hasNamedObject` refers to `NamedObject`, `hasSymbol` refers to `Symbol` and finally, `hasPersonification` refers to `Personification`. The cited ICON classes represent second-level subjects represented in the fictional representational space, therefore including both real and fictional, non-existent subjects (e.g., Medusa, the Greek mythological character appearing in various media), in compliance with the modeling of subjects in narratology (Bartalesi et al., 2017; Ciotti, 2016; Damiano & Lieto, 2013). An *invenzione* is linked to the elements that compose it through the property `composedOf`. Finally, multiple iconographic recognitions that take part in an interpretation of an artwork are linked to the interpretation using the `iconographicallyCompliesWith` property. Figure 5.4 shows the classes and properties relative to this level of recognition.

5.3.4 Fourth design iteration: Iconological Recognitions (level 3)

Iconological interpretations (third level) focus on the recognitions of intrinsic meanings[17]. An intrinsic meaning links the whole artwork or some parts of it to a cultural phenomenon or a concept that defines it. The `IconologicalRecogniton`

[17]Even if Panofsky's terminology seems to prefer the term *symbolic values* for expressing the interpreted third level aspects of the artwork, we decided to adopt the term 'intrinsic meanings' to avoid confusion with the second level symbols.

class is linked to the `IntrinsicMeaning` class [18] through the property `recognized IntrinsicMeaning`. From there, the n-ary class `IntrinsicMeaning` can be linked to a specific composition, image or artistic motif that can be the focus of the intrinsic meaning through the properties `hasComposition`, `hasImage`, `hasArtisticMotif`. Then, it is linked to the expressed concept through the property `recognizedConcept`. For the range of this property, we reuse the DOLCE class `SocialObject` because there was no need to create an ad-hoc class for this element.[19] Additionally, since an Intrinsic Meaning can also reflect some cultural phenomena, it is linked to the class `CulturalPhenomenon` through the property `recognizedCulturalPhenomenon`. Currently, `CulturalPhenomenon` has 4 subclasses, which specify the type of cultural phenomenon, namely `Attitude`, `Belief`, `CulturalValue`, and `Tendency`. These terms are taken from Panofsky's vocabulary in the description of the third level of artistic interpretation.[20] Finally, all the iconological recognitions that take part in an interpretation made on an artwork are linked to it with the property `iconologicallyCompliesWith`. A graphical rendering of this fourth iteration, representing the third level of the interpretation, can be found in Figure 5.5.

5.3.5 Ontology extension for simple interpretations

A final design phase extended the ontology to introduce property chains[21] to create shortcuts, allowing direct links between artworks and their respective first, second, and third level of interpretation subjects[22]. By leveraging these super properties, the ontology meets the requirements of ontology usability by museum institutions, which may be interested in a lighter description of artwork content when little information is given. At the same time, it fosters information retrieval when a detailed description is required, i.e., for expressing complex interpretations by Art History scholars. As the properties are defined as property chains, they can be inferred through the use of a reasoner when a thorough description using the ontology classes and properties described above is provided. The extension

[18] Compared to factual and expressional meanings expressed through a property, an intrinsic meaning needed an n-ary class for representation because of expressivity reasons (owl does not support n-ary predicates).

[19] The concepts, ideas, abstract elements that are linked to intrinsic meanings on an iconological level are very broad (Panofsky, 1972). Therefore, we decided to reuse this DOLCE class (Social Object) which conceptualises a broad set of possible entities (Gangemi et al., 2002)

[20] Although these subclasses could be formally associated with mental entities just as recognitions, they differ in their function. Recognition are modelled on a meta level of the interpretation, as they are used to describe a recognition act made by an interpreter. These subclasses are meant to be the object of the interpretation, as they are associated with the recognition of an intrinsic meaning of the artwork itself. As it will be discussed in the final section, further work will be dedicated to a more thorough description of cultural phenomena and their subclasses.

[21] https://www.w3.org/TR/owl2-primer/#Property_Chains

[22] The ontology extension is published in (Sartini & Baroncini, 2023)

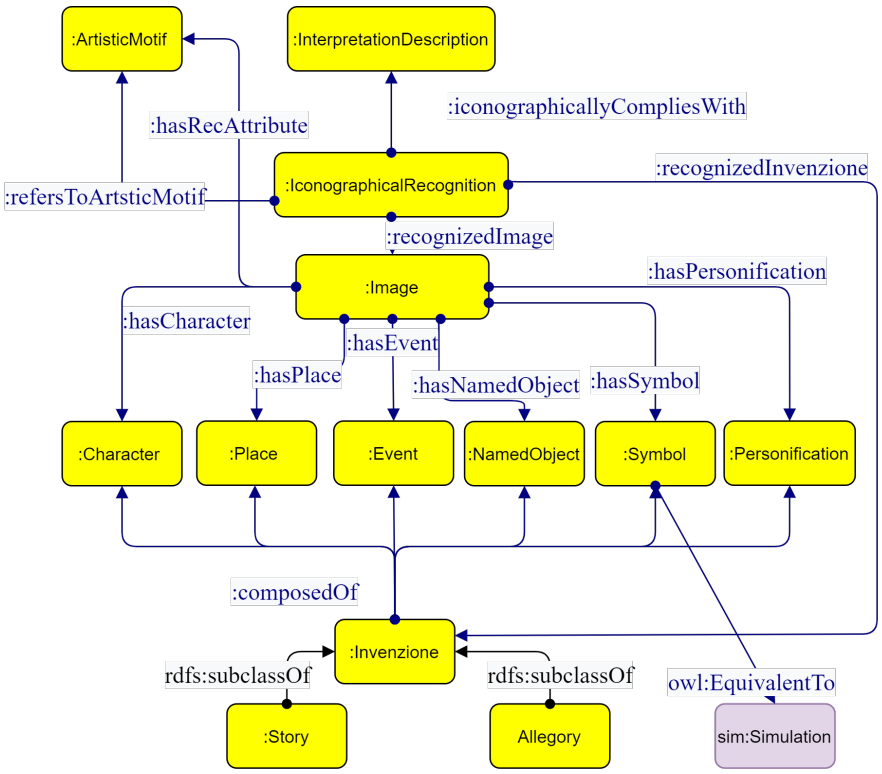

Figure 5.4: ICON ontology classes and properties describing the iconographical level of interpretation (level 2)

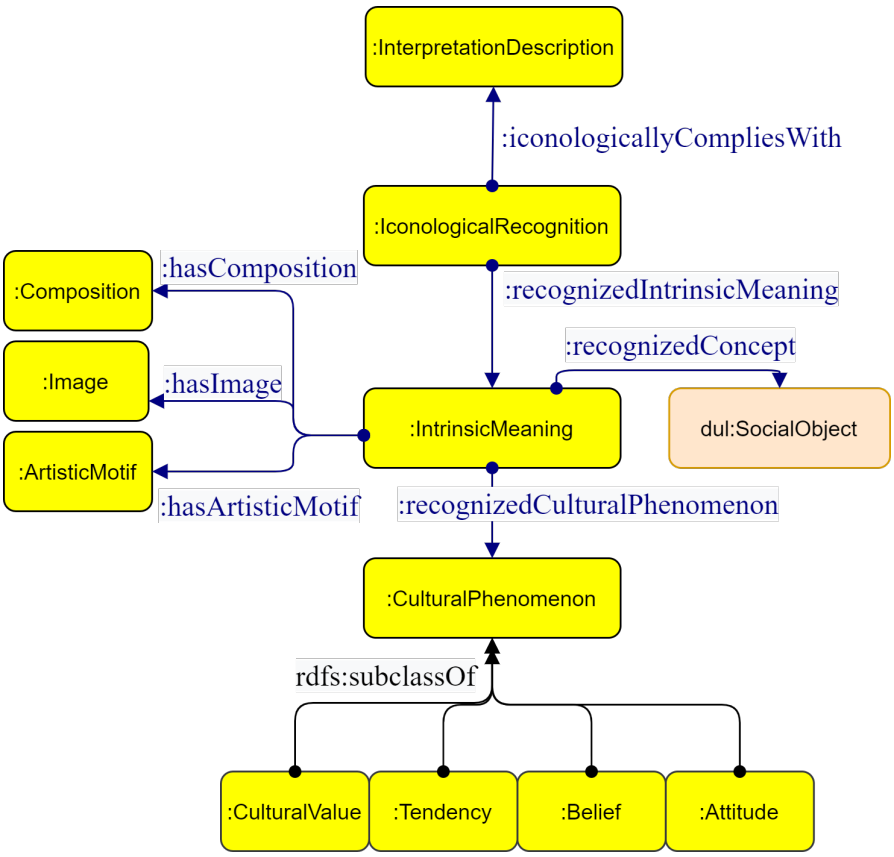

Figure 5.5: ICON ontology classes and properties describing the iconological level of interpretation (level 3)

consists of the addition of three properties, one for each level, relating instances of the class `icon:Artwork` and their respective described subjects.

`icon:preiconographicallyDepicts` links the `Artwork` to the recognized level 1 subjects, as it is expressed as a property chain of `icon:aboutWorkOfArt`, `icon:recognizedArtisticMotif` or `icon:recognizedComposition`, and `icon:hasExpressionalMeaning` or `icon:hasExpressionalMeaning` or `crm:P138_represents`. These chains come into play when the elements of the first level of interpretation are identified either through the recognition of an artistic motif or as part of a recognized composition.

The property `icon:iconographicallyDepicts` links an artwork and a second level of interpretation element that the artwork represents. This property, similar to the previous one, is made a super property of a chain that connects an artwork to instances of classes defining characters, symbols, personifications, places, events, named objects, stories, and allegories (including their stories and allegories components).

Third, we define the property `icon:iconologicallyRepresent` to link an artwork to the cultural phenomenon or abstract concept it represents, characterizing the third level of interpretation. This property serves as the super property of a chain connecting an `icon:Artwork` to instances of either `icon:CulturalPhenomenon` or `dul:SocialObject` classes. We present the chains in Manchester Syntax below:

- preiconographicallyDepicts:

```
hasRecognition o recognizedArtisticMotif o
    hasExpressionalMeaning

hasRecognition o recognizedArtisticMotif o hasFactualMeaning

hasRecognition o recognizedArtisticMotif o crm:P138_represents

hasRecognition o recognizedComposition o hasPart o
    hasFactualMeaning

hasRecognition o recognizedComposition o hasPart o
hasExpressionalMeaning

hasRecognition o recognizedComposition o hasPart o
crm:P138_represents
```

- iconographicallyDepicts:

`hasRecognition o recognizedImage o hasCharacter`

`hasRecognition o recognizedImage o hasPersonification`

`hasRecognition o recognizedImage o hasEvent`

`hasRecognition o recognizedImage o hasPlace`

`hasRecognition o recognizedImage o hasSymbol`

`hasRecognition o recognizedInvenzione`

`hasRecognition o recognizedImage o hasNamedObject`

`hasRecognition o recognizedInvenzione o composedOf`

- iconologicallyRepresents:

  `hasRecognition o recognizedIntrinsicMeaning o recognizedCulturalPhenomenon`

  `hasRecognition o recognizedIntrinsicMeaning o recognizedConcept`

5.4 Refactoring: reuse and alignment to relevant ontologies and ontology design patterns

To promote ontology interoperability and reusability, we connect to several external ontologies through means of alignments and reuse. We present our alignments and reuse by following guidelines proposed by the state of the art (V. Carriero et al., 2020; Osman et al., 2021). Our ontology selection for reuse and alignment 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 of them offer design solutions to the competency questions defined from the requirements in Section 5.2.

Due to the complexity of the field, the number of ontologies to be reused, and the heterogeneous domains from which they come, we adopted a hybrid reuse approach

(V. Carriero et al., 2020), which, depending on the specific cases explained below, considers either reusing directly the classes and properties of the aforementioned ontologies (either by importing the whole ontology or parts of it), or (indirect reuse) using them as fully extensional ontology patterns, or just as intensional patterns.

Extensional reuse happens when classes or properties of an ontology $O1$ are logically *aligned* to an external ontology $O2$, which we want to reuse with its full-fledged semantics because it is compatible, desirable, or necessary. For example, if we extensionally align a $O1$ class `Organisation` to a $O2$ class `dul:SocialObject`, we intend to inherit the semantics of DOLCE's social objects, e.g., that they are not physical.

On the contrary, we use parts of an external ontology $O3$ as purely intensional constructs when we want limited interoperability, which does not include accepting in $O1$ all the semantics provided in $O3$, because it may be partly incompatible. For example, we may intensionally align a $O1$ class `Image` to a $O3$ class `crm:E36_Visual_Item` because we might not want to inherit the axiom stating that `crm:E36_Visual_Item` is a subclass of `crm:E89_Propositional_Object`.

In order to implement this distinction, indirect reuse is designed using different mapping properties, according to the semantics they provide, and its impact into the resulting reasoning. We have used RDFS (`rdfs:subPropertyOf`, `rdfs:subClassOf`) and OWL (`owl:equivalentTo`) logical properties when we want the alignments to provide first-order extension to ICON schema and data, while we have used SKOS `skos:broadMatch`, `skos:related`, and `skos:closeMatch` for purely intensional mapping, which can be used at query time to integrate data represented with ontologies that may harm the logical integrity of ICON knowledge.

Among the reused ontologies, we have used an intensional (or 'terminological') mapping for CIDOC, VIR and FRBRoo, because we have noticed potential problems when reasoning is jointly made with both the axioms from ICON, and from those ontologies. For example, a full extensional alignment of the class `icon:Image` as `rdfs:subclassOf crm:E36_Visual_Item` would make an automated reasoner infer that `icon:Image rdfs:subclassOf crm:E89_PropositionalObject`, which is not defendable, since propositional entities typically exclude visual, musical, or other information modalities. In other words, CIDOC contains here a debatable assumption, which should be ignored when reusing data that use CIDOC as their schema. Now, if we use a purely intensional mapping: `icon:Image skos:broadMatch crm:E36_Visual_Item`, we make a commitment that can be discussed, and the triple can be used to make SPARQL-based data integration, but we will not get the inference that images are propositions.

In this section, we give a thematic overview of the classes and relations reused

for satisfying a specific task, and we refer to the documentation for further details on the single alignments. Table 5.7 shows the direct reuse of external classes and properties in ICON, and Table 5.8 shows the indirect alignments.

Table 5.7: Direct Reuse of Classes and Properties in ICON

External Element	**Type**	**Ontology**	**Usage**
Agent	Class	DOLCE	represents interpreter (with dul:includesAgent property)
Quality	Class	DOLCE	represents recognized quality of artistic motifs (linked from icon:ArtisticMotif with the dul:hasQuality property)
SocialObject	Class	DOLCE	used as the symbolic meaning linked to an icon:IntrinsicMeaning class through the property icon:recognizedConcept
includesAgent	Property	DOLCE	links icon:Recognition to the agent (dul:Agent) performing it (also a non-human agent)
givesSupportTo	Property	CiTO	links icon:Recognition to another icon:Recognition that supports it
citesAsEvidence	Property	CiTO	links icon:Recognition to an entity (owl:thing) that is the evidence on which the recognition is based
citesForInformation	Property	CiTO	links icon:Recognition to an entity (owl:thing) that is the source in which the recognition is found (e.g., a bibliographical reference)

5.4.0.1 Recognitions as situations

According to the guidelines of eXtreme Design (Presutti et al., 2009), we defined our local problem (in our local space) as the expression of recognitions through N-ary relationship classes due to the inability of expressing N-ary predicates in OWL. As explained in the previous paragraphs, our conceptualization of the `icon:Recognition` class required a good deal of contextual information (such as the agent performing it, what is recognized in the form of first, second, or third level of interpretation entities, the artwork). We have chosen the situation ontology design pattern[23] as a solution because it was designed to solve modeling issues regarding multiple contextual information connected to the same class in the form

[23] http://www.ontologydesignpatterns.org/cp/owl/situation.owl

Table 5.8: Indirect Reuse of Classes and Properties in ICON: icon element - type - external element

External Element	**Type**	**Ontology**	**ICON Element**	**Type of alignment**
E5_Event	Class	CIDOC	Action	skos:broadMatch
E36_Visual_Item	Class	CIDOC	Artwork; ArtisticMotif; Composition; Image; IntrisicMeaning	skos:broadMatch
InformationObject	Class	DOLCE	Artwork	rdfs:subClassOf
E13_Attribute_ Assignment	Class	CIDOC	Recognition	skos:broadMatch
InterpretationAct	Class	HiCO	Recognition	rdfs:subClassOf
Situation	Class	DOLCE	Recognition	rdfs:subClassOf
Description	Class	DOLCE	InterpretationDescription	rdfs:subClassOf
Simulation	Class	Simulation Ontology	Symbol	owl:equivalentTo
F38	Class	FRBRoo	Character	skos:broadMatch
E1_CRM_Entity	Class	CIDOC	ExpressionalQuality; Invenzione; NaturalElement	skos:broadMatch
E31_Document	Class	CIDOC	InterpretationDescription	skos:broadMatch
E90_Symbolic_Object	Class	CIDOC	Symbol	skos:broadMatch
IC11_Personification	Class	VIR	Personification	skos:closeMatch
Subject	Class	ArCo	Character; Personification; Event; NamedObject; Place; Symbol; Invenzione; Action; NaturalElement; ExpressionalQuality	rdfs:subClassOf
E89_Propositional _Object	Class	CIDOC	Event; NamedObject; Place	skos:related
P138_represents	Property	CIDOC	hasCharacter; hasEvent; hasExpressionalMeaning; hasFactualMeaning; hasNamedObject; hasPersonification; hasPlace; hasSymbol	skos:broadMatch
P140_assigned_attribute_to	Property	CIDOC	associatedForm; refersToArtisticMotif	skos:broadMatch
P141_assigned	Property	CIDOC	recognizedArtisticMotif; recognizedComposition; recognizedImage; recognizedInvenzione	skos:broadMatch
P106_is_composed_of	Property	CIDOC	hasPart	skos:broadMatch
K4_is_visual_prototype_of	Property	VIR	hasCopiedMotif; hasPrototypicalMotif	skos:broadMatch

of n-ary relationships. The Situation ontology design pattern is reused via the import of DOLCE Ultralight.[24] The n-ary relationship ODP is specialized by our `icon:Recognition` class, by making it a subclass of `dul:Situation`.

5.4.0.2 Interpretations as descriptions

The types of recognitions that we have presented are formalized as situations. In the Descriptions and Situations pattern[25] that is also formalized in DOLCE-Ultralight and DOLCE Zero, situations are loosely associated with *descriptions*, i.e., intensional entities that are used criteria for a situation to occur. The pattern is used in most domains: in medicine, a pathological situation depends on the diseases or syndromes that are used to interpret it, and which can have different probabilities to correspond to the actual situation; in Law, different norms may apply to a same legal case; in an everyday situation, an observer may interpret it differently according to his/her perspective, culture, or intention. In the iconographical and iconological domain, as also applied in the ArCo ontology network (V. A. Carriero et al., 2019; V. A. Carriero et al., 2021), all recognitions and high-level interpretations are based on perception criteria, which make a rationale emerge, and eventually motivate a particular interpretation with respect to others. A description is therefore a conceptual entity, constituted by parameters, roles, tasks, etc. (Gangemi & Mika, 2003), which is satisfied by a situation when it involves entities that are classified by one of the parameters, roles, tasks, etc. that constitute a description. For example, the interpretation of a painting (Named *A*) such as 'in this painting, there is a lion which symbolizes courage' is compliant with (i) a pre-iconographical recognition (recognizing an artistic motif as a carrier of the factual meaning of a lion), and (ii) an iconographical recognition (recognizing the image of the lion as the simulation of lion-courage). These recognitions would involve the recognizer, a source, the time period, as well as (potentially) additional iconographical aspects. Hence, we formalize this complex relation in terms of compliance: `InterpretationDescriptionPaintingA isCompliantWithPreiconographicalRecognition LionRecognitionInA` and `isCompliantWithIconographicalRecognition LionCourageRecognitionInA`. The property `isCompliantWithPreiconographicalRecognition` is made a sub-property of `dul:isSatisfiedBy` which links a `dul:Description` (our `InterpretationDescription` is subsumed under description) to one or more `dul:Situation` (our `Recognition` is subsumed under situation).

[24] The aforementioned ODP was reused through the DOLCE Ultralight ontology because the Situation class in DOLCE is linked to other classes that are reused in our ontology as well, such as the agent.

[25] http://www.ontologydesignpatterns.org/cp/owl/descriptionandsituation.owl

5.4.0.3 Describing artwork content

Since CIDOC CRM offers a way of describing the content of visual elements (`crm:E36_Visual_Item`, `crm:P138_represents`, `crm:E1_Entity`), we modeled the more specific elements recognized in each level of interpretation following this modeling principle as a guideline, and aligning our classes to CIDOC's ones through SKOS relations. As illustrated by Figure 5.6, all the classes representing the general subject as represented in the contest of the artwork (i.e., Artistic Motif, Composition, Image, IntrinsicMeaning) are a `skos:broadMatch` of `crm:E36`. Furthermore, the recognized subjects at every level are a `skos:broadMatch` of `crm:E1_Entity`. In this way, the patterns linking the visual elements recognized in each level and the general subject can be seen as a specification of `crm:P138`. Identification of the artwork at the abstract level (Artwork, `skos:broadMatch` of `crm:E36_Visual_Item`) is intended to make the ontology compliant with the CIDOC-CRM modeling of cultural objects, whereas the alignment of Artwork with `dul:InformationObject` is motivated by the DOLCE conceptualization of the Information Object that fits our Artwork definition. Similarly, we declare a match between the recognized subjects and a CIDOC-CRM class, when available. In this way, `icon:Action` is aligned with `crm:E5_Event`, for including also actions not directly made by human beings. `icon:Event`, `icon:NamedObject` and `icon:Place` are aligned with `crm: E89_Propositional_Object`, as such subjects may describe real or imaginary places, following the solution adopted by Schneider et al. (2021, p. 11).

5.4.0.4 Interpretation details

The class `Recognition` has been aligned with classes from HiCO, CIDOC-CRM and DOLCE, as shown in Figure 5.2. The class `hico:InterpretationAct` is intended to represent the context in which a recognition[26] is made, i.e., furnishing more information about the recognition to validate the claim. The recognition such represented can be further specified by `hico:interpretationType` and `hico:Interpretation-Criterion`. For its purpose and formal structure, `icon:Recognition` was made a subclass of it. Since also the purpose expressed by `crm:E13_Attribute_Assignment` is of documenting the context in which an assertion about a cultural object was made, it is a `skos:broadMatch` of Recognition, as Recognition is more specific than the more generic concept expressed by `crm:E_13`. Furthermore, `crm:E13` is practically used as an n-ary relationship class linking two individuals through ancillary properties, `crm:P140`, `crm:P141`,

[26] In the context of this study, since we align `hico:InterpretationAct` to Recognition, we refer to it with the term *recognition* for the clarity sake

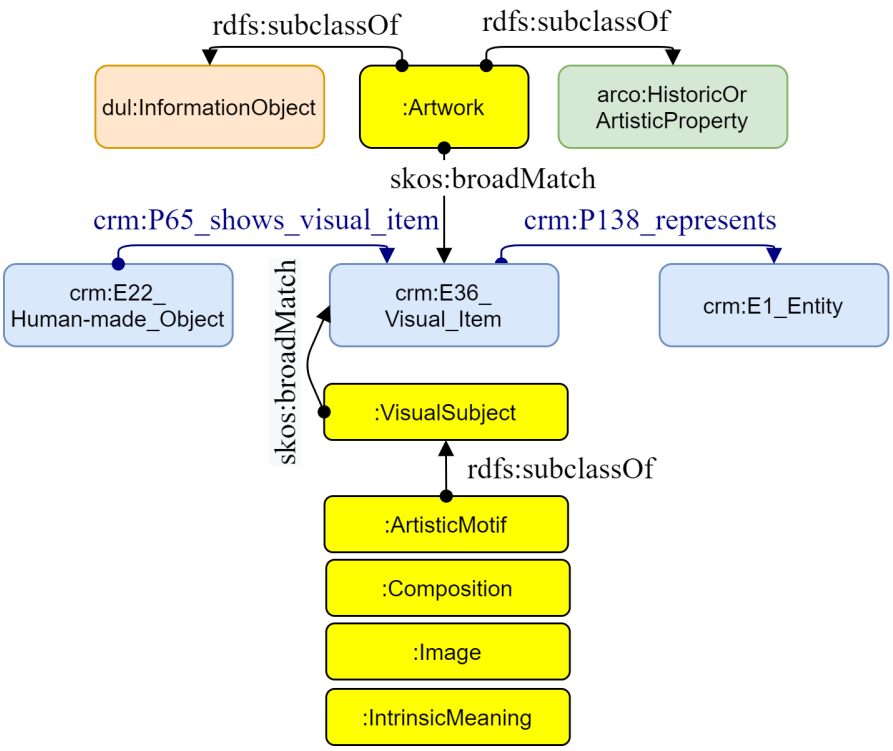

Figure 5.6: Alignment with other ontologies of the artwork description

identifying respectively the element to which the assignment is made and the assigned one. Therefore, when this logical structure is respected, the respective properties in the subclass of `icon:Recognition` are aligned to `crm:P140` and `crm:P141` through `skos:broadMatch`. Respectively, `RecognizedArtisticMotif` or `RecognizedComposition` at level 1, `RefersToArtisticMotif` and `Recognized Image` or `RecognizedInvenzione` at level 2.

By the alignment with `hico:InterpretationAct`, and `dul:Situation`, the ontology not only enhances interoperability but also inherits a variety of means for expressing further detail about each recognition act at each level. For example, the possibility to express an agent using `dul:Agent` which includes both humans and computers, the time of the recognition using the `includesTime` property of DOLCE, the interpretation criterion[27] `InterpretationDescription` class and type (HiCO) allows the user to fully document the recognition acts, giving a comprehensive representation of the subjectivity of the recognition itself.

The Motif Recognition is developed as a specialization of the VIR property `K4i_has_visual_prototype`, documenting the use of a visual prototype for an image, enriching the latter by giving the possibility to add further details about the interpretation and to highlight the direct correspondence between the portions of the copying and copied artworks. For example, the derivation of the visual arrangement

[27] in the case of our ontology, interpretation criterion is linked to every single recognition, and not with the interpretations

of the relief *Allegory of salvation* from the Roman relief depicting *Hercules and the Caledonian Boar* described by (Panofsky & Saxl, 1933, p. 228, fig. 4-5, p. 231), can be further described by recognizing that the deer in the former is derived from the boar in the latter, and so on. Our property `icon:hasPrototypicalMotif` was aligned with `skos:broadMatch` to `K4_is_visual_prototype_of`.

5.4.0.5 Subjects

As it is the closer definition of artistic subject intended as an object represented by an artwork, we align all the subjects of the ontology to the ArCo's class `arco:Subject`. Specifically, we indirectly reuse `arco:Subject` by subsuming `icon:Place`, `icon:NamedObject`, `icon:Character`, `icon:Event`, `icon:Symbol`, `icon:Personification`, `icon:Action`, `icon:NaturalElement`, `icon:ExpressionalQuality`, `icon:CulturalPhenomenon`, `icon:Invenzione`, `dul:SocialObject` to it. In doing so, we also propose a new way of attributing a subject to a work of art compared to ArCo. In fact, while ArCo directly links a subject to the physical representation of the work of art, we link it to an interpretation made on the visual representation of what is in a physical work of art. By reusing the class `arco:Subject` and not its properties, which consider the physical artwork as the domain, we also avoid possible logical inconsistencies between ArCo's description of physical artefacts and ICON description of visual items. In contrast, the representation of the subjects as manifested in the artwork (i.e., Artistic Motifs, Compositions, Images, and Intrinsic Meanings), are subclasses of `icon:VisualSubject`, which is disjoint with `arco:Subject` to underline their different nature and role. Table 5.9 displays the division between subjects and visual subjects according to the different iconographic and iconological levels.

5.4.0.6 Symbols

In an artistic interpretation, an interpreter might recognize a symbol of a specific cultural context in an artwork. For the modeling of symbols, we reuse the entire Simulation Ontology (Sartini et al., 2021). This ontology, designed to conceptualize cultural symbols, uses the n-ary `sim:Simulation` class to link together a symbol, expressed by the class `sim:Simulacrum`, its symbolic meaning, expressed by the `sim:RealityCounterpart` class, the cultural context in which the symbol denotes the symbolic meaning (`sim:Context`) and the source of the claim (`sim:Source`). We aligned our class `icon:Symbol` to the `sim:Simulation` class to allow the expression of symbolic meanings using the Simulation Ontology structure.

Table 5.9: Distinction between subjects and visual subjects depending on Panofsky's levels of interpretation

Subject	**Visual Subject**	**Level**
Action	Artistic Motif or Composition	I
Natural Element	Artistic Motif or Composition	I
Expressional Quality	Artistic Motif or Composition	I
Character	Image	II
Event	Image	II
Named Object	Image	II
Place	Image	II
Personification	Image	II
Symbol	Image	II
Invenzione	A series of images	II
Social Object	Intrinsic Meaning	III
Cultural Phenomenon	Intrinsic Meaning	III

5.4.0.7 Expression of Style

The expression of style is an important characteristic related to iconographic and iconological studies (Requirement 10, Section 5.2). Knowing the history of styles is, according to (Panofsky, 1972), a fundamental requirement for the correct interpretation of level 1 objects. Furthermore, as is evident, among others, from Warburg's studies on *Pathosformeln* and *Nachleben der Antike*, forms of style are a subject of interest in iconology. Therefore, we reused CIDOC-CRM to model it according to the solution adopted by the linked.art project[28], using the structure.

```
crm:E36_Visual_Item crm:P2_has_type <style_type>.
<style_type> crm:P2_has_type <aat:300015646>.
```

where the last object is the Getty AAT vocabulary term defining style. Although the property's domain is `crm:E1_Entity`, it is suggested to use it with `crm:E36_Visual_Item`, in compliance with linked.art directions. Even if we do not express `icon:VisualSubject` and `icon:Artwork` as direct subclasses of `crm:E36`, it is possible to reuse this pattern since ICON's classes are not disjoint with CIDOC's ones. Therefore, we reuse this existing solution to model requirement 10 of section 5.2 In this way, both the artwork itself and every portion of the image identified at each level can have its own style specification declared.

The same pattern can be adopted to specify the genre of the artwork, using the term for genre (`aat:300056462`) in place of the style one.

[28] https://linked.art/model/object/aboutness/#style-classification

5.4.0.8 Citations, sources, evidences

As shown in Figure 5.2, the CiTO ontology is directly reused to represent the source (`cito:citesForInformation`) from which the Recognition is extracted, the evidence (`cito:citesAsEvidence`) on which it is based and the supporting (`cito:givesSupportTo`) between acts of recognition. This representation is fundamental to encourage a documented description of the recognition, its reference and support.

CHAPTER 6

The Iconology dataset

The creation of the dataset containing iconographical and iconological claims was conducted as a question-driven approach to answer domain-specific questions with a quantitative approach in the art history field. This approach, based on a direct, qualitative study of the literature, allows us to better define the important features of the domain. The dataset creation was structured in phases, namely: 1) study of the domain and definition of RQ, 2) ontological modeling, 3) database design, 4) database population, 5) cleansing, 6) conversion into RDF, 7) alignment, and 8) evaluation.

The study of the domain and ontological modeling were already treated in Chapter 5, whereas the evaluation will be addressed in Chapter 8. In this chapter, we will describe the remaining phases. Furthermore, we will present the modeling of the remaining items included in the metadata, such as artworks metadata, people, and books.

6.1 Sample selection and description

Panofsky's bibliography was gathered from the published bibliography by Previtali (2009) and by the Dictionary of Art Historians[1], for a total of 79 articles and books. The criteria for selecting the books that could be described were 1) the availability of an English translation, 2) the presence of artwork depiction, and 3) the focus on Middle Ages or Renaissance/Late Renaissance art. Out of them, 20 were excluded for the lack of a translation (6), lack of depictions or theoretical theme (12), and out of topic (2), obtaining 59 references. For answering the current study research questions, we focused on a selection of the most important books for iconological studies.

[1] https://arthistorians.info/

6.2 Database creation

The database for data population consists of a Google Spreadsheet divided into 1) tables for actual data entry, and 2) tables for controlled vocabulary to be used in the descriptive tables, progressively extended as soon as the database was populated with information.

6.2.1 Database modeling

The database modeling was developed on the basis of the characteristics of the 11 case studies in Baroncini et al. (2021) and of the ontology discussed in Chapter 4. Consequently, we can define the basic requirements as follows:

- The database should include the subdivision into the ICON recognition types, according to the levels
- The type of subject for each level should be specified
- It should be possible to register relations between subjects at different levels (e.g., second level subjects composed by a group of level 1 subjects)
- It should be possible to cite other recognitions
- Recognitions shall be grouped into interpretations, and further details about the interpretation can be recorded
- It should be possible to include divergent recognitions

Following these requirements, a first draft of the tables was created and tested with the described test cases. Therefore, a table for each interpretation type was created, plus one for artworks metadata description, one for people description (artists, interpreters), and multiple tables for the controlled lists of artwork type, style, subject entry, and cited texts. As a result, we obtained a scheme of tables expressing a theoretical framework of the features of iconographical and iconological interpretations that is relevant to describe, according to Panofsky's theory. The dataset population led to further improvements. A column and a controlled list of terms for relations between artworks were added, plus tables for descriptions of books and the text they express. Each time a new feature was added to the tables, the descriptive ontological framework was enriched to express it by reusing existing ontologies. In other words, each choice for the modeling of the database implied the definition of its formal conversion to RDF. For example, we provide two tables for the description of books, as we consider them composed of 1) a

physical part and 2) the text they carry, in compliance with the CIDOC-CRM modeling to which such information is converted in a second moment.

6.2.2 Tabular framework

Tables included in the tabluar framework have different roles. Whereas part of them are tables in which the actual content is collected, other tables consist of lists of controlled terms to uniquely identify objects included in the previous tables. Main tables include the description of artwork metadata, books and the text they carry, the recognitions of objects at each level of interpretation, and of symbols. The second type of table is needed mainly to uniquely identify the objects, create controlled lists for field completion in other tables, and indicate alignments with external sources if retrieved. We include tables with controlled lists for places, people, style, period, type of cultural object, and type of relationship between cultural objects. Furthermore, another table provides the controlled terms for the subjects described at each level and their possible alignment.

In this section, we briefly describe the entities described in the main tables. Some general rules apply to the tables. Each term is identified either by its name, a number, or an alphanumeric string, according to the cases. If multiple terms for the same field exist, they are separated by the symbol "@", which will be used as a separator during the RDF conversion to create seaprate entities for each term.

6.2.2.1 Table for Cultural Objects

In this table, the metadata description of the cultural object is provided. It includes details about the author, the uncertainty of the author attribution, the title, years or century of creation, material, type (e.g., relief, painting), the place and city of conservation, and the URI identifying the book in which it is described. Furthermore, if the artwork has a relation with another one (e.g., is part of, a copy or a drawing of another cultural object) or parts of the same object (e.g., a scene of a fresco), it is possible to insert such related object by indicating its ID and further specifying the desired relation. Other columns are dedicated to the possible URL of the image if it is retrieved, its source, image rights, and the presence of an existing URI for the cultural object described (e.g., a Wikidata ID). Furthermore, other columns indicate style, period, and material, in which only terms from the respective controlled lists specified in different tables are allowed. All these characteristics are summarized in Table 6.1.

Table 6.1: Field described in the table for cultural objects description

Feature	Description
ID	A number from 1000 to 9999 that identifies both the physical object and the respective artwork. Test cases, external artworks, and types have letters added, respectively, 'test' (eg. 1001test), 'ext' (e.g., 1001ext), and 'test' (e.g., 1001test)
Author	The name of the author(s) responsible for the creation of the artwork, that should be added to the list of controlled values in the table 'People'. The names should be written in the format <name> <surname>. If multiple authors occur, the symbol "@" is used as a separator, to model the authors as separate entities in the final graph
Certainty	Indication of the author's certainty. If the author's attribution is not certain, insert a '?'.
Title	The title of the artwork
Title for checking duplicates	A string containing, if present, the author, title, date, city, and location (e.g., 'The firmament, 1440, Florence, San Lorenzo'). It allows the detection of duplicate items during the database population phase, and it furnishes a label with a thorough description of the artwork.
Date	The date of the artwork creation. It can be expressed as an interval of years (e.g., 1450-1550), as a century (e.g., XII Century) or a portion of it (e.g., late, early, first quarter etc.)
Date certainty	Indication of the certainty of the date. If the attribution is not certain, insert a '?'.
Material	The material of the cultural object (e.g., marble). It must be included in the list of materials specified in the 'Type' table.
Book ISBN or DOI	The URL identifying the article or the book from which the interpretation about the cultural object is extracted. If it is discussed in multiple bibliographic records, all of them must be inserted here.
City of conservation	The city in which the cultural object is stored, which must be included in the controlled vocabulary of table 'Places'.
Place of Conservation	The location (e.g., church, museum, collection) in which the cultural object is stored, which must be included in the controlled vocabulary of table 'Places'.
Type of relation	If a relation with another cultural object occurs, this field specifies the type of such a relation (e.g., 'part of', 'study for', 'companion piece of'). The values should be inserted in the table 'TypeOfRelation'.
Related artwork	If the related object is an artwork, this field is filled with the descriptive label. The artwork must be described in the current table.
Related Book	If the related object is a book (e.g., a manuscript), this field is filled with the respective descriptive label. The book must be described in the 'Books' table.
IDext	If a related object is described, this field contains its ID.
Specific location	field for possible further indications about the portion occupied by the cultural object on the related object. For example, a relief on the East portal of a church will have a specific location 'East portal'.
Object Type	The type of the related object (e.g., 'building)
Style	The cultural object style. It must be included among the styles listed in table 'Style'.
Period	The cultural object period (e.g., 'Renaissance'). It must be included among the periods listed in table 'Style'.
Source Link	A possible URL at which the object is described
ImageLink	A possible URL of the object image
ImageLicense	the image license
Image Rights	If indicated, the image rights holder
URI	The URI identifying in an unique way the object in another knowledge graph (e.g., Wikidata)

6.2.2.2 Tables for Book description

The book description is divided into two tables: the table 'Book' for describing the material aspects (signature, editor or printer, date and place of production), and the table 'Book Content' for the description of the text carried by the material book, of which we only describe the author and title of the text. Whereas the Cultural Object table refers to the ID of the book's material aspect (e.g., when illumination is part of a manuscript), the Book content identifier is recalled in recognition tables when the text is cited as evidence of recognition.

6.2.2.3 Recognitions and interpretations

The content of the scholar's interpretation is described in five tables. Three of them describe the subjects recognized at each level, namely the first, the second, and the third. One table is dedicated to the description of the Visual Pattern citations, in which a relation between first-level objects belonging to two artworks is established. Another is dedicated to the thorough description of symbols.

In the case multiple interpretations of the same artwork occur (e.g. when various references describe an artwork or when the art historian cites interpretations by other scholars), such interpretations are described in the table 'Interpretations.' There, interpretations concerning the same artwork are progressively numbered. Such a number is also reported in each recognition, composing the interpretation in the remaining tables in the dedicated field. During conversion to RDF, the unique identifier of each recognition will include such a number. Besides the Artwork ID (e.g., 1001ART), further features can be described in the Interpretation table, namely: the date in which the interpretation was made, the author, and the source.

As introduced in Chapter 5, an interpretation is made up of single recognitions. In the remaining tables, every row identifies a single recognition. All the recognitions can be further specified by adding 1) the citation of evidence on which the recognition is based (e.g., recognition of the personification of abundance based on Ripa's *Iconologia*) and its specific part, if any (e.g., the exact verse of a poem), 2) the bibliographic reference from which the interpretation was extracted, 3) the person responsible for the recognition, and 4) the indication of another recognition that the current recognition supports (e.g., the recognition of a visual citation from a classic relief in a Middle Ages artwork supports the phenomenon of 'reuse of classical motifs deprived of their content').

In the following sections, we report the remaining descriptions of the recognitions at each level.

6.2.2.4 Table for First Level description

In the first-level description table, the pre-iconographical recognitions are described. According to Panofsky's theory, such subjects are Natural Elements, Actions, and Expressional Qualities, which may be gathered in compositions, i.e., groups of related subjects that may be further interpreted at the second level of interpretation. Therefore, such elements are the subjects described in this table (see Table 6.2). Compositions may be further described by a Compositional Structure, i.e., the visual arrangement that they present, such as a mathematical perspective or a pyramidal composition among three characters.

In the ICON ontology modeling, we further introduced the concept of qualities that the artwork may have, namely further description of the particular appearance of the subject as depicted in the specific artwork under examination (e.g., colors, cardinality, pose). During the description of artworks, it was necessary to introduce another field for describing the text included in the representational surface, such as the written names of the depicted characters (Linguistic Object). Furthermore, we introduced the field 'Style', to be used when a style is attributed to a single object depicted rather than to the whole cultural object (e.g., Roman columns depicted in a Renaissance painting).

A progressive counting in each artwork uniquely identifies the Artistic Motifs and Compositions described, respectively the subjects as manifested in the specific artwork and their eventual gathering. The relation between each composition and its members is established through the field 'Composed by Artistic Motif', in which the counting of the respective Artistic Motifs is reported.

6.2.2.5 Table for Visual Citations description

In this table, relations between first-level subjects or compositions between two artworks involved in a Visual Citation are established. It is done by indicating, for each recognition, the copying and copied artwork ID, and the count of the respective copying and copied Artistic Motif or Composition. The counting corresponds to the one expressed in the table for the first-level description previously desribed.

6.2.2.6 Table for Second Level description

Similarly to the First-Level Subjects table, the second-level one describes the level 2 subjects recognized and their possible specifications. Table 6.3 summarises the fields included. According to the ICON modeling, such subjects are Characters, Places, Events, Named Objects, Personifications, and Symbols, which can be grouped into Stories, Allegories, or Invenzioni. During the database population, it was necessary to introduce another column for the artworks depicted in the

Table 6.2: Fields described in the table for pre-iconographical recognitions

Feature	Description
Artwork	The artwork ID (e.g., ART1001)
Artwork Interpretation	If multiple interpretations are present, this field indicates the progressive number of the interpretation, further described in table 'interpretations'
Artistic Motif Count	This field contains a progressive number which uniquely identifies every subject recognized, either Artistic Motifs or Groups of them (i.e., Compositions)
Natural Element	The objects, human beings, and natural items that can be experienced which are depicted in the artwork
Expressional Quality	An Expression Quality corresponds to the 'expressional subject matter' described by Panofsky (1972). It individuates, by way of example, emotions, an atmosphere, and the character of a person.
Quality	Specification of the subject recognized, i.e., Natural Element, Action, or Expressional quality. Such specifications may be, for example, color, pose, and age.
Action	An event as recognized at level 1 of interpretation, e.g., a man who lifts one hand or a man riding a horse. Usually, actions are included in Compositions which include the other subjects performing the action.
Linguistic Object	The writings included in the artwork surface
Style	The style of a subject recognized (e.g., 'Roman' sculpture). It must be included among the styles listed in table 'Style'.
Composition	Compositions are groups of Artistic Motifs. In this field, a progressive number uniquely identifying the compositions in each artwork is specified.
Composed By Artistic Motif	In this field, the artistic motifs which are part of the composition are indicated through their number, assigned in field 'Artistic Motif Count'.
Compositional Structure	Description of the eventual shape (e.g., pyramidal), or an ordering principle (e.g., perspective) that characterizes the positions of the objects included in the composition.

artwork under examination, to be selected from a controlled vocabulary with all the Cultural Object IDs from the respective table. Each row (i.e., each recognition) recognizes either one of the single subjects or a collective one (i.e., Allegories, Stories, and Invenzioni). If a single subject is recognized as part of a collective subject in the painting, this is stated by indicating such collective subject in the respective field. The recognition of the collective subject is also stated by describing it in a new line to allow for the description of the provenance assertion (e.g., the person responsible and evidence, if any). The eventual correspondence with level 1 recognitions is provided by indicating in the fields Artistic Motif and Composition the progressive number identifying them, assigned in the table for first-level descriptions. Furthermore, if one or more of them had a relevant role in the iconography identification, they are indicated in the field 'Recognizing Attribute'. During the data population, we introduced the fields of 'Role' and 'Source Actor', to better describe the situations in which a second-level subject was depicted in a specific role and when a Character had the facial traits of an existing person.

Table 6.3: Fields described in the table for iconographical recognitions

Feature	**Description**
Artwork	The artwork ID (e.g., ART1001)
Artwork Interpretation	If multiple interpretations are present, this field indicates the progressive number of the interpretation, further described in table 'interpretations'
Count	This field contains a progressive number which uniquely identifies every subject recognized, either Images, when a single subject (namely Characters, Places, Events, Named Objects, Personifications, Symbols) is identified, or Invenzioni, when a collective subject (namely Allegories, Stories or Invenzioni) is identified. or Groups of them (i.e., Compositions). As each recognition identifies only one Image or Invenzione, the counting identifies also the recognition.
ArtisticMotif	In this field, the progressive number identifying an Artistic Motif, previously defined in the First Level description table, is indicated, if the recognized iconography refers to it. For example, if a book recognized in ART1001, identified in the first level as a Natural Element, with Artistic Motif count 2, is then recognized as the Named Object *Bible*, the artistic motif number 2 will be indicated in this field during the iconographical recognition.
Composition	Similarly to the field of artistic motifs, an iconographical recognition may refer to a Composition of Artistic Motifs rather than singular motifs. In this case, the counting identifying the Composition is reported in this field.
Recognizing Attribute	If one or more artistic motifs had a relevant role in the recognition of the iconography (e.g., the arch and arrows help the identification of Cupid), their identifying number is reported in this field.
Character	The second-level subjects recognized as Characters according to the ICON ontology, namely a real or imaginary being identifiable as a second-level subject
Role	the role in which the character is depicted in the specific artwork. For example, in ART1199 Saint Luke is depicted in the role of a painter.
Source Actor	This field is filled when the recognized character bears also a portrait of a real human being. The source actor must be inserted in the table 'People', from which a controlled list for the current field is created. For example, the character of the Young Mago recognized in ART1003test bears the portrait of a young Lorenzo de' Medici.
Place	This field concerns those subjects representing places with a precise identity, e.g., the city of Bologna.
Event	This field concerns those events with a precise identity, such as the battle of Cascina
Named Object	This field concerns those objects with a precise identity, such as the Bible or the Venus' belt.
External Objects	In this field, we include recognizable cultural objects, such as artworks depicted in the artwork under examination. For example, the drawings by the art historian Popp shows how certain Michelangelo's sculptures (registered as the External Object) were arranged in a unique project never realized.
Symbol	In this field subjects invested with a symbolical meaning are indicated. The description rule adopted consists of writing the symbol and the symbolical meaning separated by a high dash without spaces (e.g., 'soap bubbles-futility').
Personification	subjects in which a person is depicted to illustrate a concept (e.g., Personification of Charity).
Story	A combination of subjects representing narrative elements (e.g., the story of Hercules)
Allegory	A combination of subjects invested with symbolical meanings. Usually, at least a symbol or a personification should be involved (e.g., The abduction of Europa as an allegory of Christ saving the human soul).
Invenzione	A combination of subjects with presents the characteristics both of a Story and of an Allegory.

6.2.2.7 Table for Symbols description

Symbols recognized in the Second-Levels Subjecs table are further described in this table according to the Simulation ontology. Each symbol, called *Simulation* in the ontology terminology, is split into 1) the symbol depicted in the artwork, called *Simulacrum*, and the symbolic meaning to which it refers, called *Reality Counterpart*. Further details may be provided, such as the type of Simulation which describes in a more thorough way what is the relationship between the symbol and its symbolical meaning (e.g., if it is an allusion to the meaning), or the Context in which such Simulation is used (e.g., Medieval or Renaissance). We enriched the description by providing a field for indicating the source in which the Simulation is found, such as classical poems in which an attribute of a deity is invested with symbolical meaning (e.g., the Venus's belt is a symbol of seduction in Homerus' Iliad).

6.2.2.8 Table for Third Level Description

Similarly to the previous levels, the iconological recognitions are described in the rows of this table, and a progressive number for each artwork identifies the recognized Intrinsic Meanings. The subjects related to it can be either concepts or Cultural Phenomena, according to the ICON ontology. The fields of Artistic Motif, Composition, and Image are used to indicate to which subject recognition the recognized intrinsic meaning refers if provided. For example, the recognition of the phenomenon 'Representation of classical content with contemporary formal motifs' in ART1142 refers specifically to Virgil, identified by Image 1.

6.2.3 Description rules and database population

Out of the selected books, we decided to include only the artworks having a picture in the book edition, to reduce the likelihood of erroneously selecting the artwork and to foster the retrieval of a digital image of the artwork itself. In certain cases, the cited artwork without an image was included because the art historian was providing a significant description of it, or the artwork was of crucial importance for proving his claim (e.g., the presence of a visual citation of the artwork in question). The chapter of the book (Panofsky, 1955) describing the activity of Suger, the abbot of Saint-Denis, was not included since no artworks were cited in it.

6.2.4 Database cleansing

For this study, the quality assessment was partially conducted on the tabular database. Processes for data cleansing and expression harmonization were adopted during and at the end of the manual insertion.

During the data insertion, values were progressively inserted in controlled lists, one for each type of entity inserted (e.g., people, Natural Elements, etc.). This process was fundamental for remaining as relevant as possible to the language used by Panofsky without penalizing the use of a common vocabulary. Indeed, the use of drop-down menus of already inserted controlled lists enhanced the reuse of the already inserted terminology, if appropriate, during the data entry phase. Following this, a strategy for entity disambiguation was adopted. Since artworks could be described in multiple books, it was important to detect already described ones and use a unique identifier for those artworks appearing multiple times. To this end, a conditional coloring of the cells in the column for artwork description (including an indication of author, artwork title, date, and place of conservation) was adopted. The resulting cases were manually examined and eventually merged.

At the end of the insertion, a spelling check, a singular-plural forms harmonization, and resolution of uncertainty in inserted data (e.g., deciding if 'angel' is a second or first-level subject) and in expression variation (e.g., if saying 'laying pose' rather than 'reclined position') were performed. In particular, the spelling check was 1) manually conducted by reading the controlled lists in alphabetical order for a better comparison of word variations, and 2) passing it through an English spelling check tool. Through the same technique, the use of singular and plurals was detected, and data were normalized by using the singular form of the term, accompanied by the indication of 'multiple quantities' as a qualifier of the first-level object, as it is the level for describing the actual appearance of identified subjects. Through a Python script for finding similar terms, ambiguity in data was checked. For example, it emerged that the term 'angel' was indiscriminately used both at the first and second levels of description. As it is a narrowly cultural-related entity, it was decided that it must be described at the second level only.

Before RDF data conversion, it was verified whether all the terms in the tables were included in the controlled vocabularies. If not, the single cases identified through the script were manually corrected in the source data.

6.2.4.1 Rules for normalization

To avoid mistakes in the data conversion process, rules for data entering were decided. This determined, for example, the way to write dates, uncertainty, and places. Furthermore, terms were translated into English, and old terms

used by Panofsky were updated to the current use (e.g., we wrote the current preferred name 'Angelo Bronzino' instead of 'Agnolo Bronzino'), or corrected (e.g., the Pennysilvania museum reported by Panofsky is currently called Philadelphia Museum).

6.3 RDF data creation

6.3.1 Ontological modeling

In this section, we provide the modeling adopted for describing the various aspects of the objects included in the dataset. For the artwork content description according to the levels, we directly reused the ICON modeling described in Chapter 5. For the other aspects, we tried to reuse as much as possible the CIDOC-CRM standard (version 7.1). Other ontologies were used for more specific tasks, namely: CiTO[2] for describing citations, PRO[3] for describing people temporary roles, and the foundational ontology DOLCE[4] for first-level objects details.

6.3.2 Main objects modeling

Artwork The artwork is represented both in its physical (`crm:E22_Human-Made_Object`) and visual (`icon:Artwork`) aspects, related by the property `crm:P65_shows_visual_item`. In the case in which the artwork is a building, (e.g., a church), we use the class `crm:E24_Physical_Human-Made_Thing` in spite of `crm:E22`. All the metadata about the artwork is related to the physical entity, according to CIDOC-CRM specifications, including the artwork's label and title. For the description of the content of the artwork according to the levels, we directly reused the ICON modeling described in Chapter 5, with little variation. Since the reference standard we adopted is CIDOC-CRM, we took some decisions to better integrate the ontological description in this context:

- The person responsible for each recognition is a `crm:E21_Person`, linked to the recognition class through the property `crm:P14_carried-out_by`.
- Some recognized subjects were further specified with CIDOC entities. It is the case of the written text represented in the artwork (e.g., the names of the depicted characters), which are registered as level 1 subjects but with type `crm:E33_Linguistic_Object`, following the guidelines provided

[2] https://sparontologies.github.io/cito/current/cito.html
[3] http://www.sparontologies.net/ontologies/pro
[4] http://ontologydesignpatterns.org/wiki/Ontology:DOLCE+DnS_Ultralite

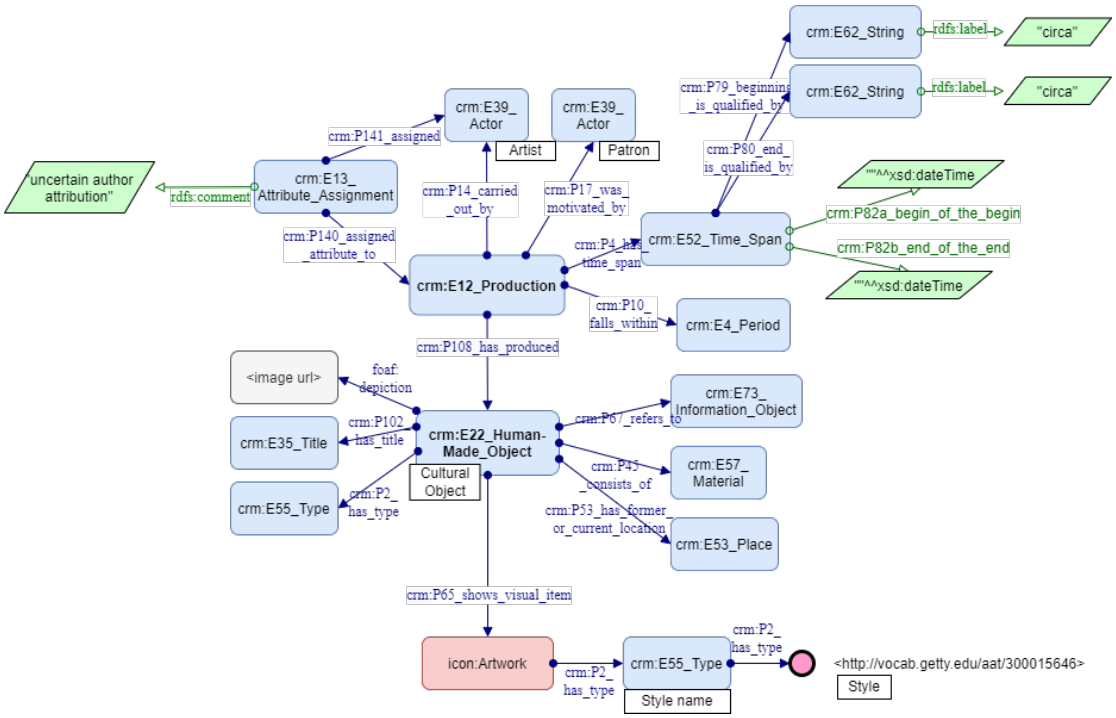

Figure 6.1: Modeling of artworks

by the Linked.art project[5]. Similarly, the artwork depicted into another artwork was modeled as follows. An image is recognized at the second level of interpretation, and it is linked with the relation `crm:P138_represents` to the entity identifying the painting, possibly further described in the dataset. This modeling was adopted for three artworks, namely ART1373, ART1364, and ART1365, respectively representing OBJ1255ext and OBJ1256ext, OBJ1374, and OBJ1375.

- The recognized third-level concepts were expressed as `crm:E28_Conceptual_Object`, in order to maintain the general rule of using the standard, when possible.

The artworks metadata that we decided to describe are 1) its production act, with the further specifications of the author (`crm:P14_carried_out_by`), the possible patron (`crm:P17_was_motivated_by`), the dates of creation through the expression of a time-span (`crm:E52_Time_Span`), the period (e.g., Middle Ages) related through the property `crm:P10_falls_within`, 2) title (`crm:P102_has_title`), 3) material of which the object is composed, such as marble (`crm:E57_Material`, 4) its location, i.e., the museum or place in which it is preserved (`crm:P53_has_former_or_current_location`), and the type (e.g., relief, oil painting, etc), represented as a `crm:E55_Type`, related by the property `crm:P2_has_type`. Further features shown

[5] https://linked.art/model/object/aboutness/

in Figure 6.1 will be described in the remaining paragraphs.

Person The modeling of people regards the artists, patrons, portrayed characters, books authors, or editors, identified by the class `crm:E21_Person`. Sometimes, the artwork's author is not a single person, but a group, such as a school, and therefore represented as `crm:E39_Group`. In certain cases, to better document the final interpretation, it was relevant to register further details, such as the role held by the person or the membership in a group. Whereas the former was expressed through PRO ontology class `pro:RoleInTime`, further described with the required time indication, the latter was described through the property crm:P107_has_current_or_former_member, relating the group to the single person. The specified role is independent of the role that the person may have when depicted in an artwork (e.g., Vulcan depicted in the role of Teacher of the Mankind). We model it by reusing PRO classes and relations as from our data emerged the necessity to describe them as temporarily assumed roles (e.g., John VIII Palaiologos, Byzantine Emperor from 1425 to 1448), represented in the Benozzo's frescoes in the Medici Chapel, used as a test case for dataset creation and modeling.

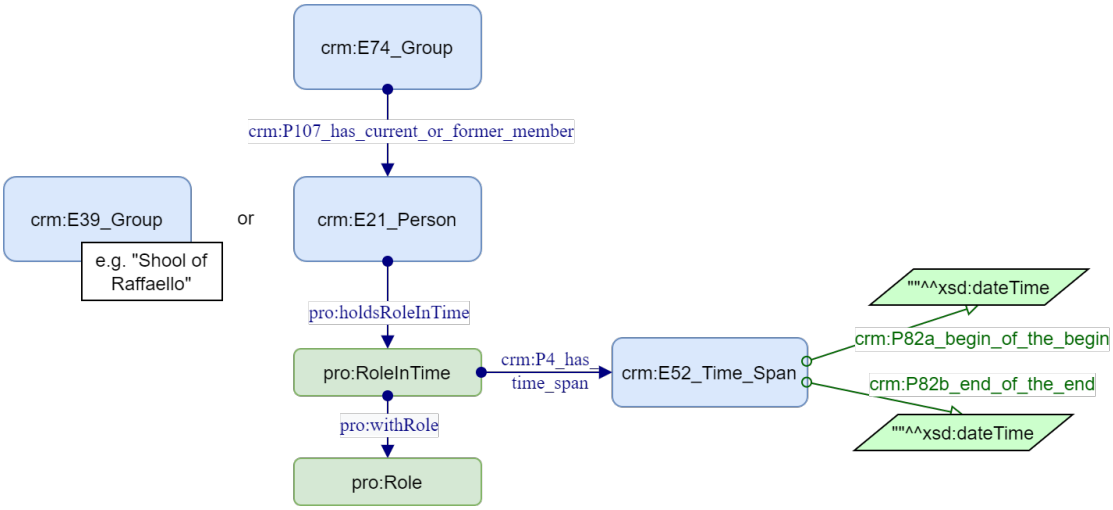

Figure 6.2: Modeling of people

Book The relevance of books was progressively understood and therefore iteratively modeled during the database population process. Indeed, books have different roles in the described dataset. They can be the manuscript or book of drawings in which the artwork is depicted, or the book carrying the text cited as a piece of evidence during an interpretation. Since, in some cases, the interpretation was citing the text in which the artwork under observation is placed, a modeling framework defining clearly the respective identity of the physical book and the text, along with their relations between each other, with the artwork, and with the interpretations was needed. For this purpose, we adopted the modeling

provided by CIDOC-CRM (see Figure 6.3). The book in its physical meaning is represented by an instance of `crm:E22_Human-Made_Object`. Similarly to the artworks' metadata description, the physical book can be further specified with its production details, along with the date (or period) and place of creation, the person responsible (in this case, with the role of editor or printer), and an identifier (`crm:E42_Identifier`, through the relation `crm:P1_is_identified_by`, used, for example, to represent the manuscript's signature. The text written in the book is described as an instance of `crm:E73_Information_Object` and related through the property `crm:P128_carries`. The illumination in its physical meaning is identified as an instance of the class `crm:E25_Human-Made_Feature`, which can be part of a folio or page, identified as another instance of the same class. Both the instances are linked to the book through the relation `crm:P56i_is_found_on`.[6]

As shown in Figure 6.4, further details can be provided about the text carried by the book. If an author is given, he/she is specified as the author of the creation event (`crm:E65_Creation`) of the text. The title is specified through the relation `crm:P102_has_title`. If, during a recognition, a specific part of the text was cited as evidence, we created a new `crm:E73_Information_Object` entity to represent only the text portion. The whole text is linked to the latter by the relation `crm:P106_is_composed_of`.

Places In the dataset, places are generally the locations in which the artwork is stored. Nevertheless, we distinguished between those places as physical locations (`crm:E53_Place`) and the museum or collection respectively as a legal body and as a physical entity (Fig. 6.5). This specification was made to indicate museums and collections as the places of the artworks without violating the range of the property `crm:P53_has_former_or_current_location`, which needs an entity of type `E53`. Furthermore, the localization into a wider region, e.g., a city or a country, was registered by the relation `crm:P89_falls_within`.

Style Following the CIDOC-CRM and Linked.art indications, style is described as a type (`E55`) of the visual entity `icon:Artwork`, having as type the AAT term for *style* (see Fig 6.1). The modeling, already reused in the ICON ontology development, is also adopted in data creation. We adopted this pattern also for describing the style of single level 1 subjects depicted by the artwork (e.g., a classical column in a Renaissance setting).

[6]Although we initially referred to the study by Bellotto (2020) for manuscript description through CIDOC-CRM, we do not directly reuse it as we encountered some inaccuracies in the use of properties. The author of the study describes the manuscript as a `crm:E71_Human-Made_Thing`. The property `crm:P165_incorporates` relates the physical manuscript (E71) to the text (E73), although the domain of this property, as specified in CIDOC-CRM documentation, is E73. A similar inaccuracy is found in the use of the relation `crm:P10_falls_within`, which must relate two temporal entities (E4), to describe the relation between the illumination and the book.

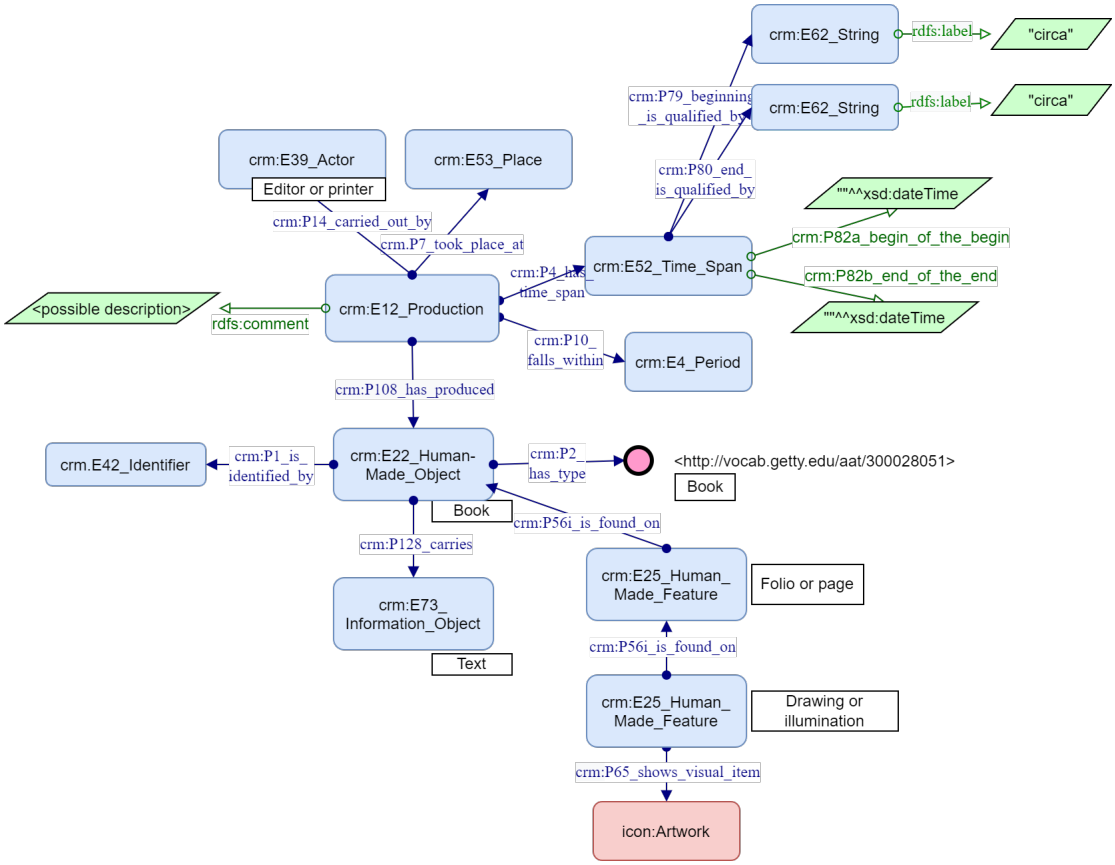

Figure 6.3: Modeling of books

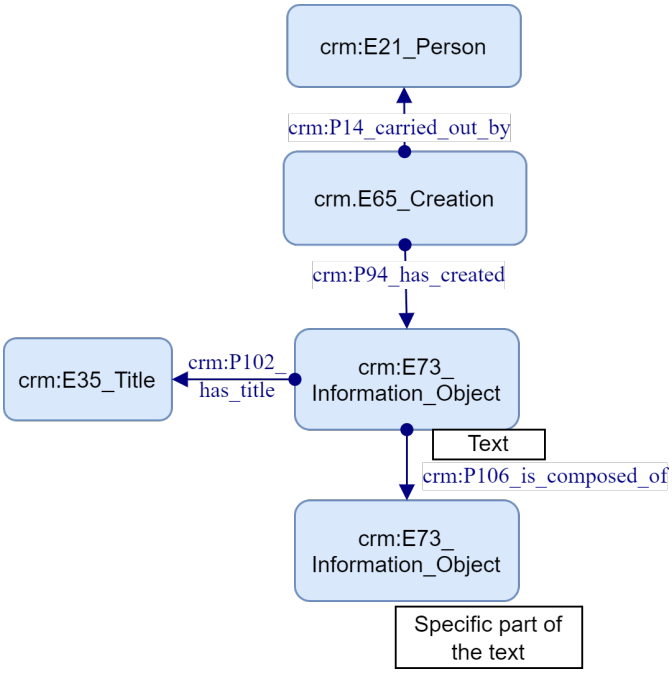

Figure 6.4: Modeling of books content

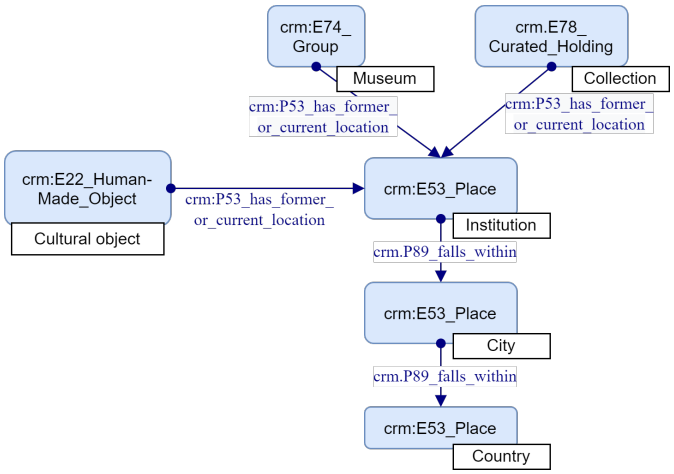

Figure 6.5: Modeling of places

6.3.3 Modeling challenges

The modeling phase, the table structure, and the conversion were iterative processes in which the gradual insertion of artworks led to improving the structure, although the core modeling could be set thanks to the test case studies. In the following, we describe some features that we included.

6.3.3.1 Artwork specifications

The major part of the type of artwork could be described according to the modeling cited above. Nevertheless, some specific types of artwork require specific modeling.

Feature, detail, or parthood During the modeling, we had to decide how to consider the relation between a portion of an artwork and the whole. We identified different typologies:

1. The art historian examines only a portion of a whole artwork, presenting also a picture of the needed detail only. The portion is not a separate entity of the artwork (e.g., the detail of one painting)

2. The art historian examines only a portion of a whole artwork, but the artwork has a complex structure divided into multiple scenes. It is the case, for example, of many frescoes or pictorial cycles, such as Raffaello's frescoes in the Loggia of Psyche (Villa Farnesina, Rome), or Michelangelo's decoration of the vault in the Sistine Chapel.

3. The apparently unique artwork is composed of parts made by different artists in different periods. We provide as an example the case described by Panofsky

(1955) of Vasari's frames surrounding already existing Medieval drawings (OBJ1427, OBJ1428, OBJ1429).

4. The artwork is part of another one having a different type. It is the case, for example, of a fresco (e.g., the cited above Raffaello's *Stories of Cupid and Psyche*) and the physical architecture (Loggia of Cupid and Psyche, Villa Farnesina, Rome). Although sharing the same surface, the two artworks are created in separate moments, eventually by different people. Another example belonging to this case is the creation of a miniature and the physical manuscript page on which it is represented.

5. The artwork identity can be recognized, it may be shown independently but it is conceptually or structurally part of a bigger one. The pales of a polyptych, the stone throne of a church, and the card of a deck, belong to this category.

After the identification of different typologies of relations between artworks, we modeled the illustrated cases as follows. In the first case, since the portion of the artwork exists only in relation to the art historian's interpretation, the artwork has a unique identity. To not diverge from Panofsky's reading, we maintained both an independent visual artwork and object when a single interpretation of it was provided, but when more parts of the same artwork were described on different occasions, the artworks ID were merged in a unique one, registering the different interpretations for the same artwork. It is the case, for example, of ART1123. Whereas in *Renaissance and Renaissance in Western Art* only Venus is described, in *Studies in Iconology* the focus is solely on Cupid.

In the second case, since the portion of the artwork can be identified, but is physically and conceptually part of a whole artwork, it is registered as a physical feature (`crm:E25`) related to the main artwork through the relation `crm:P56_bears_feature` (see Figure 6.6). In this case, we have only one production instance associated with the whole artwork and separate instances of visual artworks.

We provide the same modeling for the third and fourth cases (Fig. 6.7). As in case 2, the artistic objects are described as features. Nevertheless, since they exist due to different productions, a new instance of `crm:E12_Production` is associated with the feature. Also in this case, each physical instance (i.e., the Feature and the Human-made Object) has an instance of visual artwork associated.

The fifth case is expressed as two instances of Human-made Object (`crm:E22`), related by the property `crm:P46i_forms_part_of`, and one instance of Artwork is associated to each of them.

Pendants Artworks may be recognized to be narrowly related to other ones.

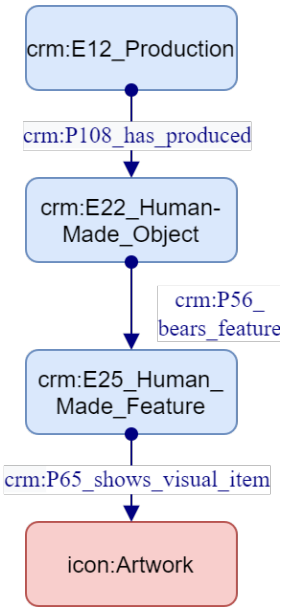

Figure 6.6: Detail of (feature-of relation)

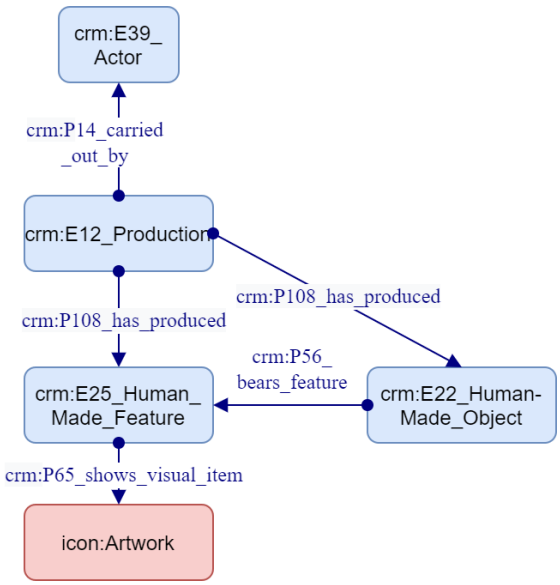

Figure 6.7: Feature of relation

According to the Getty Vocabulary definition, artworks which are 'designed, or purposefully chosen, to be displayed together'[7] are defined as *pendants*. We model a collective object representing the union of the companion pieces as an instance of `crm:E22`, as it is the recommended class for representing collective objects[8], and we then relate the two instances of physical artworks (E22) through the property `crm:P46_is_composed_of`. The aggregate object has type `aat:300179422`, identifying the typology. Four artworks of the dataset correspond to this case, identified as two couples of artworks (OBJ1522 and OBJ1468, OBJ1283 and OBJ1284). In both cases, the companion pieces are examined to better understand deeper meanings in the interpreted artwork.

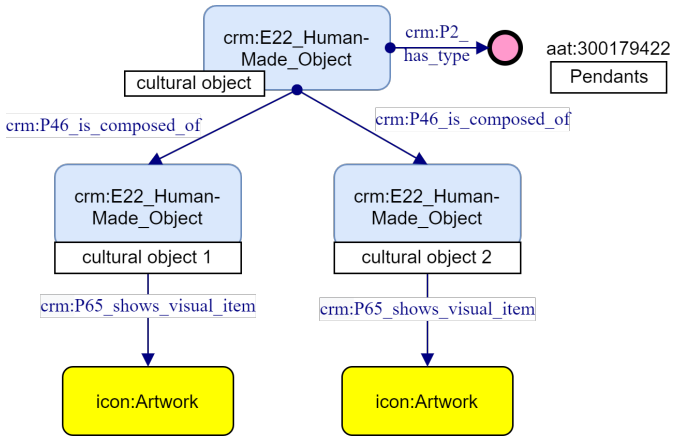

Figure 6.8: Pendant of relation

Specific location In the metadata of the described artworks further information about their specific position on another cultural object was included. For example, it was indicated that the fresco under observation is placed on the dome of a building or in the east portal of a church. Another case is the indication of the exact page of the manuscript on which the illumination is placed. In these cases, we modeled the specific location as an entity of type `crm:E25_Feature`, which is further related to the instance of the final object (e.g., the building or the manuscript) through the relation `crm:P56i_is_found_on`, expressing that the feature is borne by the final object. We modeled the chain of relations through the property `crm:P56i_is_found_on`, which indicates that the artwork is found in the specific location. Another `crm:P56i` property relates the cultural object to the specific location. This modeling is often used to express the relation between illuminations, illustrations, and reliefs, and the specific part of the whole object,

[7] http://vocab.getty.edu/page/aat/300179422
[8] https://www.cidoc-crm.org/Entity/e22-human-made-object/version-7.1.1

respectively a manuscript or a book page, or a building part, and finally to the whole object itself.

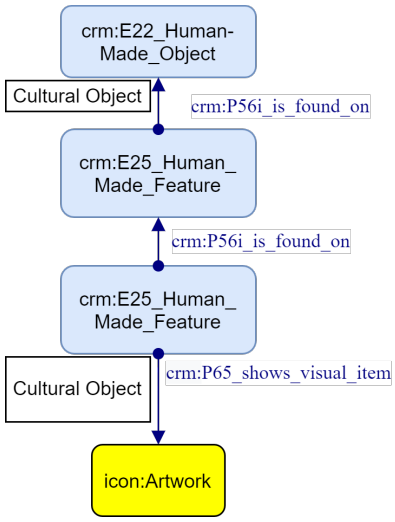

Figure 6.9: Modeling of specific location

6.3.3.2 Uncertainty

The modeling of uncertainty is a widely discussed topic in the Linked Data for Cultural Heritage community (Costa et al., 2010), as it is a concept often embedded in cultural heritage data. In this study, we decided to model the uncertainty encountered as comments expressing this fact in the registered data.

Uncertainty of author The author attribution may be uncertain, information that is usually expressed as placing a question mark after the author's name in the artwork title. We modeled this uncertainty aspect by expressing the relation between the production event and the author as an Attribute Assignment (`crm:E13`) linking the two entities and adding to it a a comment through the relation `rdfs:comment` specifying the uncertainty with the string 'uncertain author attribution'.

Date uncertainty Multiple types of uncertainty relative to the date exist. One possible case is that the interpreter is unsure of the exact year of artwork creation, usually indicated by the formula 'c.' before the year. We expressed this case through the addition of a qualifier, i.e., the string 'circa', to the E52_Time-Span entity, as shown in Figure 6.1. To the author's best knowledge, there is no term to express uncertainty in the currently available controlled vocabularies. The uncertainty can be distinguished for the starting and ending date, by choosing the suitable relation among `crm:P79_beginning_is_qualified_by` and

`crm:80_end_is_qualified_by`, respectively for uncertainty in the beginning and end date.

In other cases, date uncertainty is expressed with the terms 'ante' or 'before', to indicate that the artwork production necessarily happened before one specific date, and with 'post' or 'after', for indicating a production happened after a certain date. In those cases, only the available date, namely the start or the end of the assumed time period, is registered. Finally, it is common, in the absence of sources documenting the exact date of execution, that the creation date is indicated as a range between years or as portions of a century. We expressed the earliest and latest indicated dates as the extremes of such a time span.

Uncertainty in recognition Sometimes the author expresses uncertainty over certain observed facts. This happened in a couple of cases, in which the art historian recognizes a first level object but she/he is uncertain about the precise identification of it. In describing Lotto's *Venus and Cupid* (ART1002test), Christiansen recognizes a wreath, made of laurel or myrtle (ART1002test). Furthermore, Panofsky, while describing Titian's *Sacred and Profane Love* (ART1329), recognizes two animals, hares or rabbits, in both cases symbols of animal love and fertility (Panofsky, 1972, p. 150). Since in both cases, the choice among alternatives was not the discriminant for further interpretations, we registered only one of the options, preferably the most probable.

6.3.3.3 Artwork content

In this section, we provide an overview of the challenges raised in the description of artwork content. They range from the description of style and visual arrangements to figures embedding the characteristics of multiple iconographies and the presence of meanings related to collective objects.

Style of objects represented In some cases, the art historian observed a specific style not of the overall painting but of an object depicted in it (for example, the presence of classical columns in a Renaissance painting). The style modeling adopted in the ICON ontology (see Section 5.4.0.7) can be applied to every part of the visual surface of the artwork and, consequently, to each subject recognized as manifested in the artwork. In this way, it is possible to state that a subject has a specific style without asserting that all the artwork belongs to the same style.

Meaningful visual arrangement of subjects The perspective or visual composition in which subjects are arranged can lead to a deeper understanding of their meaning. It happens, for example, in the description of two representations

of the planets and their children (Bodl. ms. or. 133. (Arabic), XIV Century. *The seven planets and their children*, folio 25b, ART1031; *The children of Saturn*, 1420, Padua, Palazzo della Ragione, ART1032). While describing them, Panofsky says that the visual arrangement in a scheme bears a scientific attitude characteristic of the Arabians illustrations, afterwards abandoned in the Northern countries. In this case, we described the grid scheme as the `icon:CompositionalStructure` of the first level subjects and cited the Artistic Motifs involved in the visual arrangement as narrowly related to the phenomenon recognition, through the relation `icon:hasArtisticMotif`. Although it is challenging to model the cultural phenomenon recognition as directly based on the visual arrangement, we implicitly linked the fact by the relation with the recognized first-level objects.

Reference to iconographical types A few times Panofsky himself was conducting real artworks interpretations according to the knowledge of an 'art type' (e.g., Zeus with certain attributes). This approach is a challenge to the adopted modeling, since the recognized subjects are always dependent on a specific artwork. The solution adopted consists in the creation of an artwork record to which the type is associated. The ID of the artwork includes the string 'type' to specify that it does not correspond to any specific work of art.

Characters depicted in the guise of other ones In Western art, it is frequent to find subjects that express two characters. Mythological portraits, i.e., portraits of real people represented as mythological characters, are included in this typology. In other cases, they can be represented simply as characters part of the scene, as it happens in Benozzo Gozzoli's frescoes of the Medici chapel (`OBJ1003test`), where the three Magi and other participants bear the portraits of historical characters (e.g., Lorenzo de' Medici). Another case is when two iconographies are overlapped, e.g., between a character and an iconographical type. This can be seen in Lotto's Venus and Cupid (`OBJ1002test`), where Cupid is depicted with the attributes of the iconographical type of *puer mingens*.

We modeled this special case with the ICON ontology by relating two Characters to the same Image, through the relation `icon:hasCharacter`.

Complex meanings deriving from a set of artworks As iconology often explores the meaning of complex objects, it is often the case that the discovered meaning belongs to a set of artworks considered in their whole and taking into account their relations. It is the case, for example, of the interpretation of Correggio's decoration of the St. Paul chamber in Parma (OBJ1507)(Panofsky, 1961) or of the Medici Chapels in the San Lorenzo Basilica (OBJ1212ext) (Panofsky, 1972).

The solution adopted consisted in providing a description of the overall object only at the deeper level, and specifying level 1 and 2 descriptions for separate artworks identifying the single part (e.g., the walls of the chamber).

6.3.3.4 The concept of evolution

The domain of iconology and iconography addresses the variation of iconography, symbols, and cultural phenomena over time. This analysis includes the variation of what is depicted and ideas and texts on which the artworks may be based. Therefore, in the considered data, we could observe the following nuances of the concept of evolution:

- the evolution of iconography and symbols. The art historian notes variations in how the iconography is depicted and its symbolic meaning in the considered artwork. Consequently, s/he interprets the various visual evidence in which the variation occurs and analyzes it by assessing their spatial and temporal coordinates

- the evolution of ideas. The artwork's specific realization is evidence of an external change in society, such as a new identity, taste, religious practice, or idea

- the evolution of the text on which the depiction is based. In some cases, not only the actual notion expressed by the text or its interpretation changes. Especially when dealing with handwritten texts, words may vary and generate divergent visual interpretations of their content. An example is the illustrated rendering of Petrus Berchorius *Metamorphosis Ovidiana*. Several illuminations of the manuscripts of the tradition of its text present variations motivated by errors in copying the text. In this way, for example, Venus emerging from the sea is represented in ART1121 (ms. fr. 373, folio 207, Bibliotheque Nationale de France, Paris) with a sea goose rather than a sea shell due to the error in the text of Berchorius, in which 'aucam marinam' is written in place of 'concam marinam' (Panofsky, 1969, p. 86).

- the evolution of the artwork itself. When an artwork is restored or damaged, the actual content depicted may vary. It is the case of Lorenzo Lotto's *Venus and Cupid*, in which the restoration unveiled some precious details adorning Venus.

These characteristics make modeling the concept of evolution embedded in various aspects previously mentioned challenging. Several studies have attempted to address the complexity of the topic at hand. One such study, conducted by

Van den Heuvel and Zamborlini (2021), proposes a method for modeling time as *storylines*, defined as multidimensional networks that change over time. This concept is based on Kubler's theory of time (Kubler, 2008) and provides valuable insights. However, so far, no ontology has been released. Therefore, we plan to address this issue in future work, and we propose alternative solutions for the current project.

This thesis addresses the concept of evolution through the perspective of material artifacts (e.g., artworks, manuscripts). In other words, each artwork constitutes evidence of a different stage of evolution, whether of iconographies or cultural phenomena, in accordance with the definition provided by Panofsky of artworks as *documents of* a phenomenon (Panofsky, 1955). Accordingly, the time of the change is provided by the artwork's date. Furthermore, different aspects of the same cultural phenomenon are represented by a new cultural phenomenon. For example, the phenomenon *evolution of the iconography of Saturn* may be related to other specific aspects of it, such as *The accepted type of Saturn in late medieval art emphasizes his cruelty: Saturn devouring his child, Saturn castrate*. The relation is established by ascribing both phenomena to the same artwork in which the specific variation of the phenomenon occurs. In this way, a periodization of the phenomena can be provided on the basis of the dating of each piece of evidence. Nevertheless, this modeling choice has some limits. The overall argumentative discourse of the art historian, tracing a line of evolution, is currently not explicitly modeled in data. However, the evolution of iconography, symbols, and concepts can still be traced through a quantitative analysis that relates the subject to the dating of artworks.

6.3.4 Conversion into RDF

The conversion from the tabular format to an RDF graph was performed through a Python script realized for the purpose. The library used is RDFlib[9], a library for parsing, creating and updating RDF graphs. Through the script, the information in each table column was accessed and converted according to the required modeling for each object described in Section 6.3.1. During this phase, the instances included in the modeling but not explicitly described were created (e.g., the instantiation of a `Production` entity for each Cultural Object). Accordingly, rules for creating URIs were established.

6.3.4.1 URI creation rules

To perform entity disambiguation, each identified instance must have a unique and specific URI. In the case of this dataset, we opt for using 'non-opaque' URI,

[9] https://rdflib.readthedocs.io/en/stable/

namely, it should be possible for a human reader to easily understand the described content without recurring to the label. Although this facilitates the usability and comprehension of data from humans, it requires precise URI creation rules to be adopted during the conversion and to verify errors in URI assignment (e.g., the same URI is used for two different entities). This could happen, for example, if a string is used to describe two entities of a different type (e.g., 'Cupid' as a character and as a Story). For this reason, we set URI creation rules and a final RDF data verification through SPARQL queries.

We adopted different URI creation rules, namely 1) creation of non-opaque URIs composed by entity labels, 2) non-opaque URI with further specification of their type, and 3) identifiers composed by a progressive numeration and a string identifying their nature. In some cases, a composition of creation rules was adopted.

Non-opaque URIs In the cases in which ambiguity was impossible to take place, and there was a need to make it easier to refer to the same entity without erroneously creating multiple URIs for the same entity, we created non-opaque URIs from the objects label. The URI was created starting from the instance label. Spaces and punctuation marks were replaced by a dash, letters were written in lowercase, and non-UTF-8 characters were converted to the respective ones (e.g., 'è' was converted to 'e').

This approach was adopted for the URI creation of the major part of identified subjects.

Non opaque URIs with type specification For disambiguation between entity types, in certain cases, we included the entity type in the URL, following the format `<d:+type+'/'+name_uri>`, where 'type' is a lowercase string identifying the type, and 'name_uri' is the non-opaque URI created with the rules described above. Table 6.4 provides an overview or entities following this rule along with real examples extracted from the dataset.

In the case of conversion of places, a further information about the city in which a building is placed is provided in the place name (see 'berlin' in the provided example in Table 6.4), to avoid that homonymous places are merged.

Identifiers composed by a string and a number The identification through progressive numbers is preferred for identifying those objects that are rarely repeated during the dataset population, such as books and cultural objects. Indeed, the search for duplicates was performed, and the respective IDs were merged after a manual verification of similar entities. Other entities for which an identifier was created in this way are: Artworks, Texts, Recognitions, Interpretations, Production

Table 6.4: Entities for which the type is indicated in the URI

Type of object	**String for type specification**	**Example**
Style	'style'	`d:style/classicizing`
Period	'period'	`d:period/florentine-renaissance`
Type of Relation	'relation'	`d:relation/project-for`
Cultural Object Type	'type'	`d:type/oil-on-wood`
Material	'material'	`d:material/stained-glass`
People	'person'	`d:people/ae-popp`
Quality	'quality'	`d:quality/standing-pose`
Linguistic Object (E33)	'linguistic-obj'	`d:linguistic-obj/preserpine`
Story	'story'	`d:story/the-abduction-of-europa`
Allegory	'allegory'	`d:allegory/allegory-of-good-counsel`
Invenzione	'inv'	`d:inv/the-advent-of-the-celestial-venus`
Personification	'pers'	`d:pers/divine-justice`
Role	'role'	`d:role/teacher-of-mankind`
Place	'place'	`d:place/berlin-staatliche-museen`

and Creation Events, manifested objects (i.e., Artistic Motifs, Compositions, Images, Intrinsic Meanings), and Cultural Phenomena. In these cases, the identifier is composed of an incremental number associated with a string expressing the type of entity described (e.g., an artwork). In fact, associating an ID during the insertion phase helps in maintaining separate entities for the description of homonymous instances having the same type, as happened in the Cultural Object or Books field. Consequently, all the entities related to them (e.g., visual Artwork, Creation or Production events, manifested objects, Recognitions, and Interpretations) follow the same rule. It is particularly important that the reification classes (i.e., manifested objects, recognitions, creation, and production events) have a unique identifier, guaranteed by the progressive numbering.

Besides the type, the ID such defined provides additional information about the resource. The initial part of it may indicate 1) the nature of the described object, such as *OBJ* for cultural objects, *ART* for the artworks that they represent, *BIBL* for bibliographic records, 2) the entity to which they may be narrowly related, as it is in the case of manifested objects, in which the first part of the ID is the identifier of the artwork in which they are recognized. Table 6.5 summarizes the ID creation rules adopted for each entity, whereas manifested subjects and recognitions are described in Table 6.6.

In order to better maintain the relation between the physical and visual instances describing an artwork, the same progressive number is maintained (e.g., the Artwork of the cultural object `d:OBJ1001test` is `d:ART1001ext`). Furthermore, the artworks used as a test for database modeling are indicated by the suffix 'test', whereas not described cultural objects, which were added for providing more contextual information about the actually described works, have the suffix 'ext'. Those artworks included only to describe the main characteristics of a type, to which the art historians refer while describing other subjects, are identified by the suffix 'type'. These suffixes are present both in the Cultural Object ID and in the respective Artwork ID.

Describing manifested subjects (i.e., the recognized subjects specifically as they appear in the analyzed artwork) requires a best practice to avoid entity repetition. In this case, their ID is composed of the Artwork ID (e.g., ART1001), a string identifying their type, and a progressive number. Whereas the number is inserted in the database, the type string is added in the conversion phase[10]. As shown in Table 6.6, the string 'AM' identifies Artistic Motifs, 'COMP' Compositions, 'IMG' Images, and 'INTRINSIC' Intrinsic meanings.

Similarly, Recognitions at each level are composed by the Artwork ID, followed by a string identifying their type and a progressive number.

[10]the semantics in the tabular framework are already expressed by the column name

Table 6.5: Identifiers creation rules

Type of object	**ID creation rule**	**Further specifications**	**Example**
Cultural objects	'OBJ'+ <number> + eventual specification	'test': artwork used as a test case 'type': artwork solely describing a cited iconographical type 'ext': external artworks not further interpreted	d:OBJ1001test d:OBJ1005type d:OBJ1515ext d:OBJ3102
Artwork	'ART'+ <number> + eventual specification. Number and eventual specification are the same of the respective Cultural Object	'test' 'type' 'ext'	d:ART1001test d:ART1005type d:ART1515ext d:ART3102
Book	<number> + 'bibl'		d:1097bibl
Text	<number> + 'bibl' + - + 'content'		d:1050bibl-content
Production event (E12)	'PROD' + <objectID>		d:PROD1141 d:PROD1051bibl
Creation event (E65)	'CRE' + <objectID>		d:CRE1014bibl-content

Table 6.6: Identifiers creation rules (Manifested subjects and recognitions)

Type of object	**Identifying string**	**Example**
Artistic Motif	'AM'	d:ART1001-AM1 d:ART1512-1-AM1
Composition	'COMP'	d:ART1001-COMP1 d:ART1512-1-COMP1
Image	'IMG'	d:ART1001-IMG1 d:ART1512-1-IMG1
IntrinsicMeaning	'INTR'	d:ART1001-INTRINSIC1 d:ART1512-1-INTRINSIC1
Preiconographical Recognition	'PREICREC'	ART1001-PREICREC1 ART1512-1-PREICREC1
Iconographical Recognition	'ICREC'	d:ART1001-ICREC1 d:ART1512-1-ICREC1
Iconological Recognition	'ICONOLREC'	d:ART1001-ICONOLREC1 d:ART1512-1-ICONOLREC1
Interpretation Description	'DESC'	d:ART1001-DESC d:ART1512-DESC1

6.3.5 External vocabulary and alignment

Following the principle of interoperability, we aligned the entities to external vocabularies and identifiers when possible. According to the intended use, we adopted different strategies:[11]

- For terms intended to consistently describe the resources or identified by a unique specific vocabulary, e.g., their type or style, we directly reused the terms of a vocabulary. When the respective term in the chosen vocabulary was not found, an internal URI was created following the URI creation rules described above.

- For entities (e.g., artworks, people), an internal URI was created and aligned to external entities, when available. This approach allows an alignment of a single entity to multiple Knowledge Bases and helps in maintaining a robust internal description.

- When possible, the super-class or the higher hierarchy terms of the aligned external vocabulary were retrieved. For example, when a Character was

[11]The code used for the alignment is made available in the project repository: https://github.com/SofiBar/IconologyDataset

aligned to an Iconclass term, the belonging of the character to the Iconclass category was registered (e.g., 'Bible').

In the following sections, we further describe each identified typology.

6.3.5.1 External vocabulary

The direct reuse of vocabulary (first case of the previous list) was adopted for the characteristics of cultural objects, namely: the type of artwork (e.g., 'oil painting'), the material and style. All of them were reconciled with the Getty vocabularies. The strategy adopted for each type is illustrated in Table 6.8. For each entity type, the hierarchy name was identified (e.g., `aat:300010357` expressing the 'Materials' hierarchy), and a string match between the database terms and the labels of the Getty's terms under the specific category was performed.

A direct reuse of current ontologies and vocabularies was attempted for cultural object relations. Despite some of them finding a respective in CIDOC-CRM (i.e., 'detail of' and 'feature of' described with `crm:P56i_is_found_on`, and part of described with `crm:P46i_forms_part_of`), the majority could potentially be aligned with Getty CONA associative relations[12]. Nevertheless, we could not reuse them, since they are not released in LOD yet.

For this reason, we provide the alignment we adopted, envisioning a better reconciliation should the CONA relations be released as LOD. Table 6.7 gives an overview of the cultural object relations, providing, if present, 1) the directly reused property from CIDOC-CRM, and 2) the possible alignment with Getty Vocabularies associative relations.

6.3.5.2 Reconciliation

The creation of an internal URI aligned with external resources was the preferred approach for the majority of the entities described. This allows one to have a stable internal definition of the uniqueness of the resources and align the same instance with multiple vocabularies.

The alignment was made through a script that checked for the entities, each time customized according to the type of object treated and the entity searched. The process followed a semiautomatic approach, as the results were then manually checked, and the accepted ones were added to the database, and linked to the correspondent entities through the relation `owl:sameAs`.

Table 6.8 provides an overview of the alignments made for each type of entity, highlighting the strategy adopted and the Knowledge Bases to which they were

[12]Available at: https://www.getty.edu/research/tools/vocabularies/guidelines/cona_3_5_associative_rels-copy.html

Table 6.7: Cultural object relations aligned with CONA associative relations

Dataset relation	**CIDOC-CRM relation**	**CONA associative relation**
copy of		4311 copy after
derives from		
detail of	crm:P56i_is_found_on	
feature of	crm:P56i_is_found_on	4603 incorporated in
model for		4125 model for
part of	crm:P46i_forms_part_of	4215 part of same whole
project for		4131 plan for
reconstruction of		4103 native reconstruction
reintegration of		
related to		
removed from		
study for		4115 study for
took as model		
first version of		

aligned.

As the alignment of artworks was particularly challenging, due to the inconsistency with which artworks titles are registered between knowledge bases, we aligned only those artworks having an author specified. In particular, we used regular expressions to look for entities having the same author label and containing the title in the title label on Wikidata.

Characters and Allegories were reconciled with Iconclass. Several attempts were made to set a search to retrieve the most precise number of alignments. This led to the choice of performing a fuzzy match of the character over the single keywords related to each Iconclass term. For this aim, we used the Fuzzywuzzy Python library.[13] As four types of similarity are offered, we calculated the means of all of them, setting the threshold to 95 points of similarity, over a total of 100. We adopted the same approach to calculate all the fuzzy string matching, adjusting case by case the minimum value required.

The places were aligned with Getty's Thesaurus of Geographic Names (TGN). In particular, we performed SPARQL queries in which the name of a city had to be part of the 'city' hierarchy. Furthermore, we retrieved the Country in which the city is located by indicating the type (e.g., 'Republics'). In this way, it was possible

[13] https://pypi.org/project/fuzzywuzzy/

to align both cities and countries. We further extracted Wikidata's entities which are already aligned to the Getty ones through the relation `wdt:P1667`.

Furthermore, we aligned people to Wikidata, specifying the type `wd:Q5`. We further extracted the respective terms on VIAF and ULAN through the alignments already expressed in Wikidata.

Using a similar approach, symbols were aligned to HyperReal.

To further enrich data, periods were manually aligned to the PeriodO period gazetteer[14] (Rabinowitz et al., 2018) and to Chronontology,[15] as they add further information about period dates and geographical area.

6.3.5.3 Hierarchy of aligned objects

As subjects recognized in artworks are thoroughly described, a wide range of different instances is created, making difficult the retrieval of similarities among them. For this reason, when feasible, we retrieved the higher hierarchy terms of the aligned ones. In detail, we focused on Characters, Stories and Allegories.

As some characters were registered with a term identifying a specific aspect characterizing them (e.g., Venus Coelestis), the relation with the closest character (in this case, Venus) was registered through the relation `skos:broader`. The same approach was adopted for stories (e.g., the story 'An eagle tears at Prometheus' liver' is related to 'Story of Prometheus'). Some types of objects were better specified with a further retrieval of the higher hierarchy to which they belong in the aligned vocabulary. It is the case, for example, of Characters and Stories. We enriched the results by retrieving the higher hierarchy of mythological characters from Wikidata. Since they are the most represented subjects, we wanted to perform more detailed analyses on them. They were retrieved by searching for terms that were instances of (`wdt:P31/wdt:P279*`) mythological characters and classical deities or of their subclasses.

We tried to align cultural phenomena to existing classification, e.g., the Warburg's iconographical index categories. Although the index constitutes a fundamental reference point for iconological studies, an alignment to their categories doesn't help to better define the nature of each phenomenon, as it reflects the scholars's interpretations rather than providing a systematic, typological classification of cultural phenomena. For this reason, we provided a typological generalization of the cultural phenomena and their thematic core, specified as a type (`crm:P2_has_type`).

[14] https://perio.do/
[15] https://chronontology.dainst.org/

Table 6.8: Overview of the alignments made and matching strategy

Type	**Criterion**	**Aligned with**
Artwork	String match of artwork title and exact match of author	Wikidata
Artwork	Look for the artwork title, artist's name and surname with regular expressions	Zeri&Lode
Characters	fuzzy match with the single key words > 95	Iconclass
Places(cities with countries)/Getty	search the string with type "city" (gvp:placeType [rdfs:label 'cities'@en])and retrieve the country in which it is located ([rdfs:label 'republics'@en]). If the fuzzy match is > 95, add it to the dictionary	Getty
Places	Extract the Wikidata entities from the Getty ones (wdt:P1667: Getty Thesaurus of Geographic Names ID)	Wikidata
People	entity of type Q5, string search	Wikidata
People	Retrieve alignments through the already reconciled wikidata entities - wdt:P245 (ID ulan), wdt:P214 (ID viaf)	VIAF/ULAN
Material	search for a term in the 'material' hierarchy (aat:300010357) having exactly the same string	AAT
Style and Period	exact string matching with the label of a term belonging to the hierarchy 'style'	AAT
Period	manual alignment	Chronontology, PeriodO period gazetteer
Symbols	fuzzy match with symbols and symbolical meanings	HyperReal

Figure 6.10: (1) Allegory of Salvation, c. XIII A.D., Venice, St. Mark's Basilica, external wall; (2) Hercules carrying the Erymanthian Boar, III Century, Venice, St. Mark's Basilica, external wall.[17]

6.4 A case study description

In this section, we provide a description of a case study included in the dataset, to better illustrate the modeling decisions and description techniques adopted.[16] To this end, we present two reliefs on the external wall of Saint Mark's Basilica in Venice. In detail, we give insights into how extracting the information from the scholar's free text is conducted, how the data is inserted in a tabular database, and how they are finally represented according to the ontological modeling presented in the previous chapters.

This example is considered by Panofsky and Saxl (1933). They noticed that two reliefs (see Figure 6.10) belonging to different historical periods and representing different subjects have clearly the same shapes and general arrangement. One of them (Figure 6.10, right) is recognized by them as a late Roman one, representing the story of Hercules Carrying the Herymanthian boar, whereas the other one, from the XIII Century A.D., represents an allegory of Christ saving the Christian souls.

According to the authors' interpretation, the medieval artist used the Roman relief as a model for representing a new subject, closer to their culture. Since, during the Middle Ages, there were other cases in which a classical artwork was used as a model for contemporary subjects, deprived of its classical content, they claim

[16]The example here described was made available as part of the ICON ontology documentation: https://w3id.org/icon/docs/

[17]Images were respectively retrieved from 1) Europeana; image license: CC BY-ND; rights owner: University of Bologna 2) www.canalgrandevenezia.it; image license: CC BY-NC-SA 3.0 IT

that these artworks are pieces of evidence of a cultural phenomenon happening. They conclude that Medieval Western art was unable to retain a classical prototype without destroying its original meaning since they were not capable of realizing the unity of classical form and classical subject matter.

The table shown in Figure 6.11 summarizes how this content can be described at each level of interpretation.

In the Roman relief (see Figure 6.10, right), a nude man with lion skin carries on his shoulders a boar. In the lower right corner, another figure takes shelter in a barrel. At the second level of interpretation, these level 1 subjects can be grouped and recognized: the man is Hercules with his attribute 'lion skin' and the man in the barrel is King Eurystheus.

The same interpretative process can be applied to the other artwork (see Figure 6.10, left), where the man is recognized as Christ, the deer as the symbol of the Christians' souls, and the dragoon in the corner is recognized as a symbol of the Devil. His actions of carrying the deer and of stepping on the Devil are respectively symbolically interpreted as the act of saving the souls and defeating the Devil. In the last level, the cultural phenomenon described above is recognized as belonging to the medieval artifact.

Each level of interpretation described will be formally represented in the following sections.

6.4.1 Information extraction and subdivision into levels

As mentioned above, the process of information extraction was manually conducted. In this section, we describe how it was conducted for the selected case study. The text was read and interpreted, and the subjects identified were divided into levels. The text reproduced in Figure 6.11 is the text section in which Panofsky and Saxl focus on the description of the current case study. After the first reading, the text was interpreted through the lenses of the three-layered Panofskian approach, obtaining the identification of phrases describing the possible subjects. In the figure, the levels of interpretation are distinguished through the use of different colors. Secondly, the type of object described, and their relations were identified (e.g., the compositions gathering level 1 objects, corresponding to a second-level subject). For example, Hercules is understood as the composition of a nude man and a lion skin. The correspondence between levels is indicated in the table by the indication of the same number. As it is possible to notice, the expression of the free text into structured data implies a more concise way of expressing the concepts (see the expression of the cultural phenomenon), and the normalization of some terms (e.g., 'drapery' became 'cloth'). The assignment of Artistic Motifs

and the subject type occurs during the database population phase, described in the next section.

«[...] an antique relief of Hercules could not be imitated without changing its mythological subject. The lion's skin was replaced by a fluttering drapery, the boar became a stag, the terrified Euristheus was left out, and the hero was made to stand upon a vanquished dragon. As the human soul was often symbolized by a stag, the result of these changes was that the classical hero had been transformed into the Saviour conquering evil and saving the souls of the Faithful. From this example we learn that mediaeval Western art was unable, or, what comes to the same thing, was unwilling, to retain a classical prototype without destroying either its original form, or, as here, its original meaning.»
E. Panofsky and F. Saxl, 1933, p.228

Level	Hercules and the Erymanthian Boar	Allegory of Salvation
1	nude man with a lion skin (1) carrying on shoulders a boar (2) Man taking shelter in a barrel in the lower right corner (3)	Nude man with a cloth (1) carrying on shoulders (2) a deer (3) and stepping on (4) a dragoon in the lower right corner (5)
2	Hercules (1) Erymanthian Boar (2) King Eurystheus (3)	Christ (1) deer as symbol of souls (3) symbolic act of saving souls (2) dragoon as symbol of the Devil (5) symbolic act of defeating the Devil (4)
3		The medieval mind was incapable of realizing the unity of classical form and classical subject matter

Figure 6.11: summary of the process of the scholar's text understanding and abstraction and how this content can be described according to the three-level approach, applied to the selected case study

6.4.2 Registration in the tabular database

The record in the database is described by different tables. In the table containing the cultural object metadata, introduced in Section 6.2.2.1, the known information about the reliefs is registered. In this case, as shown in Table 6.9, we report the artwork title, date, type, location (Venice, St Mark's Basilica), specific location (external wall), the DOI for identifying the text by Panofsky and Saxl (1933), the period, the aligned URI of the second artwork retrieved from Europeana. We further include the URL of the image and the license specification.

In the second place, the content recognized as belonging to different levels is expressed in the respective tables.

6.4.2.1 Artwork content interpretation

The content previously recognized as belonging to the first level of interpretation is examined, and it is established 1) which elements are the actual subjects, 2) of which type they are, and 3) which are qualities describing them.

First and second level recognitions The relief representing 'Hercules carrying the Erymanthian Boar' shows, at the first level of interpretation, a man wearing a lion's skin carrying a boar on his shoulders. We know that this group of subjects will be recognized as the character Hercules at the second level of interpretation.

Table 6.9: Cultural object Table filled with the case studies metadata

Feature	**Cultural Object 1003**	**Cultural Object 1004**
ID	1003	1004
Title	Hercules carrying the Erymanthian Boar	Allegory of Salvation
Title for checking duplicates	Hercules carrying the Erymanthian Boar, III Century, Venice, St. Mark's Basilica	Allegory of Salvation, XIII Century, Venice, St. Mark's Basilica
Date	III Century	XIII Century
Type	relief	relief
Book ISBN or DOI	https://doi.org/10.2307/1522803	https://doi.org/10.2307/1522803
City of conservation	Venice	Venice
Place of Conservation	St. Mark's Basilica	St. Mark's Basilica
Type of relation	feature of	feature of
Related artwork	St. Mark's Basilica, XI-XIII Century, Venice, St. Mark's Basilica	St. Mark's Basilica, XI-XIII Century, Venice, St. Mark's Basilica
IDext	1153ext	1153ext
Specific location	external wall	external wall
Object Type	church (buildings)	church (buildings)
Period	Late Roman Antiquity	Middle Ages
URI		https://www.europeana.eu/en/item/22/_21521

Then, the man taking shelter in a barrel will be recognized as King Euryshteus at the second level. Lastly, the overall scene will be interpreted as the Story of Hercules carrying the Erymanthian boar to king Eurystheus. These subjects are uniquely identified as Artistic Motifs, i.e., the specific manifestation in the artwork of the subjects represented, which allows us to distinguish, for example, the man representing Hercules from the man representing King Euryshteus.

These level 1 subjects can be grouped in different compositions. A composition is a unity of artistic motifs that have a sense together, which will eventually be further recognized at the second level of interpretation. Compositions can be also used for indicating the objects involved in an action as well, and it is not required that every composition is further interpreted in the second level of interpretation. As a result, we obtain three compositions:

- composition 1: man, lion skin, action of carrying on shoulders, boar
- composition 2: man, taking shelter, barrel
- composition 3: all the Artistic Motifs identified

The first man, the lion skin, and the boar are identified as Natural Elements, grouped, with the action of 'carrying on shoulders', in composition number 1 by indicating their Artistic Motifs count separated by a '@' (see Table 6.2). Such composition will be recognized as the Character Hercules at the second level of

interpretation, with the lion skin (artistic motif 2) as recognizing attribute (Table 6.3). The character of the Herymanthian Boar is directly related to the boar's Artistic Motif (number 3). Similarly, the man and the barrel are registered as Natural Elements, and, grouped with the action of taking shelter in a barrel, form the composition number 2 that will be recognized as King Eurystheus at the second level of interpretation. We included a final composition, which gathers all the Artistic Motifs previously described, that will be recognized as an Invenzione representing the story of Hercules and the Erymanthian boar. All the previously identified Characters (i.e., Hercules, the Erymanthian Boar, and King Eurystheus) are declared as part of this story.

The description of the other artwork (ART1004) takes place with the same process. The subject type was discriminated among Natural Elements and Actions, as no Expressional Quality is recognized in this case. Consequently, a composition of the Natural Elements 'man', 'cloth', and 'deer', plus the action 'carrying on shoulders', is recognized, and further interpreted at the second level of interpretation as Christ. Specific motifs composing it are further recognized. The deer (motif number 3) is recognized in Image 2 (defined in the field 'Count') as symbol of the souls, and the action of carrying on shoulders expresses the symbolical meaning of saving souls through the recognition of a symbol related to Image 3. Furthermore, the scholars described Christ as stepping on the dragoon. The act of stepping on is included in another composition with the dragoon, which is further recognized as the symbolical action of defeating the devil (image with count 4 in Table 6.12). The dragoon, recognized as a Natural Element, and identified with number 6, is identified with a symbol of the Devil. Finally, the Allegory of Salvation is assigned to a composition gathering all the Artistic Motifs, and all the second-level subjects (in this case, Symbols and Characters) are declared as part of it.

Formal Motif Recognition Since Saxl and Panofsky claim that the artists making the medieval relief on the Saint Mark's Basilica wall copied the shape and general arrangement from the Roman one, a Formal Motif Recognition occurs. Therefore, the composition identifying the scene represented can be related to the composition of the classical scene. In addition, more recognitions can identify the exact correspondence between:

- Christ and Hercules
- the deer and the boar
- the dragoon and King Eurystheus
- the action of carrying on shoulders in both the artworks

Table 6.10: Expression of the recognized first-level objects in tabular format

Artwork	**Artistic Motif Count**	**Natural Element**	**Quality**	**Action**	**Composition**	**Composed By Artistic Motif**
ART1003	1	man	nudity			
ART1003	2	lion skin				
ART1003	3	boar				
ART1003	4			carrying on shoulders		
ART1003	5	man				
ART1003	6	barrel				
ART1003	7			taking shelter in barrel		
ART1003	8				1	1 @ 2 @ 3 @ 4
ART1003	9				2	5 @ 6 @ 7
ART1003	10				3	1 @ 2 @ 3 @ 4 @ 5 @ 6 @ 7
ART1004	1	man	nudity			
ART1004	2	cloth	fluttering			
ART1004	3	deer				
ART1004	4			carrying on shoulders		
ART1004	5			stepping on		
ART1004	6	dragoon				
ART1004	7				1	1 @ 2 @ 3 @ 4
ART1004	8				2	5 @ 6
ART1004	9				3	1 @ 2 @ 3 @ 4 @ 5 @ 6

Table 6.11: Expression of the visual citation performed by the Medieval relief of the Roman one

Count	**from Artwork 1**	**Artistic Motif Count 1**	**Composition Count 1**	**to Artwork 2**	**Artistic Motif Count 2**	**Composition Count 2**	**Gives Support To**
1	ART1003		1	ART1004		1	ART1004-ICONOLREC1 @ ART1004-ICONOLREC2
2	ART1003	3		ART1004	3		ART1004-ICONOLREC1 @ ART1004-ICONOLREC2
3	ART1003	4		ART1004	4		ART1004-ICONOLREC1 @ ART1004-ICONOLREC2
4	ART1003	2		ART1004	2		ART1004-ICONOLREC1 @ ART1004-ICONOLREC2
5	ART1003		3	ART1004		3	ART1004-ICONOLREC1 @ ART1004-ICONOLREC2

The relation between the subjects and their part established by the recognition of a visual citation is reported in Table 6.11, which reflects the structure of the appropriate table in the database. The Roman relief (ART1003) is indicated as the artwork from which the formal motif was taken, and the Medieval one as the copying artwork (column 'to Artwork 2'). Several recognitions thoroughly describe which parts were cited. Recognition 1 expresses the citation of the arrangement of the characters Hercules and Christ and their attributes, grouped respectively in composition 1 at the first level of interpretation. Further recognitions establish punctual correspondence between the Artistic Motifs identifying the boar and the deer (Artistic Motif number 3 in both the artworks), the lion skin and the cloth (number 2), the action of carrying on shoulders (number 4), and the overall arrangement of the whole scene (composition number 3). These recognitions give support to the recognition of the cultural phenomena identified at the last level. Therefore, the 'gives support to' column is filled with the identifier of such third-level recognitions.

Table 6.12: Expression of the second level subjects recognition in a tabular format

Artwork	Count	Artistic Motif	Composition	Recognizing Attribute	Character	Symbol	Story	Allegory
ART1003	1		1	2	Hercules		Hercules and the Erymanthian Boar	
ART1003	2	3			Erymanthian Boar		Hercules and the Erymanthian Boar	
ART1003	3		2		King Eurystheus		Hercules and the Erymanthian Boar	
ART1003	4		3				Hercules and the Erymanthian Boar	
ART1004	1		1		Christ			Allegory of Salvation
ART1004	2	3				deer-souls		Allegory of Salvation
ART1004	3	4				carrying on shoulders-saving souls		Allegory of Salvation
ART1004	4		2			stepping on-defeating devil		Allegory of Salvation
ART1004	5	6				dragoon-devil		Allegory of Salvation
ART1004	6		3					Allegory of Salvation

Iconological recognition As shown in Table 6.13, at the third level of interpretation, two cultural phenomena are identified: one expressing the concept in detail ('Medieval Western art was unable to retain a classical prototype without destroying its original meaning'), and the other one expressing the general classical motif recovery which is comprised in the previous one. As the phenomenon belongs

to the Medieval relief, it is related only to it, while the Roman relief is indicated as the evidence supporting it.

Table 6.13: Expression of the third level subjects recognition in a tabular format

Artwork	Count	Concept	CulturalPhenomenon	Image	Evidence
ART1004	1		Medieval Western art was unable to retain a classical prototype without destroying its original meaning		ART1003
ART1004	2		Classical motifs recovery		ART1003
ART1004	3	Redemption		1 @ 2 @ 3 @ 4 @ 5	

6.4.3 Cultural Object information modeling

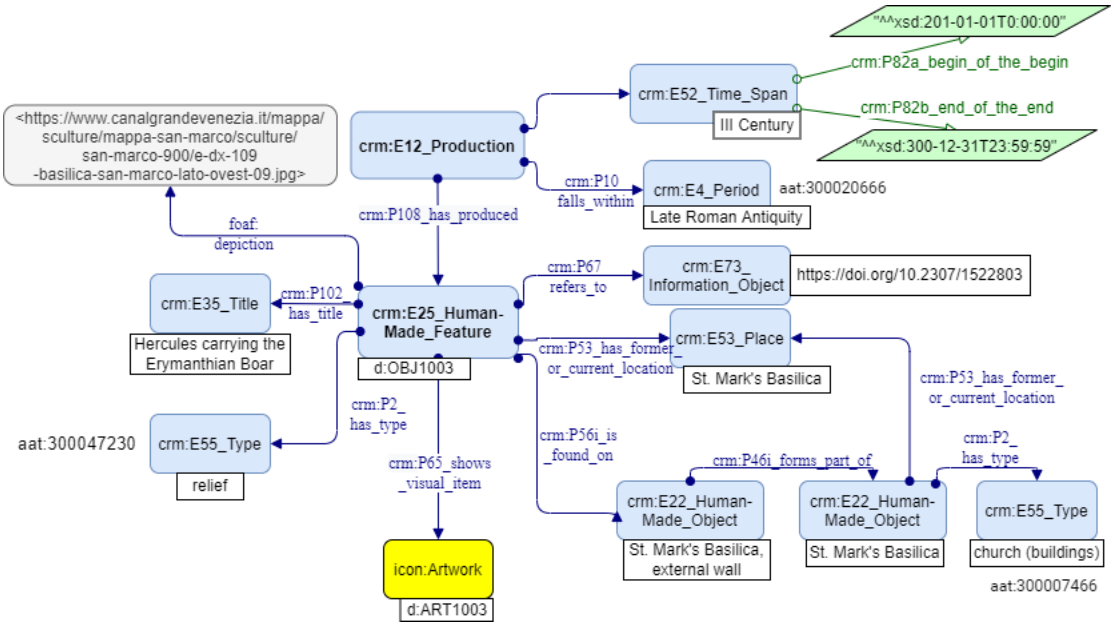

Figure 6.12: Modeling of the metadata of the Roman relief

During the conversion, the object IDs were created from the tabular data illustrated. Consequently, from the cultural object and artwork common identification number in the database (1003, 1004), the object URI (e.g., `d:OBJ1003`),the artwork (e.g., `d:ART1003`), and the production one (e.g., `d:PROD1003`) are created.

Figure 6.12 shows how the metadata listed in Table 6.9 are converted to the RDF format. As it is registered that the cultural object is feature of the Basilica's external wall, it has as type `crm:E25 Human Made Feature`, which is found on (`crm:P56i`) the external wall, found on the St. Mark's Basilica. The cultural object has as type `aat:300047230`, which is the Getty's AAT term for 'relief'. Furthermore, the title (`crm:E35`), and the location (`crm:E53`) are specified,

and the book on which the image is found (Panofsky & Saxl, 1933) is indicated by the relation `crm:P67_refers_to`. The time (III Century) and period (Late Roman Antiquity) are related to the cultural object production instance (`crm:E12`). From the time-span label, the beginning and end date, specified as strings, were automatically extracted.

6.4.4 Pre-iconographical Recognition

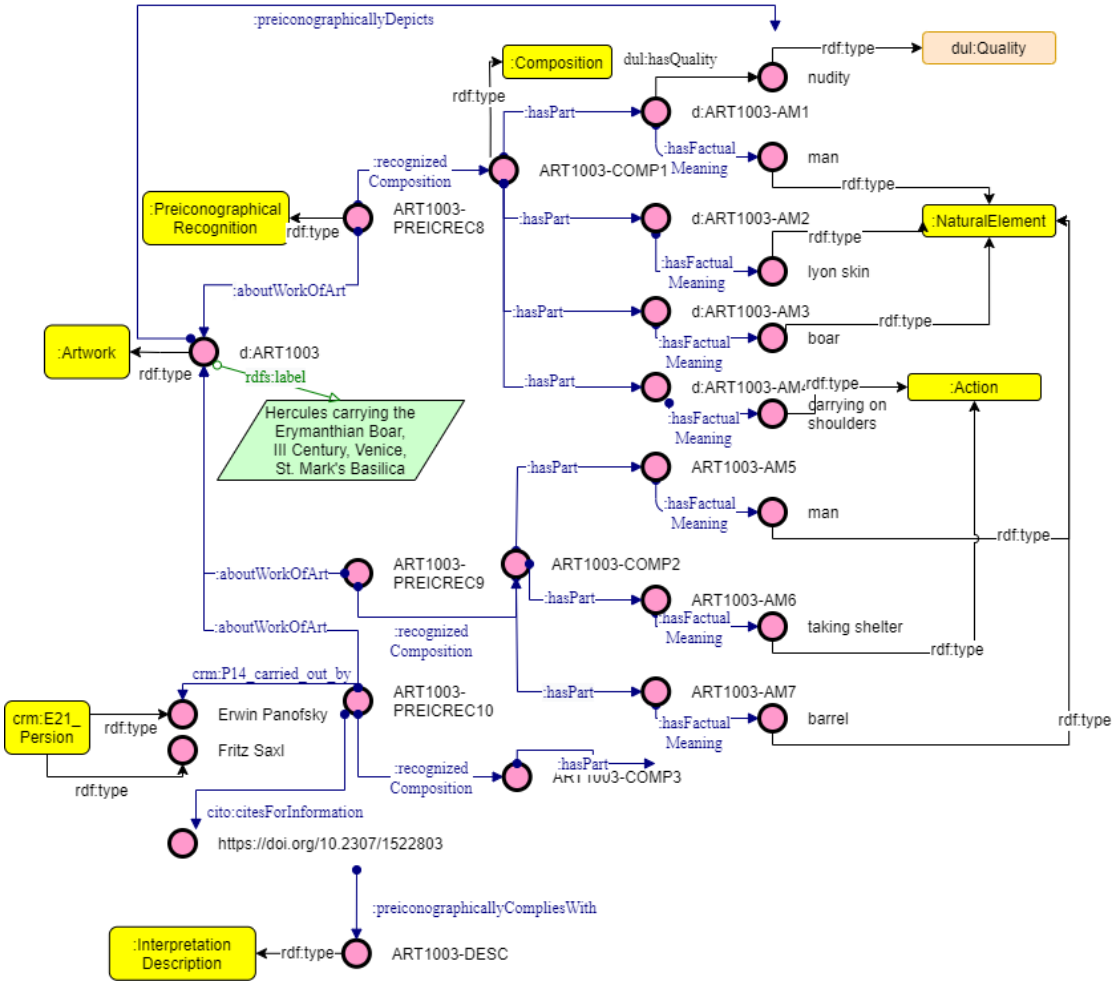

Figure 6.13: Modeling of the pre-iconographical Recognitions of the Roman relief

The recognitions illustrated in Section 6.4.2.1 are modeled as shown in figures 6.13 and 6.14. The authors Panofsky and Saxl (crm:P14_carried_out_by) recognize (PreiconographicalRecognition) a Composition, which is composed (hasPart) by ArtisticMotifs. The Artistic Motifs are then related to the general subjects and qualities that are recognized. For example, the quality (dul:Quality) 'nudity' is linked (dul:hasQuality) to the Artistic Motif related to the subject 'man' (ART1003-AM1). In this way, it is possible to affirm that only the man as represented in this artwork has 'nudity' as quality, and not all the representations of

'man' in other artworks are likely to have the same quality. All the recognitions are about (`icon:aboutWorkOfArt`) the artwork under examination (`ART1003`), and are declared as part of the overall interpretation ART1003DESC, having type `icon:InterpretationDescription` through the relation `icon:preiconographically-CompliesWith`. To foster the retrieval of subjects, the property chain `icon:preiconographicallyDepicts`, acting as a shortcut, relates the artwork to the subjects recognized at the first level of recognition.

The same approach can be used to describe the second relief of the example.

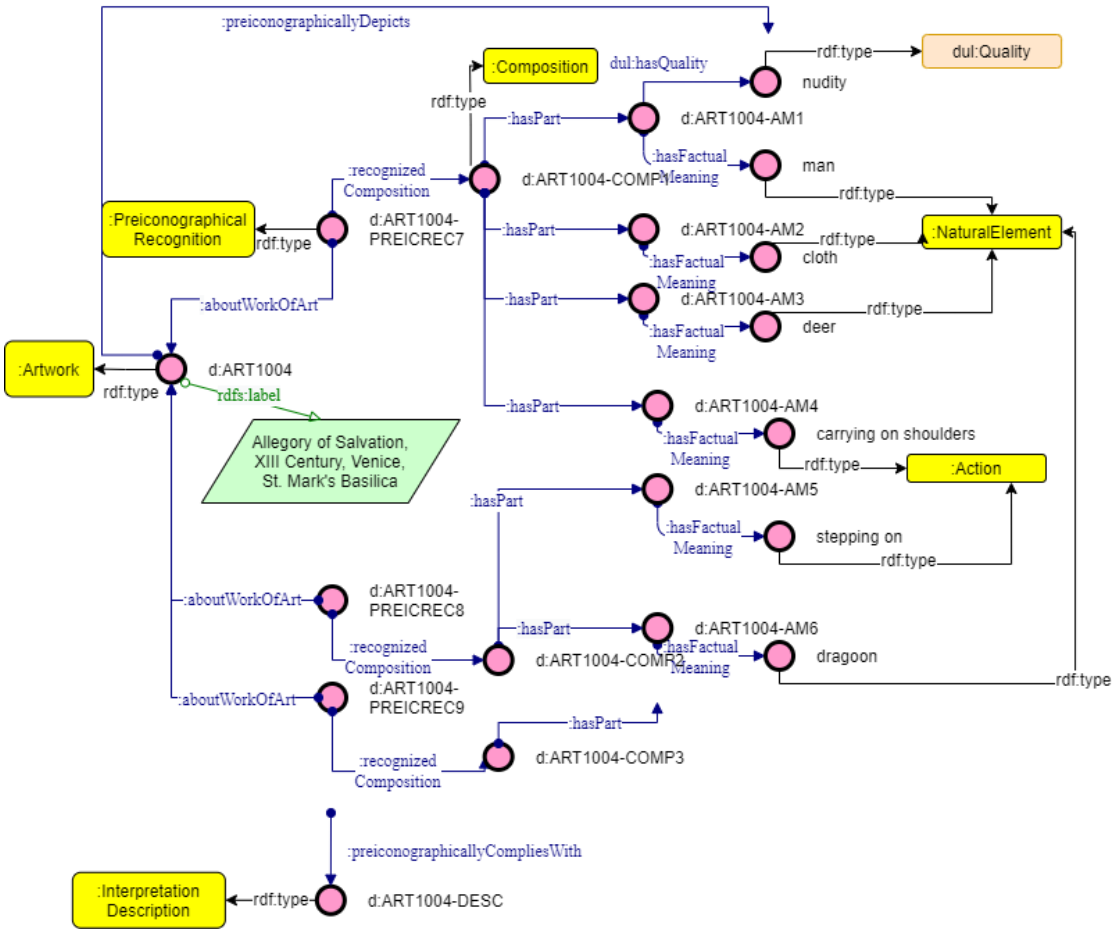

Figure 6.14: Modeling of the pre-iconographical Recognitions of the Medieval relief

6.4.4.1 Formal Motif Recognition

As shown in the scheme, each `FormalMotifRecognition` is about (`aboutWorkOf Art`) the copying artwork. In this case, Saxl and Panofsky are always the authors of the recognition. The parts of the artwork taken as a prototype (e.g., the Roman relief, here indicated as ART1003), namely Artistic Motifs and Compositions, are related to the recognition with the property `hasPrototypicalMotif`, whereas the corresponding parts of the final copying artwork (i.e., the Medieval relief) are

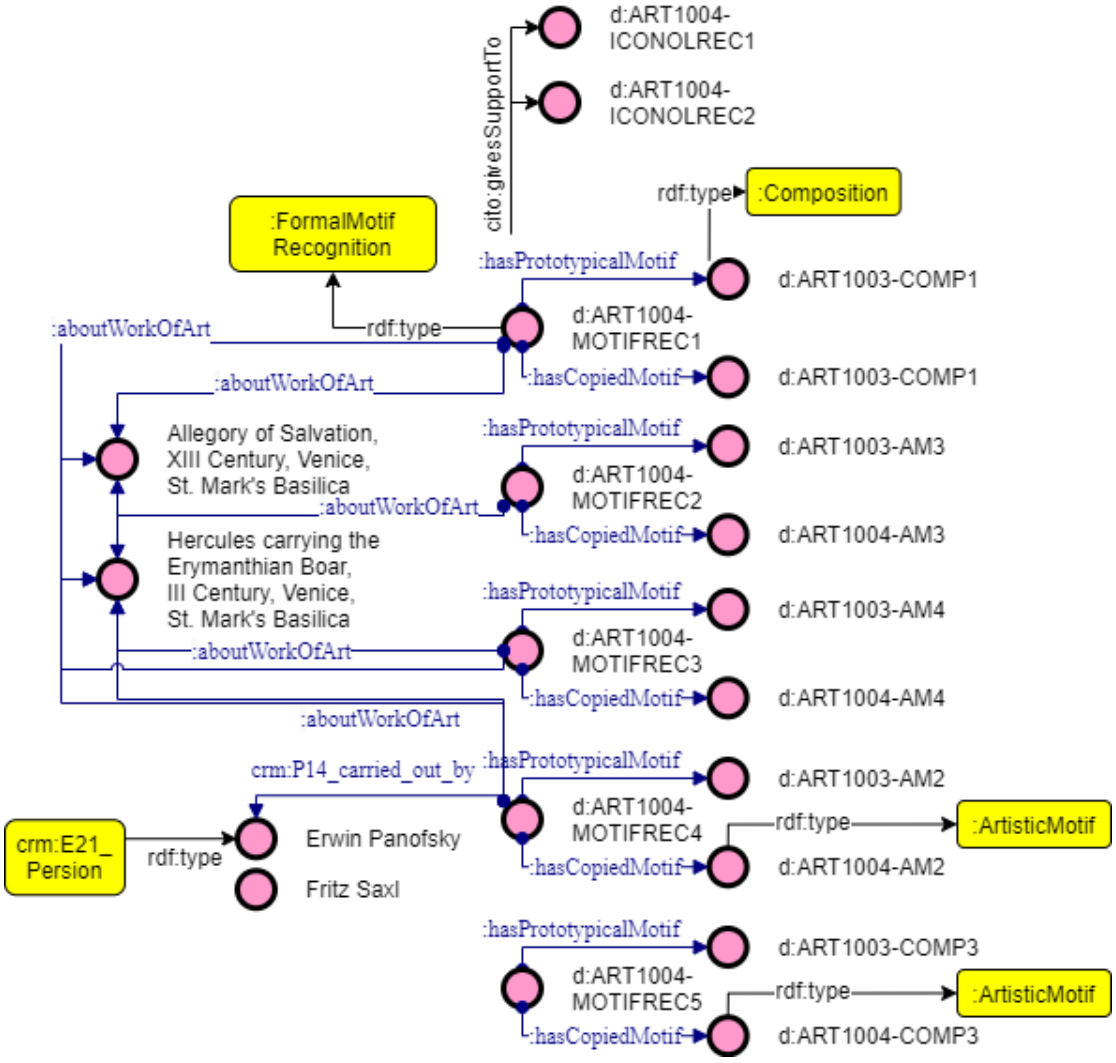

Figure 6.15: Modeling of the formal motifs Recognitions of the case studies

related through the property `hasCopiedMotif`. In this way, the direct correspondence between level 1 levels can be expressed. All recognitions are listed in the InterpretationDescription of the copying artwork (ART1004).

6.4.5 Level 2 modeling

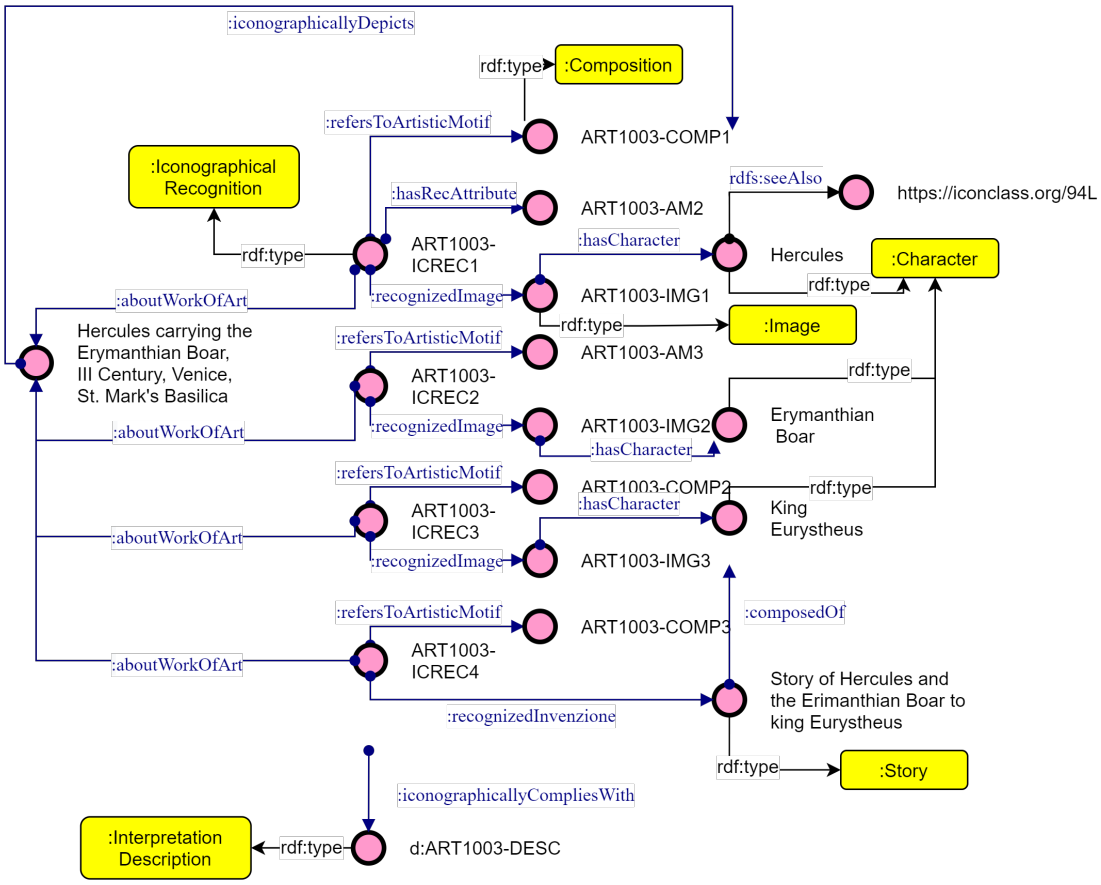

Figure 6.16: Modeling of the Iconographical Recognitions of the Roman relief

Each level 2 subject is identified and linked to its specific manifestation in the artwork considered, i.e., an Image. Each image can be linked to the Artistic Motifs or Compositions already identified at the first level of interpretation by means of an `IconographicalRecognition`. Each recognition links an Image or an Invenzione (i.e., Story or Allegory) to an Artistic Motif or a Composition, to create a direct correspondence between the subjects identified at each level.

In this case, composition 1 is recognized as the character Hercules, which is a character corresponding to Image 1. The Artistic Motif identifying the lion skin is related to the image by the relation `hasRecAttribute` since it is the attribute that allows character recognition. Then, Artistic Motif 3, identifying a not further defined 'boar', is linked to Image 2 and interpreted as the Erymanthian Boar, i.e., a level 2 Character. In the same way, Composition 3 is interpreted as the

Character King Eurystheus. These characters are recognized as belonging to the 'Story of Hercules carrying the Erymanthian boar'. The story is then linked to the artwork and to the level 1 composition identifying the scene represented by means of another `IconographicalRecognition` recognizing the Invenzione (`recognizedInvenzione`).

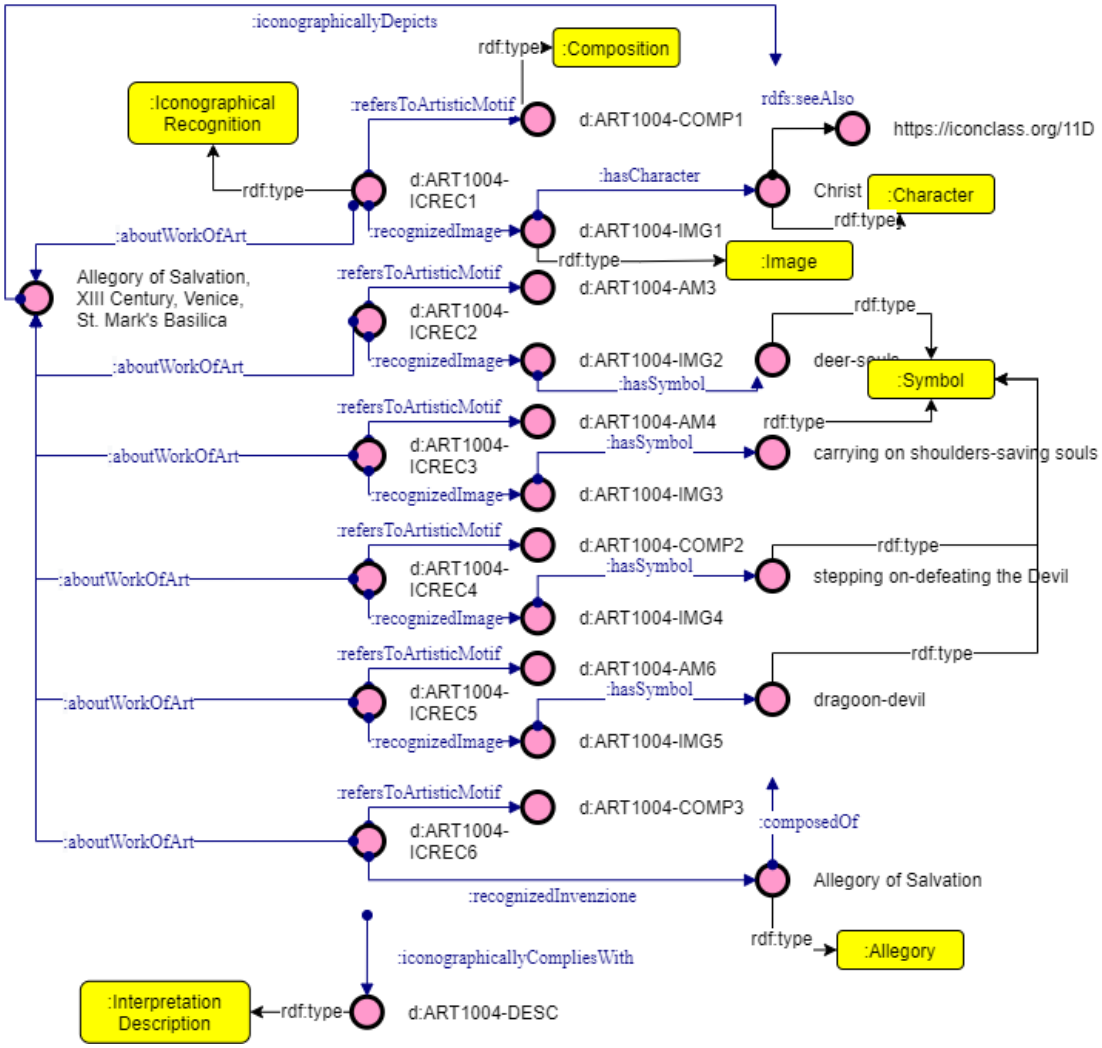

Figure 6.17: Modeling of the Iconographical Recognitions of the Medieval relief

The same structure can be applied to describe the second example. In this case, different types of second-level subjects are recognized, such as Symbols. For example, Natural Elements (deer, dragoon) and Actions (carrying on shoulders, stepping on) are interpreted at the second level of interpretation as Symbols. Finally, the composition identifying the complete scene represented is interpreted as an Allegory of Salvation, and all the second-level subjects recognized are part of it (Allegory, `composedOf`, subjects), as can be seen in the scheme in Figure 6.17. We recognize the scene as an allegory because of the presence of symbols.

6.4.6 Level 3 modeling

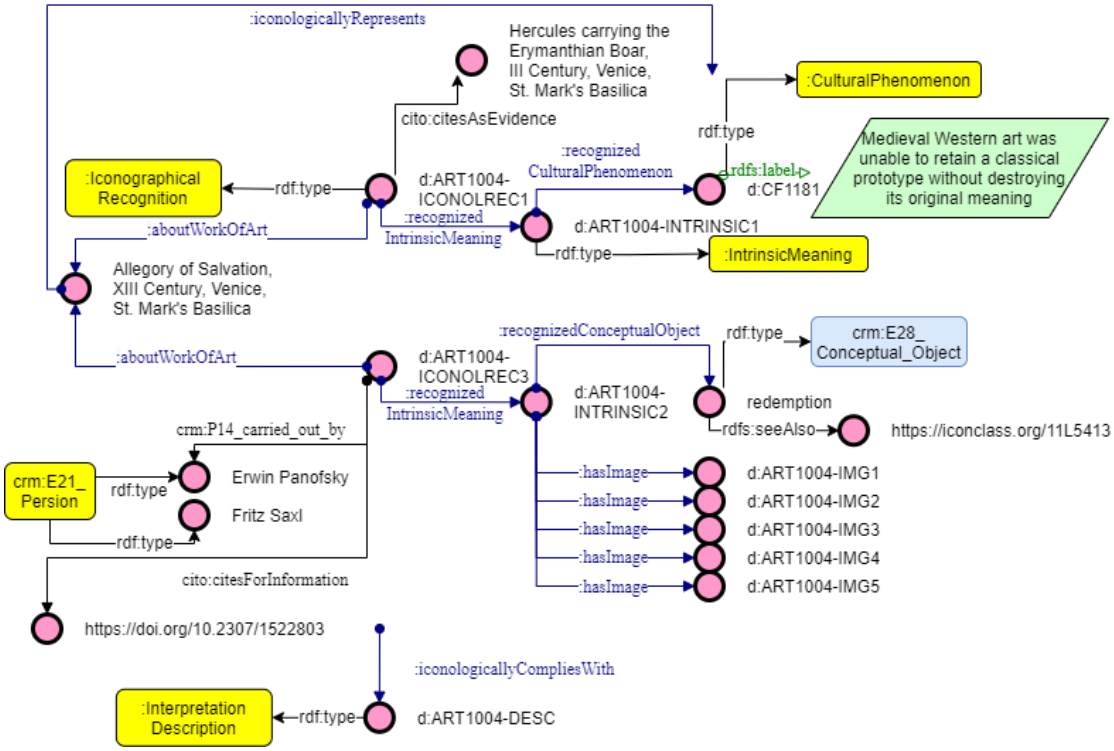

Figure 6.18: Modeling of the Iconological Recognitions of the Medieval relief

In the example described, it is recognized by Saxl and Panofsky that both the artworks are involved in a common cultural phenomenon. In detail, the phenomenon regards specifically the Medieval relief, since it concerns the way in which Medieval artists deal with the Classic art. Therefore, we can say that the phenomenon is embedded in the Medieval relief, and the roman one is evidence of it.

In addition, this recognition is supported by the fact that there is a copy of formal motifs done by the Medieval artist, who took the Roman relief as a prototype for his artwork. Therefore, it is possible to affirm that the Formal Motif Recognitions made at level 1 support the level 3 interpretation, identifying the cultural phenomenon.

At the third level of interpretation, it can also be recognized that the Allegory of Salvation expresses the concept of redemption as one of the meanings belonging to Christian beliefs. Another recognition linking this more profound meaning to the artwork is considered.

This interpretative situation can be modeled as follows (see the scheme for its representation). The authors make a third-level recognition, i.e., an Iconological-Recognition, and recognize that the Artwork representing an Allegory of Salvation has an Intrinsic Meaning (IconologicalRecognition, recognizedIntrinsicMeaning, In-

trinsicMeaning). This Intrinsic Meaning is then linked to the `CulturalPhenomenon` of itself. The Intrinsic Meaning is indeed the specific manifestation of the general phenomenon recognized. The Motifs Recognitions already conducted at level 1 that are described in the scheme of a Formal Motif Recognition (level 1), support this interpretation by means of the relation `cito:givesSupportTo`. The `IconologicalRecognition`, then, cites as evidence (`cito:citesAsEvidence`) the Roman relief.

Following, another iconological recognition links to the artwork another intrinsic meaning, which is related to the concept of 'redemption'. Since this interpretation can be seen in the second-level subjects, we can claim that this meaning is specifically related to them. Therefore, we link the recognition to specific images through the relation `hasImage`.

CHAPTER 7

Analysis

In this chapter, we illustrate the analysis conducted over the Iconology dataset. It is divided into three parts. First, we provide a qualitative overview of the dataset. Second, the domain-specific questions formulated in Section 5.1.2 are quantitatively addressed. Finally, new inquiries are explored.

7.1 Quantitative overview

We first provide an overview of the dataset to better frame the results described in the following sections.[1] Most artworks that have an explicit date (68, 5%) are dated between 1100 and 1700 (Fig. 7.1.a.). Therefore, the dataset mainly represents Western Medieval, Renaissance and Late Renaissance art. In addition, the main themes covered belong to the classical repertoire. Half of the artworks are linked to a phenomenon related to classical antiquity ('Reception of classical antiquity', fig. 7.1.b), and the most frequent second-level subjects are mythological characters (fig. 7.1.d). Although the data set is focused on Panofsky's interpretations, the interpretations of other authors are included (Fig. 7.1.c). The second represented author, Fritz Saxl, is a co-author of one of the source texts. Other authors are included because of preliminary case study results, which were used as a reference for the modeling choices. Panofsky sometimes cites interpretations by other authors; some further information was integrated by the thesis author based on the direct observation of the artworks described by Panofsky in the selected literature.

[1]The analyses presented in this Chapter are available on an online dashboard: https://iconology-dataset.streamlit.app/

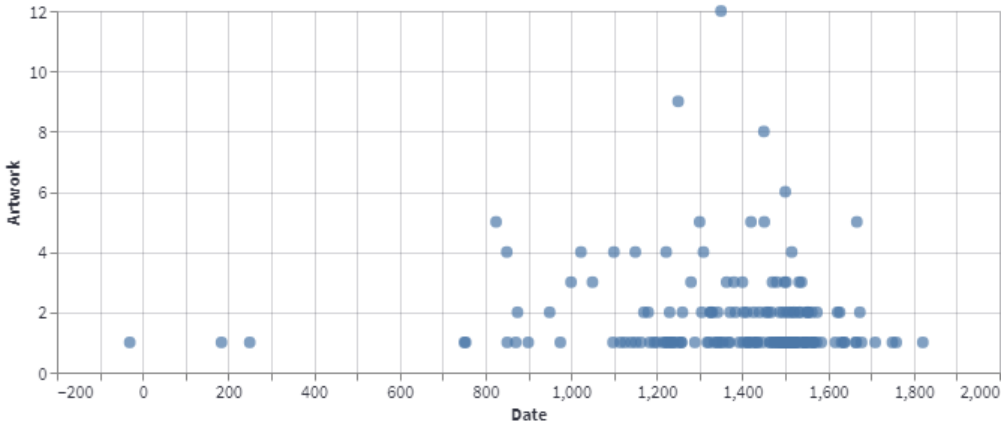

(a)

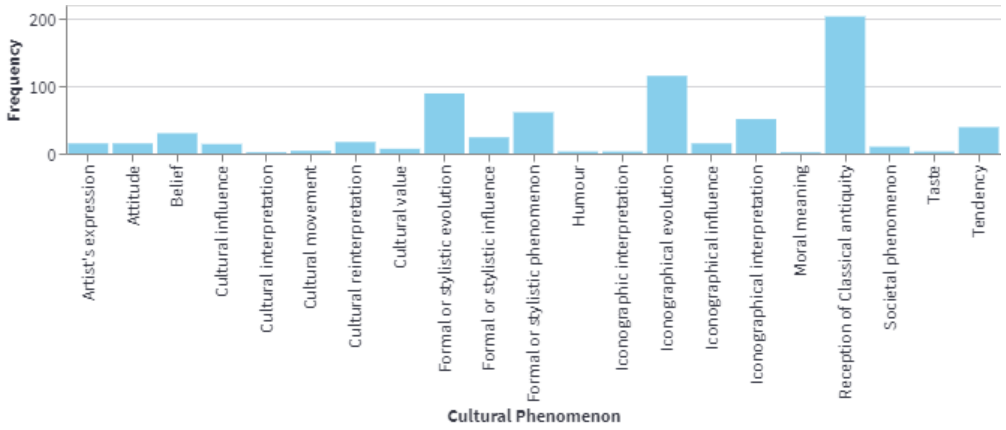

(b)

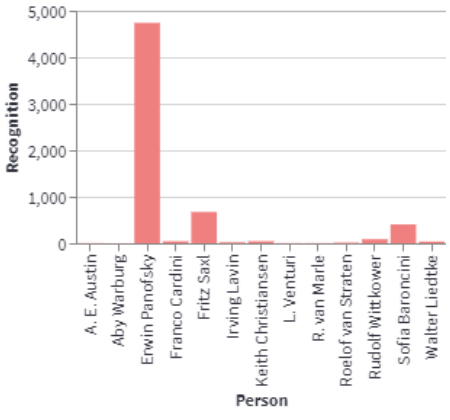

(c)

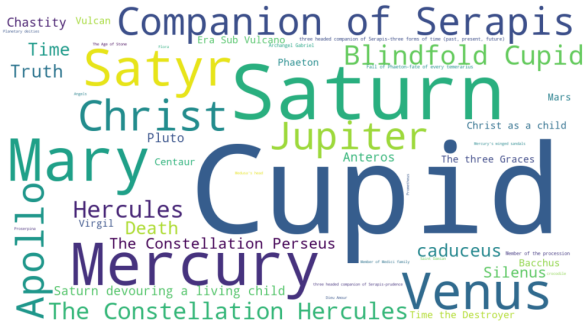

(d)

Figure 7.1: a) Frequency of artworks per date; b) frequency of cultural phenomena types; c) the number of recognitions per person responsible; d) term cloud of second level-subjects

7.2 Domain-specific research questions

To address the research questions quantitatively, we further specialized the 12 domain-specific research questions identified in Section 5.1.2 into 20 subquestions that could be directly performed as SPARQL queries. All questions (Q), subquestions (SQ) and the source from which they were extracted are presented in Table 7.1.

Table 7.1: Overview of research questions extracted from the literature, their specification in sub-questions performed in SPARQL queries.

RQ Question from literature	**Sub-questions**
Cultural phenomena	
Q1. Which cultural phenomena are witnessed by the artworks?(Warburg, 1999)	SQ1. How many artworks have a cultural phenomenon associated?
Q2. What sources and visual aspects motivate their presence?(Cardini, 2001; Warburg, 1999)	SQ2. At what frequency do level 1 and 2 subjects occur with each cultural phenomenon? SQ3. What is the frequency with which a piece of evidence supports the recognition of a cultural phenomenon? Of what type is it? (texts, artworks, etc.) SQ4. Is there any co-occurrence between cultural phenomena and other types of subjects? SQ5. What types of recognition support the recognition of a cultural phenomenon?
Q3. How does the representational evolution of subjects witness the emergence of cultural changes?(Panofsky & Saxl, 1933)	SQ6. What are the artworks that have both a cultural phenomenon in common and a different style associated?
Symbols	
Q4. How does the usage of symbols evolve?(Panofsky, 1955; Wittkower, 1977)	SQ7. Which and how many symbols express the same concept? SQ8. What different symbolical meanings can the same symbol have? SQ9. How does the usage of symbols change over time and in different contexts?

Q5. Is the symbolic meaning motivated by a specific source?(Christiansen, 1986)	SQ10. What are all the symbols motivated by a specific source? Do they differ from the corresponding subjects who do not cite a piece of evidence?
Q6. Is the deeper conceptual meaning motivated by sources?(Gombrich, 1972)	SQ11. Do deeper meanings cite as support textual sources?
Iconographies	
Q7. How does the usage of iconographies, their meaning, and attributes evolve over time?(Panofsky, 1955)	SQ12. What are the representational variations of iconography? What are the sets of level 1 subjects composing the recognized level 2 subject in artworks? SQ13. What attributes have symbolic meaning? SQ14. How does the representation of iconography vary over time? Of which level 1 subjects is it composed?
Q8. What are the attributes that allow us to recognize a subject?(Panofsky, 1955; van Straten, 2012)	SQ15. What attributes allow us to identify representations? What is their frequency? SQ16. Which are the most common and rare among the attributes marked as recognizing?
Q9. How does the representation of iconography vary?(Panofsky, 1972)	SQ17. What are the most common level 1 subjects not marked as recognizing?
Evidence	
Q10. What were the known textual sources to which the artwork refers, and what does this knowledge tell us about the thinking of the time?(Warburg, 1999)	SQ18. What were the known textual sources to which the artwork refers? Is the artwork involved in a cultural phenomenon?
Visual citation	
Q11. How do visual shapes migrate and reappear across cultures?(Warburg et al., 2020; Wittkower, 1977)	SQ19. What artworks cite the visual pattern of others?

Q12. Is a visual citation the evidence that documents a cultural phenomenon?(Warburg, 1999)	SQ20. In which case are those artworks involved in a visual citation also associated with a cultural phenomenon?

For each sub-question, we present how the questions were addressed and their results, showing 1) the performed SPARQL queries and 2) the results retrieved.

7.2.1 Cultural Phenomena

This section seeks to quantitatively analyze several aspects of cultural phenomena, reflecting on the core inquiries of iconological studies. First, the possibility of retrieving all the artworks with a cultural phenomenon is assessed (SQ1) to address the central question of what artworks witness a cultural phenomenon (Q1), constituting the core question of this type of study. Second, the sub-questions expressing Question 2 seek to quantitatively express one of the first analyses made by the iconological interpreters, namely the detection of the aspects that suggest the presence of a deeper meaning. We express this research interest by looking for 1) correlations of phenomena with other visual aspects, such as the subjects depicted, and 2) the support given by other interpretations or the presence of cited pieces of evidence. In detail, for the first aspect, we will examine the most frequent subjects at the other levels which occur with phenomena (SQ2) and co-occurrences in artworks of subjects and phenomena types (SQ4). The second aspect will be inquired by retrieving interpretations citing as support a piece of evidence (SQ3) and by highlighting eventual interpretations citing as support other ones (SQ5).

The last question of this section (Q3) concerns understanding subjects' evolution, which may witness the emergence of cultural changes. Assuming that artworks with different styles belong to different cultures, we examine whether artworks sharing common phenomena belong to different styles (SQ6).

7.2.1.1 SQ1. How many artworks have a cultural phenomenon associated?

The SPARQL query listed in 7.1 aims to retrieve all artworks associated with a cultural phenomenon. This way, an overview of the artworks involved in a cultural phenomenon can be provided.

Upon analyzing the Iconology Dataset, it was determined that **of a total of 428 artworks considered, 327 exhibit an association with a cultural phenomenon**. To further elaborate, as illustrated by the pie chart in Figure

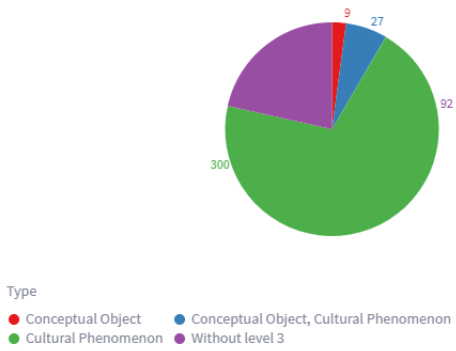

Figure 7.2: Number of artworks having a third-level subject, either a cultural phenomenon, a concept, or both of them

7.2, most of these artworks (specifically 300) are exclusively related to third-level subjects classified as cultural phenomena. In contrast, the remaining 27 artworks are simultaneously associated with a cultural phenomenon and a third-level concept.

7.2.1.2 SQ2. What is the frequency of level 1 and 2 subjects occurring with each cultural phenomenon?

This sub-question pertains to domain question Q2, which is focused on discerning the sources or visual attributes within artworks that drive the identification of a cultural phenomenon. Consequently, we approached this inquiry through a quantitative analysis, specifically by investigating the occurrences of subjects in conjunction with cultural phenomena. Since the registered phenomena are distinguished with a high level of detail, each is related to a low number of artworks. Hence, we performed the query on the 20 types of cultural phenomenon identified instead of the 379 individual phenomena to obtain more insightful results.

This investigation utilizes two separate queries for first- and second-level subjects, shown in Listings 7.2 and 7.3.

Upon interpreting the results, it should be considered that the same cultural phenomenon can belong to different types and that artworks can have multiple associated phenomena. Therefore, the same artwork can be repeated in different types.

For this analysis, we present three visualizations as bar charts. The first shows the frequency of the top 5 level 1 subjects that appear together with the type of phenomenon (Figure 7.3), followed by the same analysis repeated with second-level subjects (Figure 7.4). A third set of graphs reports the appearance percentage of level 1 and 2 subjects on the total artworks in which the phenomenon is registered. For this graph, we selected only the most represented phenomena types, defined as

the ones occurring in at least 50 artworks.

Considering level 1 subjects, the most represented one is the Natural Element *man*, as it is the most frequent of 18 out of 21 phenomena types, and it reaches the highest absolute frequency for the *reception of the classical antiquity* type, with a total of 111 artworks (53% of the total, as shown in figure 7.5(a)). *Wings* and *woman* are the following most frequent first-level subjects, registering respectively the maximum number of co-occurrences with the type *Iconographical Evolution* (29 occurrences of *wings*) and with *Reception of Classical Antiquity* (42 occurrences of *woman*). The consistent presence of the male figure is confirmed by the fact that it is the subject having the highest percentage of appearance in the most represented phenomena types. In fact, Figure 7.5(a) shows that, in all types of phenomena, the subject *man* has the highest percentage, with a minimum of 26% for the type *Formal or stylistic phenomenon* and a maximum of 57% for the type *Iconographical Interpretation*.

Moving the analysis to the iconographical level, most second-level subjects reported in Figure 7.4 belong to the classical culture. Therefore, classical deities such as Cupid, Jupiter, and Mercury appear in multiple phenomena types, registering high scores of frequency (see *Cultural Interpretation*, *Formal or Stylistic Evolution*, *Cultural Interpretation*, *Iconographical Evolution*, *Iconographic Interpretation*, *Reception of Classical Antiquity*).

As expected, the most represented type of phenomenon, namely the reception of classical antiquity, has a very high number of level 2 subjects, with the mode represented by the 16 co-occurrences of Saturn. Similarly, all the most frequent subjects for this type of cultural phenomenon are classical deities. However, the frequency of these level 2 subjects slightly decreases when related to the total number of artworks having the type of cultural phenomenon (206), showing only an 8% frequency of Saturn, the most frequent subject.

On the contrary, the high frequency of the subject registered for the type *Iconographical Evolution* partially emerges on the total percentage, showing that Saturn has a 14%.

This analysis shows that it is possible to retrieve the most frequent subjects of artworks grouped according to the type of cultural phenomenon associated with them. **Rather than showing a correlation of a single phenomenon type with the subjects, the results suggest some common tendencies**. For example, the strong representation of the male figure at the first level emerges, reflected in male characters at the second one, and the preference for subjects related to classical mythology.

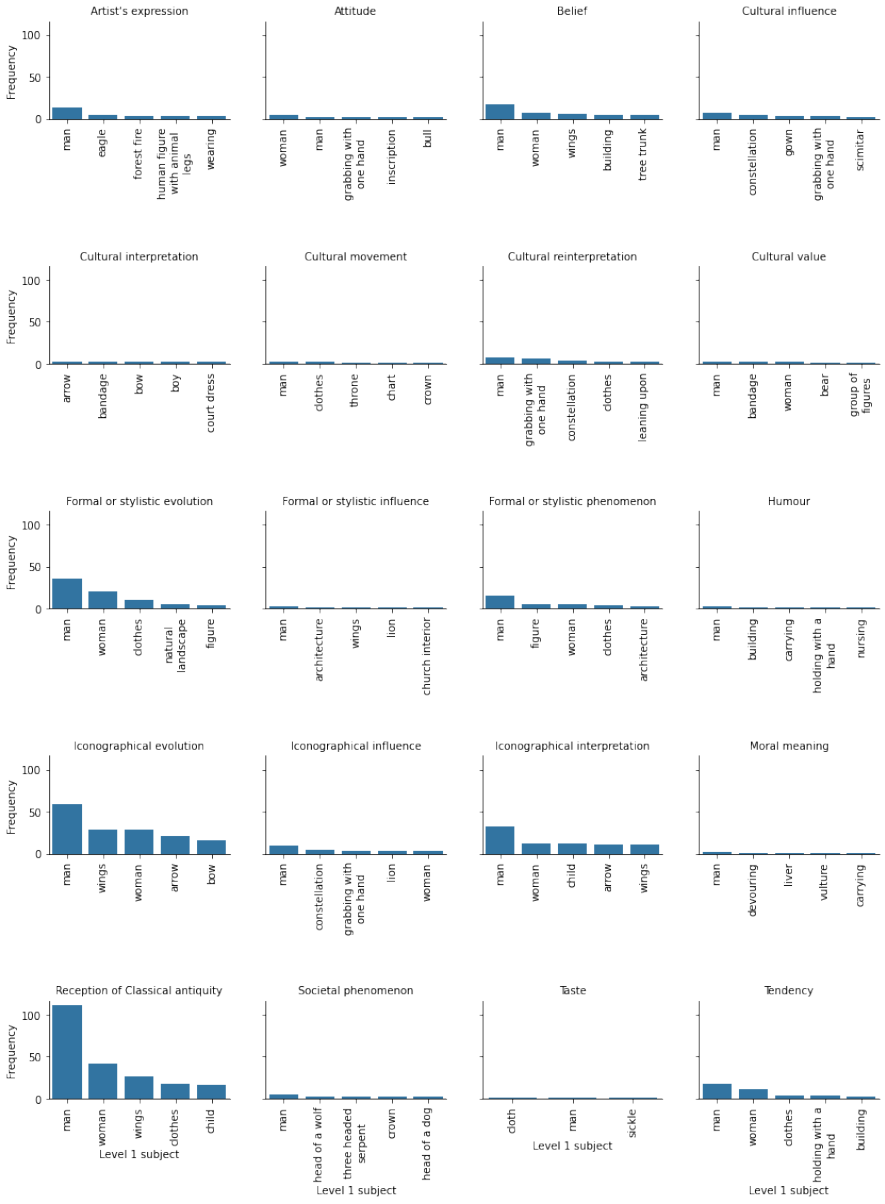

Figure 7.3: Results of SQ2, showing the co-occurrence of the most frequent 5 level 1 subjects, defined as the number of artworks in which they appear together.

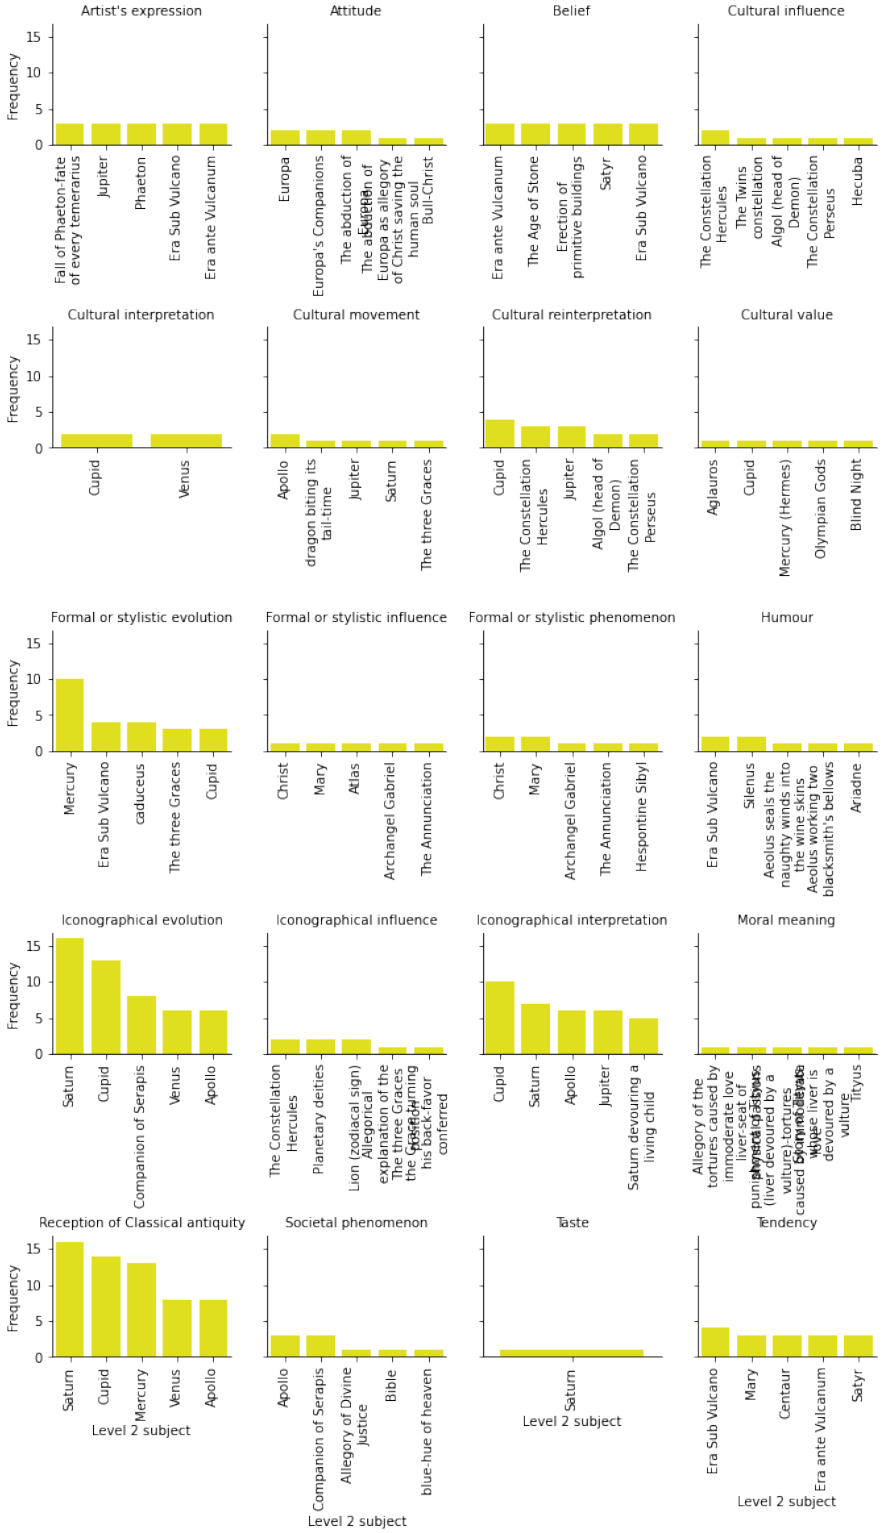

Figure 7.4: Results of SQ2, showing the co-occurrence of the most frequent 5 level 2 subjects, defined as the number of artworks in which they appear together.

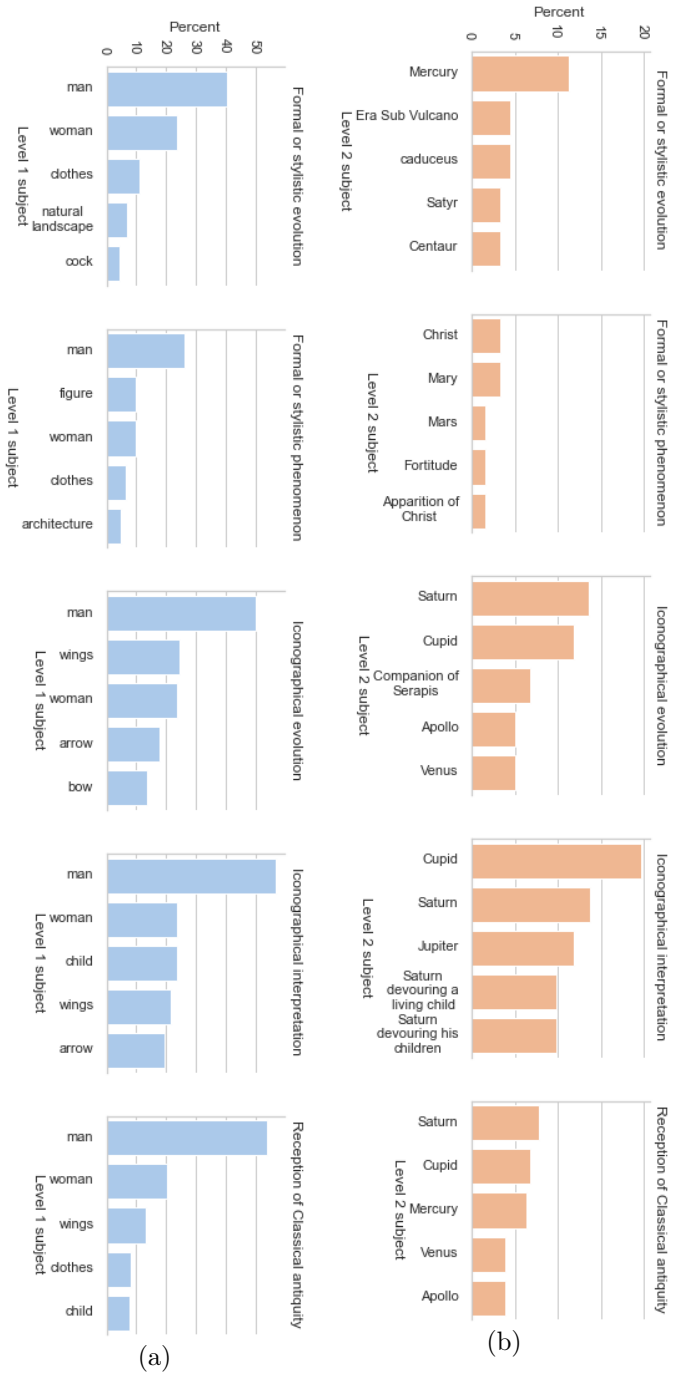

Figure 7.5: Results of SQ2, showing the percentage of the times the subject appears in artworks having the specified phenomenon type over the total of artworks in which the phenomenon is registered.

7.2.1.3 SQ3. What is the frequency with which a piece of evidence supports the recognition of a cultural phenomenon? Of what type is it? (texts, artworks, etc.)

This sub-question expresses the inquiry into the sources motivating the recognition of a cultural phenomenon so often cited during the analysis.

To quantitatively address it, we calculate the total number of Iconological Recognitions identifying a cultural phenomenon citing a piece of evidence, and we relate it to the total number of cultural phenomena Recognitions. Listing 7.5 shows the query to retrieve recognitions with these characteristics. In order to recall the total number of cultural phenomena recognitions, we repeat the same query, removing the `?rec cito:citesAsEvidence ?evid` SPARQL pattern.

As a result, **only 14% of such recognitions cite a piece of evidence** (see figure 7.21(b)).

We further examine the cited piece of evidence type by performing the query shown in Listing 7.4.

In the dataset, if a specific part of a text is cited, `cito:citesAsEvidence` property relates the recognition both to the specified passage (e.g., vv. 1-25) and to the text to which it belongs (e.g., Virgil's Aeneid). Consequently, to avoid duplicate results for the cases in which the specific portion of text is provided, we filter them out from the retrieved results through the SPARQL filter `FILTER NOT EXISTS ?evid2 crm:P106_is_composed_of ?evid`. Since the P106 relation represents that a text (here indicated by the variable ?evid2) has as part some portions of a text (?evid), we are stating that the piece of evidence we retrieve should not be of such type through the `FILTER NOT EXISTS` formula. For the same reason, we repeat the filtering whenever we retrieve evidence.

As shown by figure 7.6(a), **artworks are the most cited evidence for third-level recognitions of cultural phenomena**. While 84% of them cite as evidence an artwork, only 11% of such recognitions cite a textual source as evidence. Therefore, these results suggest that **deeper cultural interpretations are generally not directly based or supported by textual sources.**

7.2.1.4 SQ4. Is there any co-occurrence between cultural phenomena and other types of subjects?

As the iconological interpretation moves from the results of the iconographical analysis, this sub-question computationally expresses the need to seek in the iconographical study the hints for a more profound comprehension by highlighting the eventual correlation between subjects and phenomenon type registered in the artwork interpretations.

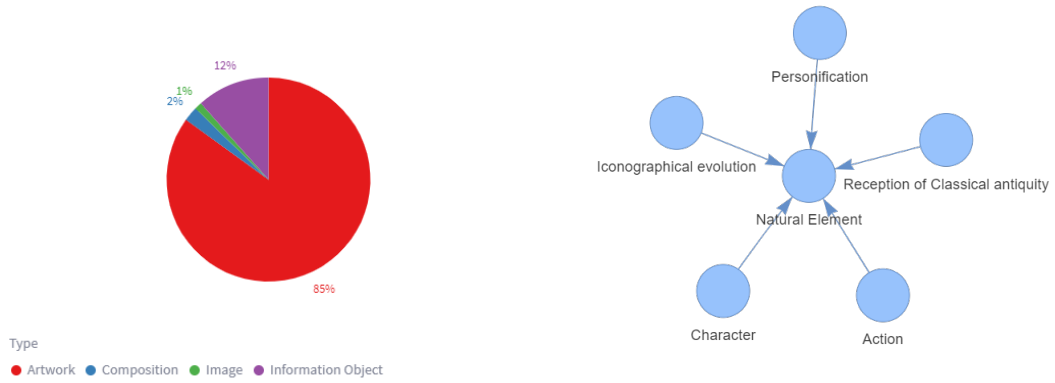

(a) (b)

Figure 7.6: a) Results of SQ3, showing the type of cited evidence; b) Results of SQ4, showing the co-occurrence of subject type (support: 0.20, confidence: 0.90)

In order to retrieve the co-occurrences of subject types, we extract, for each artwork, the types of subjects represented. For a better differentiation of cultural phenomena, we extract the types that were assigned to them. The SPARQL query to retrieve the type of subject represented in each artwork is shown in Listing 7.6.

Association rules can help discover frequencies and patterns in data. *Support* is an indicator that identifies the frequency with which a set of items appears together in the data. *Confidence* is the measure that indicates in which measure it is likely that an item of the same set is present if another element is present. The closer these measures are to 1, the more the rule is supported.

We measured the co-occurrence value according to these parameters using the Python library Mlxtend (Machine Learning extensions) (Raschka, 2018)[2]. These parameters can be dynamically changed on the online dashboard[3]. In this thesis, we analyze the set of relations obtained by setting a high grade of confidence and support, respectively, to 0.20 and 0.80, to see if there is a strong co-occurrence of subject types. It results that the types appearing more than 20% of the time are Actions and Natural Elements for level 1, Characters and Personifications for level 2, and Iconographical Evolution, Reception of Classical Antiquity, and Formal or Stylistic Evolution for level 3 (see Table 7.2. The most frequent ones are the Natural Element type, appearing 80% of the time, and Reception of classical antiquity (support: 0.63).

The sets obtained are the basis for calculating the associated items that each item is likely to have (confidence). Figure 7.6(b) graphically represents the asso-

[2] https://rasbt.github.io/mlxtend/

[3] Available at https://iconology-dataset.streamlit.app/Dataset_Overview, Section *Co-occurrences*

Table 7.2: Overview of the subject types having a support score greater than 0.20

Support	**Set**	**Support**	**Set**
0.80	'Natural Element'	0.34	'Iconographical evolution', 'Natural Element'
0.63	'Reception of Classical antiquity'	0.28	'Iconographical evolution', 'Reception of Classical antiquity'
0.54	'Natural Element', 'Reception of Classical antiquity'	0.27	'Formal or stylistic evolution'
0.49	'Character'	0.27	'Iconographical evolution', 'Natural Element', 'Reception of Classical antiquity'
0.49	'Natural Element', 'Action'	0.27	'Natural Element', 'Character', 'Action', 'Reception of Classical antiquity'
0.49	'Action'	0.27	'Character', 'Action', 'Reception of Classical antiquity'
0.47	'Natural Element', 'Character'	0.24	'Natural Element', 'Iconographical evolution', 'Action'
0.37	'Natural Element', 'Action', 'Reception of Classical antiquity'	0.4	'Personification', 'Natural Element'
0.37	'Action', 'Reception of Classical antiquity'	0.24	'Iconographical evolution', 'Action'
0.36	'Character', 'Reception of Classical antiquity'	0.24	'Personification'
0.36	'Iconographical evolution'	0.21	'Formal or stylistic evolution', 'Reception of Classical antiquity'
0.35	'Character', 'Natural Element', 'Reception of Classical antiquity'	0.21	'Iconographical evolution', 'Character'
0.35	'Natural Element', 'Action', 'Character'	0.21	'Iconographical evolution', 'Natural Element', 'Character'
0.35	'Action', 'Character'	0.20	'Natural Element', 'Formal or stylistic evolution'

ciations, relating the antecedents to the consequents with a directed arrow. It is evident that only two phenomena types, namely the reception of classical antiquity and iconographical evolution, present a strong relation with Natural Elements.

7.2.1.5 SQ5. What types of recognition support the recognition of a cultural phenomenon?

As illustrated in Section 5.1.2, recognition of cultural phenomena can be motivated by understanding other characteristics of the artwork under examination or related ones, such as subjects, visual patterns, concepts, and other phenomena. For example, by taking the real case study presented for Q2, the recognition of the overall meaning of the Medici Chapel as a political manifesto of the future family power is based on the recognition that the Magi bear the portraits of the rulers.

With this sub-question, we aim to retrieve which type of recognition is explicitly illustrated as the starting point for a deeper cultural phenomenon interpretation. In the dataset, this type of relation is represented by the relation `cito:givesSupportTo`. Therefore, we retrieve all types of recognition that support the identification of a cultural phenomenon by performing the query illustrated in Listing 7.7.

The results show that only visual motif recognitions and iconological recognitions support a phenomenon identification, with, respectively, 57 and 46 recognitions. Consequently, **at the other levels, only visual motif recognitions are explicitly highlighted as having a key role in supporting a deeper understanding of the phenomenon. The fact that nearly half of the results concern citations among recognition of the same level suggests that there is a narrow relation among recognized phenomena.**

Table 7.3: Overview of the recognitions giving support to the recognition of a cultural phenomenon

Supporting Recognition type	Supported Cultural Phenomenon Recognition	Frequency
Formal Motif Recognition	Iconological Recognition	57
Iconological Recognition	Iconological Recognition	46

Nevertheless, another strategy of tracking the dependency is adopted in the dataset, namely relating the intrinsic meaning with the specific part of the artwork to which the phenomenon refers explicitly. Despite being slightly different from an explicit recognition supporting another one, it can give hints on which subjects are related to a phenomenon, according to the author of the recognition. Listing 7.8 shows the query performed.

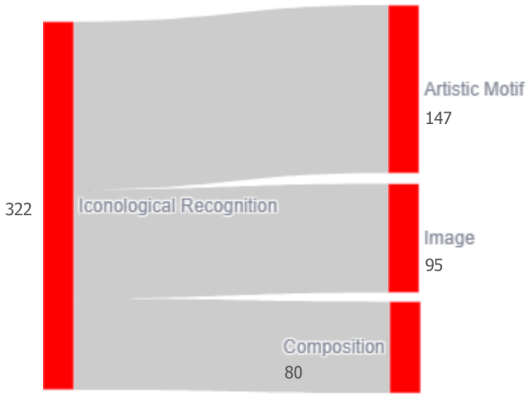

Figure 7.7: SQ5 results: Sankey diagram showing the frequency with which recognitions of Cultural Phenomena refer to a specific part of the artwork

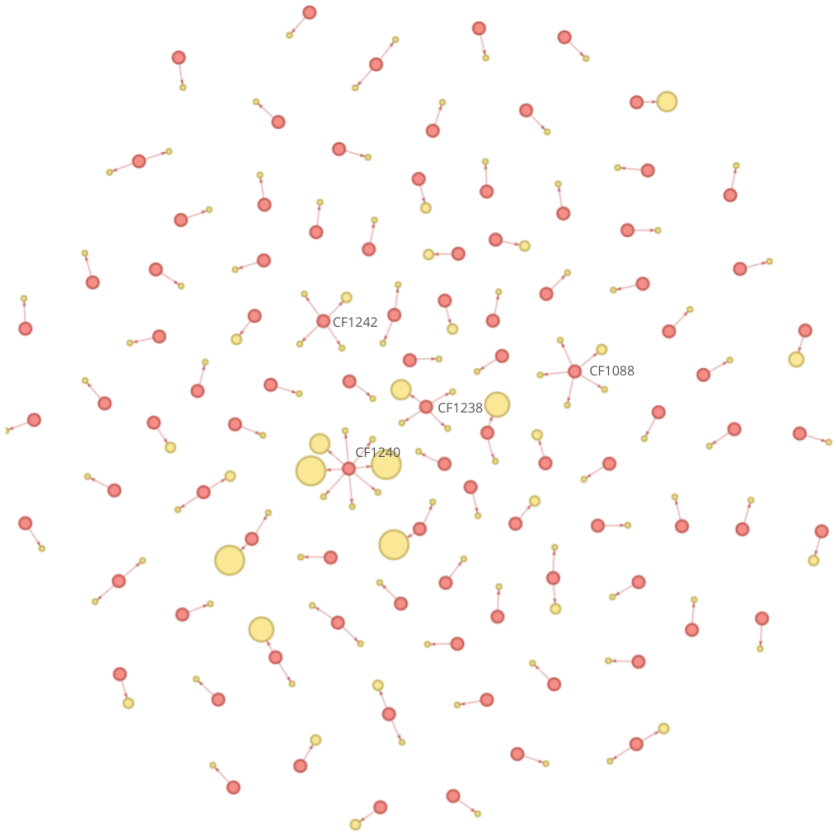

Figure 7.8: SQ6 results: network of cultural phenomena associated with artworks having a specific style

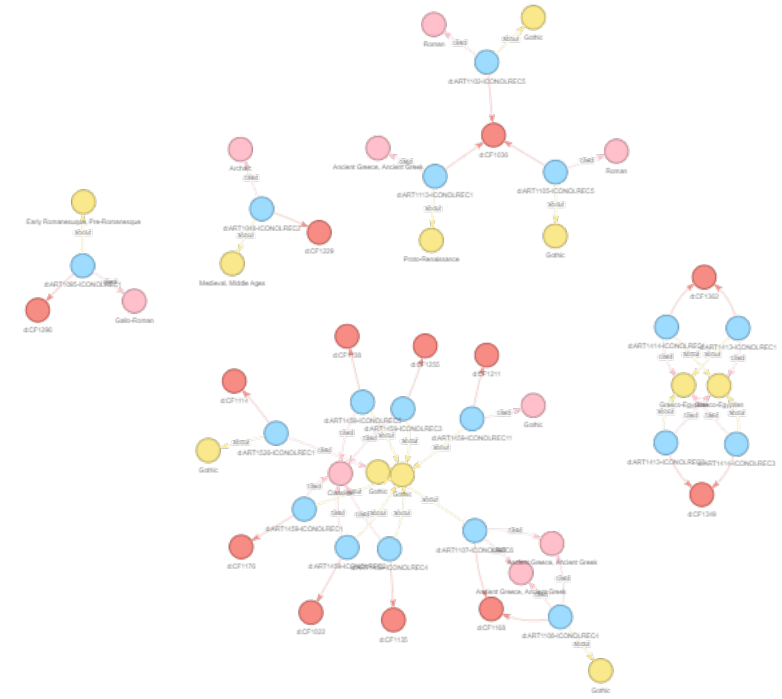

(a)

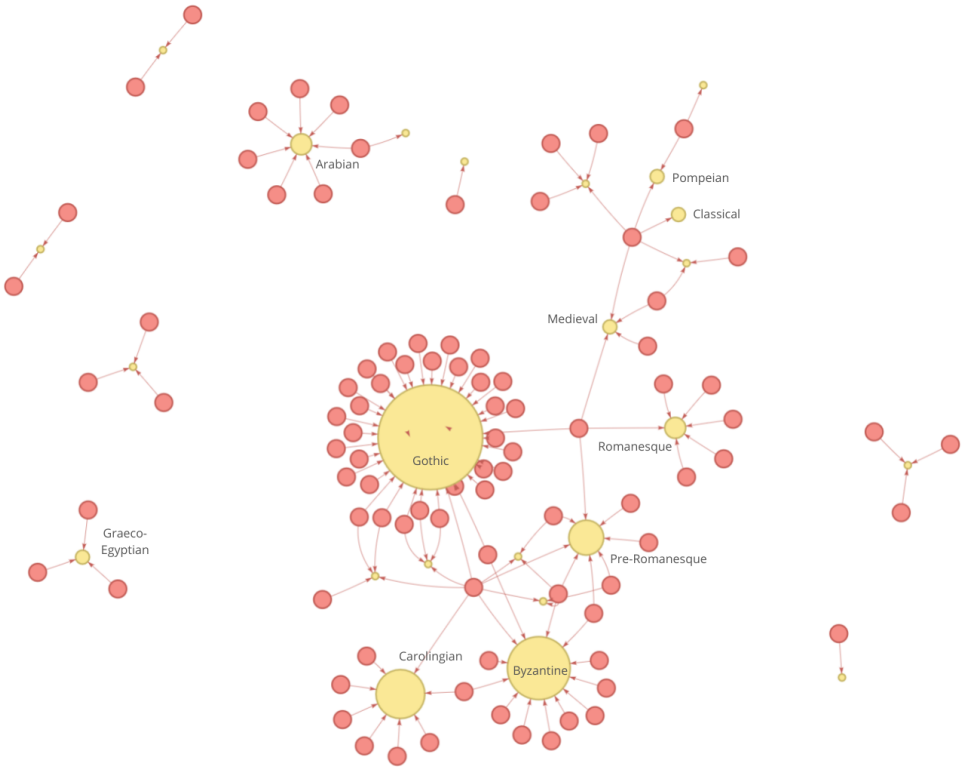

(b)

Figure 7.9: a) SQ6: artworks having both a cultural phenomenon and a style, citing artworks having a style associated, b) SQ6: artworks with a cultural phenomenon and a style, grouped by common style

Figure 7.7 shows that **recognitions that specifically refer to a portion of the artwork reference quite equally images, artistic motifs, and compositions**. This specification frequently occurs, registering a score of 322 occurrences of this relation. In detail, this result concerns 142 unique recognitions (over a total of 853 cultural phenomena recognitions), and 253 cited portions of artworks.

7.2.1.6 SQ6. What are the artworks that have a cultural phenomenon in common and a different style associated?

A good number of iconological studies address the theme of transmission of the characteristics and iconographies of the artworks in a different culture, highlighting their process of reworking and resemantization. We computationally express this inquiry by retrieving those artworks having a common cultural phenomenon and a different style specified, since it witnesses that the art historian(s) expressing their interpretation addressed a remaking of artworks through cultures. Listing 7.9 shows the query performed to identify such artworks.

Despite a good number of results being recovered, it must be noted that only 76 out of 428 described artworks have a specific style, consequently narrowing down the analysis to 18% of the artworks.

To better envision the relations between cultural phenomena and the artworks with associated styles, we retrieve all the artworks and phenomena that have these characteristics and visualize them in a network. In Figure 7.8, we represent the results, distinguishing cultural phenomena (red) from artworks (yellow)[4]. Artworks having the same style and cultural phenomenon are represented in a unique node, and the number of artworks belonging to it is represented by the node size.

Some clusters emerge from the network. The cluster having the highest number of nodes (9) and of artworks, grouped around the phenomenon *Renascences of classical antiquity in Western Art before the Renaissance*(CF1240), puts in relation styles very close to each other, namely the ones preceding the Renaissance movement in which types of classical citations started to appear (e.g., Gothic, school of Burgundy), including one artwork of Byzantine style, and 6 from the Carolingian period. The overall period covered spans from the VIII to the XIII century. Despite the classical style being theoretically related to this phenomenon, it does not emerge from this type of query, as a classical cultural object cannot embed the phenomenon of its own citation from a poster period. In order to highlight this type of relation, we further enrich the network with eventual artworks cited as evidence having a style associated, which will be discussed in the following.

[4]An interactive version of the network is available at https://iconology-dataset.streamlit.app/Analysis_of_cultural_phenomena > 'Style' > 'Artworks related to common cultural phenomena'

The second relevant cluster, with a total of 5 relations and 6 artworks, is formed by the phenomenon *Evolution of the Iconography of Time*. It includes various styles from different periods and cultures, including variations of a classical style (classical, late Roman, Pompeian), a style of German area, and the medieval one, highlighting the evolution of iconography through cultures.

Similarly to the first cluster examined, a great amount of the remaining ones show pre-Renaissance styles, as they belong to more specific phenomena related to the main one described above. Such phenomena include the classical renascenses in the Pre-Romanesque context (CF1238), the first phase of the Proto-Renaissance (CF1296), the birth of intrinsic classicism born by the encounter of Gothic and Proto-Renaissance (CF1167) and fully developed in the School of Reims (CF1117), and the aspect of the *monumentalization* of the Byzantine Style (CF1214).

Other clusters highlight the phenomenon in general terms, focusing on the representation of classical content with contemporary motifs (CF1242) and on the phenomenon of classical motifs recovery (CF1036).

Interestingly, a few other phenomena are represented, such as the influence of the Arabian East on a Late-Gothic artwork (CF1162), classical variations of Eros and Anteros (CF1043) among two artworks marked as Roman and Pompeian, and the expression of the values of a culture in the theory of proportions (CF1346), shown, respectively, in Byzantine and Gothic artworks.

As the art style specification provided reflects the one given by the art historian making the interpretation, **data seems to suggest that Panofsky overrepresented this information for the artifacts involved in the study concerning the early renascence of classical style in the proto-Renaissance period, as style is a crucial aspect for this type of analysis**. On the contrary, the style information is not explicitly stated for the major part of the remaining artworks.

We further examine the **citation relations** that may exist among artworks that have an associated style. As anticipated, the visual citations of previous styles do not emerge from the cultural phenomenon network. Hence, we extracted all the recognitions of phenomena about artworks having a style associated that cite as evidence an artwork having a style. The results are shown in figure 7.9(a), in which recognitions (blue), phenomena (red), artworks (yellow), and cited artworks (pink) are represented.[5]

As expected, **the majority of the groups examine an artwork from a more recent period that cites a previous one**. Of the 12 examined artworks, 9 belong to one of the styles of the Middle Ages, and recognitions concerning them

[5]An interactive version of the network is available at https://iconology-dataset.streamlit.app/Analysis_of_cultural_phenomena > 'Style' > 'Style of cited artworks'

cite mostly styles related to the classical period (Roman, Greek, Archaic), as 7 out of 8 belong to this group.

By grouping the artworks by style, it is possible to see cultural phenomena that are potentially shared among the same stylistic group and to make hypotheses for interconnected phenomena. Expert knowledge, registered in the form of structured data, can, in this way, be the starting point for further research hypotheses. Figure 7.9(b) shows a graph of the results of the previous analysis in which artworks that have a common style are represented as one unique node. Their number is indicated by the size of the node. In this way, it is possible to relate the cultural phenomena belonging to artworks having the same style, seeing which phenomena are potentially of interest for the interpretation of artworks of the same style.

7.2.1.7 Summary

In this section, various key aspects of iconological analysis are treated and quantitatively expressed. SQ1 demonstrated that it is possible to retrieve all the artworks that witness a cultural phenomenon. SQ2 and SQ4 were designed to find correlation of phenomena with other types. Rather than highlighting co-occurrences, SQ2 shows that the most represented subjects in the group of artworks in which the same phenomenon type is present depicts a majority of male figures (level 1) and classical deities (level 2), confirming the results obtained in the introductory analysis. Poorly significant results could be established by SQ4, in which subject and phenomenon types co-occurring in artworks were analyzed.

The analyses focusing on evidence and support reported relevant results. Surprisingly, a small number of phenomena recognitions cite as evidence a textual source (SQ3), suggesting that such interpretations are not generally directly motivated by a text. Furthermore, the analysis of the recognitions supporting a third-level one (SQ5) highlights that, among the other levels, only the recognition of a visual citation (i.e., formal motif recognitions) directly supports the phenomenon understanding. Therefore, no further hints on which other subjects are considered fundamental for the phenomenon of recognition are provided. On the contrary, it seems that other phenomena recognitions support recognitions of the same type, suggesting a narrow relation among recognized phenomena.

The last analysis of this section (SQ6) examined the style in relation to cultural phenomena. Although limited to 18% of the artworks, it showed relevant results. In the networks, clusters of artworks having similar styles emerge, showing the consistency of a phenomenon in contemporary variations. On the contrary, some clusters included artworks from distant styles, reflecting diachronical inquiries, such as the evolution of iconography.

Additionally, we examined the network of artworks with a specified style, visually citing others with a style. As the majority of them concern artworks belonging to more recent styles that cite classical ones, the tendency of reusing visual motifs from classical artworks is supported by data.

7.2.2 Symbols

In this section, we quantitatively explore two types of questions motivated by the real cases presented in Section 5.1.2. First, the evolution of symbols is explored, considering either the changes in how concepts are represented or what meanings a symbol may express, along with the possible influence that the presence of a cited source may have on it. Secondly, we explore the influence of textual evidence on the recognition of third-level concepts.

7.2.2.1 SQ7. Which and how many symbols express the same concept?

This sub-question aims at detecting one aspect of a symbol variation, namely the forms in which a symbolic meaning can be expressed by different symbols and how many times the same symbol expresses the same concept. This query allows us to detect the most common and rare symbols and constitutes the starting point for further investigation of how the representation of concepts varies and evolves in visual arts. Listing 7.10 shows the query performed on the dataset, where the variables `?simulacrumL` and `?realityL`, respectively, refer to the symbol and its symbolic meaning, following the Boudrillard terminology (i.e., *Simulacrum* and *Reality Counterpart*) used in the Simulation Ontology, adopted to describe symbols of the dataset.

Figure 7.10(b) shows the results, representing the symbolic meanings on the left and the symbols on the right.[6] Only those **meanings related to more than one symbol** are shown, obtaining **18** symbolic meanings with these characteristics **out of the total of 144 symbolic meanings (12%)**. The most well-represented symbol variations count three symbols per meaning and relate to the meanings of the three forms of time (the past, the present and the future), and love.

7.2.2.2 SQ8. What different symbolical meanings can the same symbol have?

Similarly, art historians studying the evolution of symbols are interested in knowing which symbolic meanings the same symbol can express. Hence, we retrieve all the symbols having different meanings through the same SPARQL query shown

[6] See the interactive version at https://iconology-dataset.streamlit.app/Analysis_of_concepts_and_symbols > Overview > Symbolic meanings

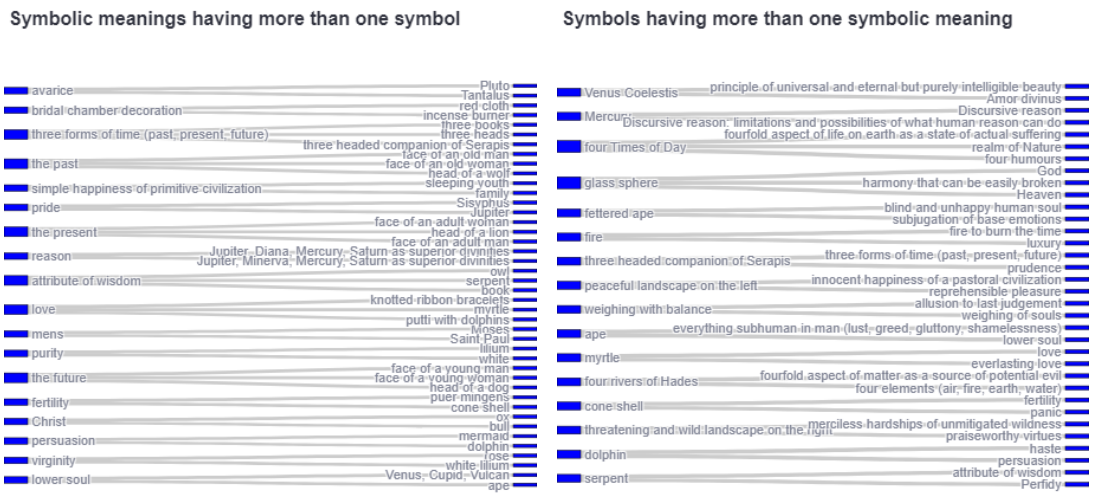

(a) (b)

Figure 7.10: 1) Symbolic meanings having more than one symbol associated (SQ7), and 2) symbols expressing more than one meaning (SQ8)

in Listing 7.10. However, a different grouping is applied to the results to obtain the symbolical meanings of the same symbol.

Results show that **only 16 symbols out of 150 have multiple meanings (11%)**. The ones registering a higher variation, with three meanings associated, are two classical deities (Venus Coelestis, as interpreted by the Neoplatonic theory, and Mercury), the personification of the times of a day, and the glass sphere.

7.2.2.3 SQ9. How does the usage of symbols change over time and in different contexts?

67% of the artworks have a date specified. As the diachronic evolution of a symbol is of great value for this kind of study, we retrieve the date of creation of artworks having a symbol and, if present, the period of the artwork production and the context associated with the symbol by querying the dataset with the SPARQL code shown in Listing 7.11.

Of the 178 symbols, 102 are associated with an artwork's date. Despite the query being performed successfully, **the diversity of symbols that emerged from the results of SQ7 and SQ8 prevents us from exploring the diachronic evolution of a specific symbol or symbolic meaning**. Therefore, we explore them through a comprehensive visualization shown in Figure 7.11(a).[7] In the chart,

[7]See the interactive version at https://iconology-dataset.streamlit.app/Analysis_of_concepts_and_symbols >Over time > Overview.

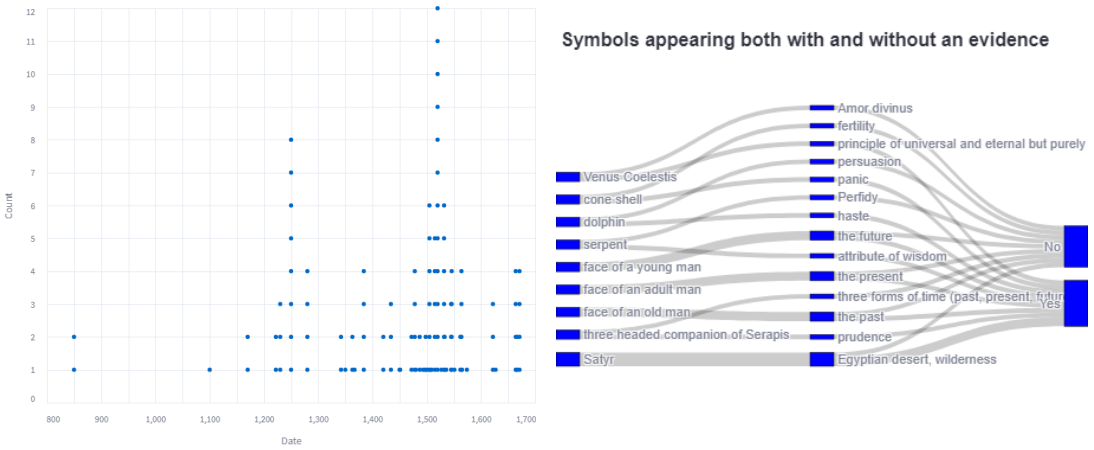

(a) (b)

Figure 7.11: 1) Symbols over time and their frequency for each date (SQ9), and 2) symbols having and not having a piece of evidence (SQ10)

symbols, represented by dots, are ordered according to the year of creation of the artwork, and a description of each symbol appears when passing over it with the cursor in the interactive visualization available on the online dashboard. For the sake of clarity, we stack the individuals sharing the date. As the date is usually specified as a range of probable (e.g., a century) or actual periods of creation, we computed dates for such ranges of time as the average of the starting and ending dates of each period. We adopted this criterion in every visualization involving a time. From the graph, it emerges that the majority of the symbols belongs to the XVI Century, as it is the most represented period in the dataset (see figure 7.1).

7.2.2.4 SQ10. What are all the symbols motivated by a specific source? Do they vary from the corresponding subjects not citing a piece of evidence?

This question aims to detect those symbols that were interpreted based on a specific source. The query in Listing 7.12 retrieves all symbols, along with evidence associated with the recognition itself or the symbol.

As we assume that the availability of evidence influences the interpretation of the symbol's meaning, the analysis focuses on those symbols that have and do not have a piece of evidence as support. It emerges that there are **49 occurrences of a symbol and a piece of evidence against 137 times in which the symbols do not have any evidence associated**.

In graph 7.11(b), we extracted the symbols (i.e., Simulacra) that appear in both groups, as they appear in recognitions in which they have and do not have evidence. We linked them to the meaning and eventual presence of a cited source. **Only nine**

symbols are shared among the two groups, and five present a meaning variation associated with a variation in the presence of evidence. For example, the symbol 'dolphin' has the meanings 'persuasion' and 'haste'. Whereas the latter meaning is related to a piece of evidence, this does not happen for the former.

This type of analysis can give a better overview of the influence that a source may have in attributing a specific meaning to a symbol.

7.2.2.5 SQ11. Do deeper meanings cite as support textual sources?

Like phenomena, third-level concepts express the understanding of artworks at a deep level. Therefore, they deserve an analysis of the basis on which the meaning is stated. To this end, we retrieve all the third-level recognitions recognizing a concept and citing as evidence a text. We compare them to 1) the total number of such recognitions citing evidence performed through the query shown in Listing 7.13, and 2) the total number of level 3 concept recognitions.

Type	Total	Percentage
Recognition of a third-level concept citing as evidence a text	5	9,7%
Recognition of a third-level concept citing some evidence	7	13,4%
Recognition of a third-level concept	52	

Table 7.4: Results of SQ11 showing the number of third-level concept recognition citing as evidence a textual source

As a result, only five recognitions of this type cite a text, out of 7 recognitions citing a piece of evidence, over 52 recognitions of this type (see table 7.4).

7.2.2.6 Summary

In this section, we explored how data can provide information for studying symbols and concepts. Although it was possible to perform all the analysis, only a few results were provided when analyzing the variation of symbols and symbolical meanings (SQ7, SQ8), showing variants for a limited number of items. Consequently, it was not possible to analyze the diachronical evolution of single symbols, limiting the analysis to an overall overview (SQ9). Nevertheless, some relevant results were obtained. SQ10 shows that it is possible to quantitatively retrieve and visualize the variations of a symbolic meaning having and not having a piece of supporting evidence. SQ11 highlights that it is possible to retrieve more profound concept

recognitions that cite a piece of evidence despite the limited results confirming the tendency of not making extensive use of text citations already registered for cultural phenomena recognitions.

7.2.3 Iconographies

The study of iconography is the core subject of iconographical studies, as demonstrated by the various types of aspects explored by domain studies presented in Section 5.1.2. Hence, several sub-questions are formulated to express this domain-specific well-established interest.

Question 7 concerns the evolution of the iconography, considering the level 1 subjects from which it is composed, the eventual meaning associated with the iconography or parts of it, and this kind of variations observed over time. We address these topics in three sub-questions. The first one (SQ12) retrieves all the iconographies (i.e., level 2 subjects) having a description at the first level, and we compare the iconographies presenting a variation in how they are represented. The second sub-question (SQ13) retrieves all the level 1 subjects that are part of iconography and have a symbolical meaning associated. Finally, SQ14 repeats the analysis performed in SQ12, but analyzes the variations of iconography over time, by retrieving the date of artworks.

Question 8 aims at analyzing the level 1 subjects that are marked as recognizing, namely the *attributes* that allow an observer to identify the specific character or artistic theme represented. SQ15 focuses on such recognizing attributes, aiming at retrieving all the relevant level 1 subjects that contribute to its recognition and exploring their frequency. SQ16 goes further in the interpretation since it aims to explore which attributes of this type are the most common and rare for each iconography.

Finally, Question 9 focuses on the expression of level 1 subjects not marked as recognizing. Indeed, iconological studies use to focus on unconventional representational details to discover deeper meanings. Query SQ17 aims at highlighting recurring level 1 subjects that may be usually obscured by the major frequency of recognizing attributes.

7.2.3.1 SQ12. What are the representational variations of iconography? I.e., what are the sets of level 1 subjects composing the recognized level 2 subject in artworks?

This sub-question aims to retrieve all the sets of level 1 subjects by which a level 2 subject is composed, in order to examine the variations that occur in the representation of an iconography in different artworks.

The relation `icon:recognizedArtisticMotif` relates the recognitions of second-level subjects with the first-level ones by which they are composed. In order to retrieve them, we perform two queries. The first one retrieves all the iconographies referring to either an Artistic Motif or Composition (Listing 7.14), whereas the second one retrieves the actual level 1 subjects linked to the artistic motifs or compositions obtained with the previous query (see Listing 7.15).

In order to give a fair representation of variations, we filtered results by keeping only those iconographies appearing more than one time. As a result, we obtained 155 iconographies, out of a total of 574 (27%). To better see the level 1 variations, we maintained the level 1 subjects group from which the iconography is composed. As a way of example, we report in table 7.5 the set of level 1 subjects of which the Character *Christ* is composed, along with the total occurrences of the iconography, the artworks in which the variation is found, and the percentage of times in which the variation appears over the total of the iconography occurrences. On average, level 2 subjects have 2.7 different sets of level 1 subjects per iconography.

Iconography	**Level 1 subjects cluster**	**Number of occurrences**	**Artwork**	**Iconography occurrences (total)**	**Percentage of variation occurrence**
Christ	'man'	2	ART1162 ART1537	10	20%
Christ	'cornerstone of a church'	1	ART1017test	10	10%
Christ	'dead body of a man'	1	ART1543	10	10%
Christ	''deer', 'man', 'carrying on shoulders', 'cloth"	1	ART1004	10	10%
Christ	''ground', 'standing on the ground"	1	ART1233	10	10%
Christ	''holding with the right hand', 'halo of rays', 'man', 'sceptre', 'lifting one hand"	1	ART1458	10	10%
Christ	''man', 'crown', 'cross"	1	ART1264	10	10%
Christ	''condemning gesture', 'man"	1	ART1538	10	10%
Christ	''child', 'halo"	1	ART1187	10	10%

Table 7.5: Example of results obtained by performing SQ12

Table 7.6: Overview of the iconographies having a description at the first level and appearing more than once (SQ12)

Description	**Count**
Number of iconographies having a level 1 description and appearing more than once	155
Total level 2 subjects	574
Average of level 1 different cluster per subject	2,75

As a significant number of iconographies have a level 1 variation, we considered providing a visualization of each iconography to allow a qualitative analysis of the

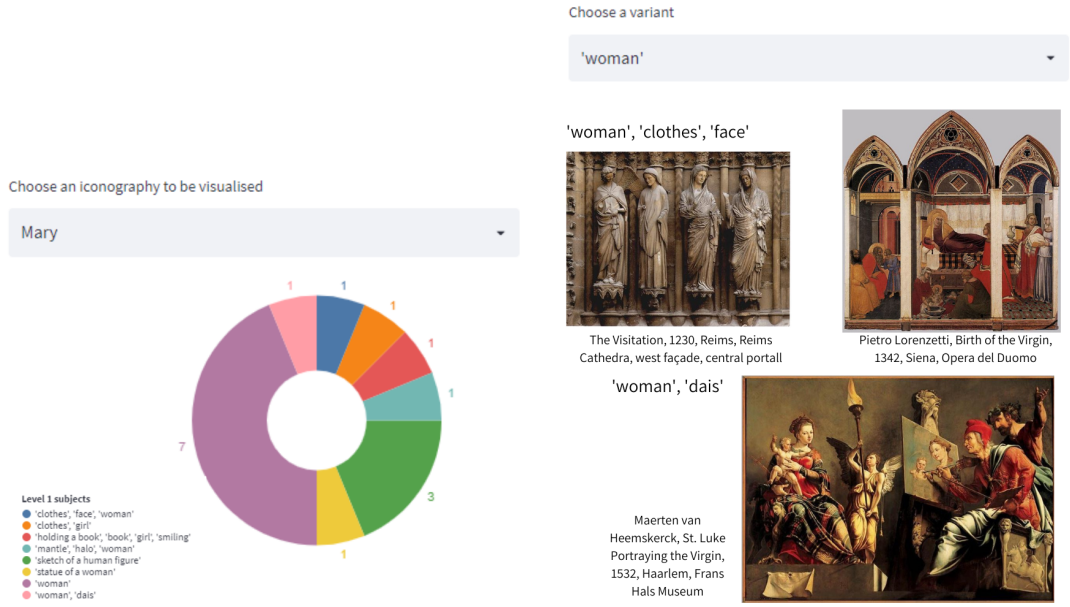

(a) (b)

Figure 7.12: a) Interface for having an overview of a single iconography variations (SQ12) and b) for exploring the artworks in which the variation is found

variations. Therefore, we developed an interactive interface where users can select the iconography to explore.[8] A pie chart displays the frequency of level 1 subject clusters that the selected iconography has (see figure 7.12(a)). Furthermore, it is possible to see in which artwork the iconography appears with each variation and, if available, see the artwork image. Figure 7.12(b) illustrates such an interactive interface, summarizing some results retrieved by selecting different level 1 clusters.

Despite being challenging to provide a quantitative overview of the overall iconographical variations which provides meaningful insights into data, **this type of inquiry can provide an overview of the variations of a single iconography which can support qualitative iconographical studies**.

7.2.3.2 SQ13. What are the attributes having a symbolic meaning?

In the current analysis, we aim at retrieving the iconographies having attributes with a symoblical meaning. To this end, the SPARQL query (see Listing 7.16) retrieves all the iconographies in which one or more of the level 1 subjects of which they are composed is involved in a symbol. It has to be specified that the level 1 subject appears both in the iconography and in the symbol in the context of the same artwork. This query is made possible by the distinction of general level 1 subjects and their uniquely identified manifestation in the artwork, namely the

[8] Available at https://iconology-dataset.streamlit.app/Analysis_of_iconographies > Variation of iconographies.

Artistic Motif. Such a query allows us to examine both the level 1 subject taken into account and the related symbolical meaning.

Results are shown in figure 7.13. Iconographies are displayed on the left, the level 1 subject is on the center, and the symbol and symbolic meaning are on the right.

We obtained 28 iconographies belonging to this type, 38 level 1 subjects having a symbolic meaning, and 41 symbols to which they are associated. On average, there are 1,3 level 1 subjects with a symbolic meaning per iconography. Relations in the figure allow us to make further considerations.

In three cases, the level 1 subject has different meanings, respectively a cloth, a male head, and a woman. This occurs when there are multiple Artistic Motifs in the same artwork recognizing the same level 1 subject. Indeed, the symbol description in the right column slightly differs (e.g., white and blue cloth).

Furthermore, this analysis allows us to see that several attributes having a symbolic meaning are shared among iconographies.

In detail, the subject *dragoon biting its tail*, with the symbolic meaning of time, appears both in the iconography of Saturn and of Time the Destroyer, for a total of 3 different artworks. Similarly, the wings, the symbol of the four seasons, appear both in the iconography of Time and in the one of the *Temps*.

The remaining subjects appearing in multiple iconographies concern the cases in which a single person embodies two iconographies (for example, Cupid can be recognized also as a personification of Love). This case concerns 1) the glass sphere, associated with the bride represented as Venus, in ART1339[9], 1) the flame, associated with the Salamander and the personification of the Right expressed by the same man in ART1013test,[10] 3) the myrtle wreath, related to Cupid, depicted as a personification of love, (ART1338)[11] and 4) Giuliano de' Medici, depicted with the iconography of a votary of Jupiter, having a scepter, in ART1375.[12]

To summarize, **this sub-question allowed us to delve into iconographies related to a symbol, exploring not only the symbolic meaning variations but also the symbols shared by multiple iconographies**.

7.2.3.3 SQ14. How does the representation of iconography vary over time? Of which level 1 subjects is it composed?

To answer this question, we retrieved the iconographies along with their level 1 subjects, artwork's ID and date with the SPARQL query in Listing 7.17.

[9]Tiziano Vecellio, *Education of Cupid*, 1560-1565, Rome, Galleria Borghese

[10]Parigi, Bibliothèque Nationale de France, Marco Polo, *Livre des merveilles*, XV sec, ms 2810, folio 24r. Anonymous, *salamander*. Source gallica.bnf.fr / BnF

[11]Tiziano Vecellio, *Allegory of the Marquis d'Avalos*, Paris, Louvre

[12]Michelangelo, *Statue of Giuliano de' Medici*, Florence, S. Lorenzo

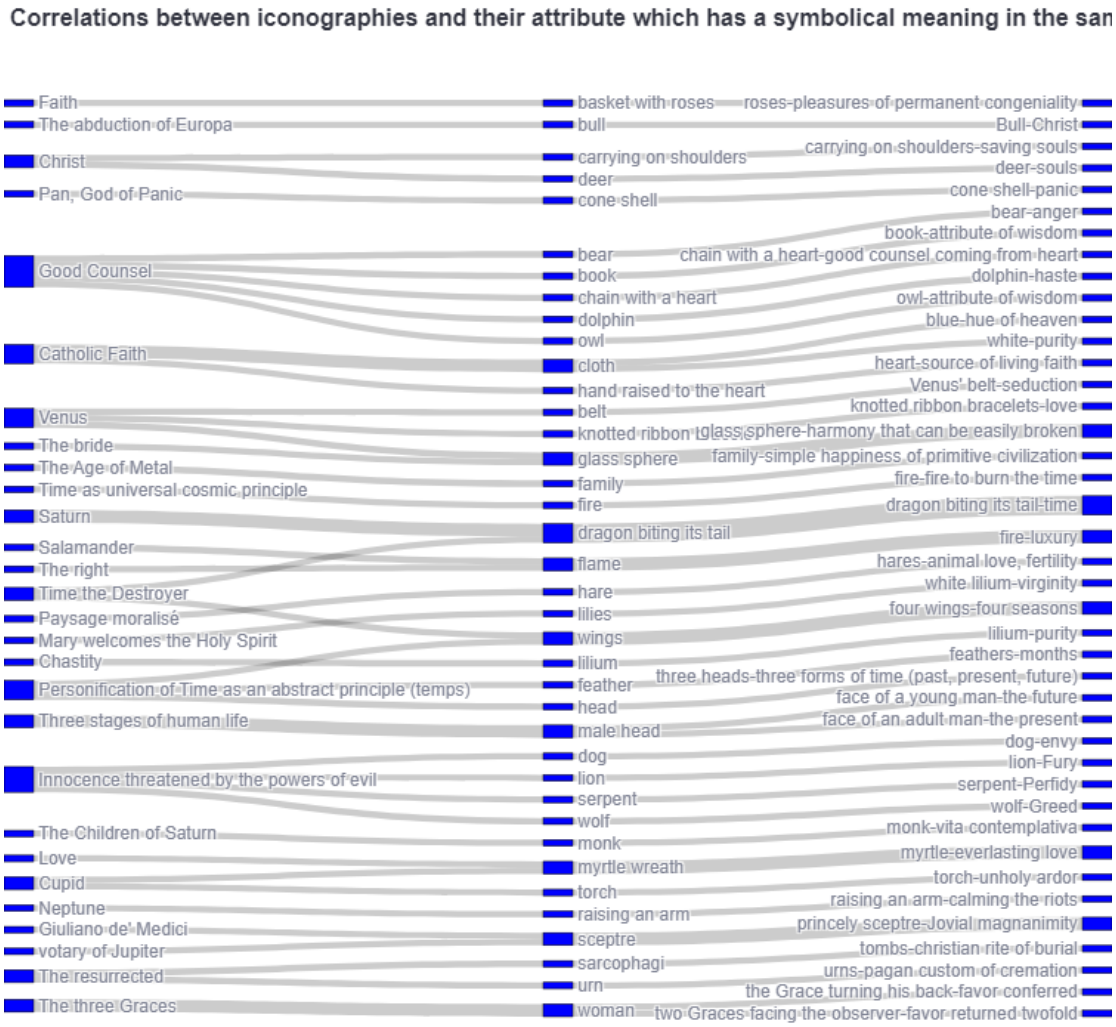

Figure 7.13: Overview of the iconographies with a level 1 subject which is part of a symbol (SQ13)

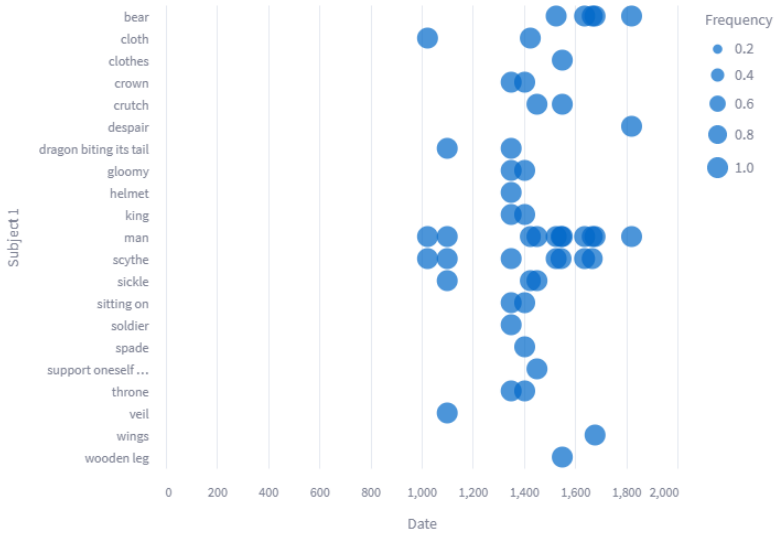

Figure 7.14: Level 1 subject variation over time of the iconography of Saturn (SQ14)

Whereas in SQ12 we retrieved the cluster of subjects, in this case, we do not group them according to the co-occurrence in the same artwork to better visualize the frequency with which level 1 subjects appear over time.

We filtered the results selecting only the iconographies appearing in more than one artwork, obtaining a total of 86 iconographies. Figure 7.14 shows how the evolution of an iconography over time can be examined through the example of the iconography of Saturn.[13] The attributes' frequency over the same year is registered through the dot size.

Similarly to SQ12, this sub-question provides a meaningful access point for historians interested in exploring the evolution over time of how an iconography is represented.

7.2.3.4 SQ15. Which attributes allow us to identify representations? What is their frequency?

With this query, shown in Listing 7.18, all the iconographies having at least one level 1 subject marked as recognizing are retrieved.

To analyze the frequency, we filtered the results by selecting only those iconographies that appeared more than once. We obtained 17 iconographies, out of the total of 155 iconographies having a level 1 description (see table 7.6). For each of them, we extracted the attribute and its frequency over the total frequency of the appearance of the iconography, namely how many times the subject is manifested

[13] On the online dashboard, the ID of the artworks in which the iconography appears is displayed. It is available at https://iconology-dataset.streamlit.app/Analysis_of_iconographies > Variation of iconographies > Over time.

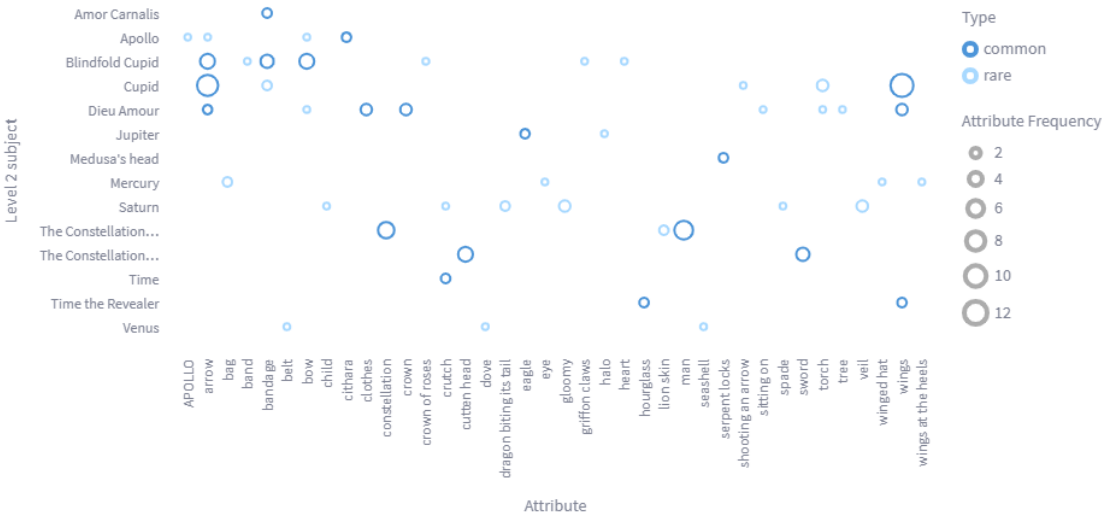

Figure 7.15: Most rare and frequent recognizing attributes for each iconography (SQ16)

in an Image. Table 7.7 shows the example of Cupid extracted from the results. The character appears 13 times with an attribute marked as recognizing. As expected, the wings are marked in the 92% of the images, followed by arrow and bow. The act of shooting an arrow is recognized only in one case.

This type of result can be analyzed for all the iconographies retrieved.

Level 2 subject	**Recognizing attribute**	**Frequency**	**Total of Images**
Cupid	wings	12	13
Cupid	arrow	10	13
Cupid	bow	7	13
Cupid	torch	3	13
Cupid	bandage	2	13
Cupid	shooting an arrow	1	13

Table 7.7: Frequency of recognizing attributes of Cupid (SQ15)

7.2.3.5 SQ16. Among the attributes marked as recognizing, which are the most common and rare?

To answer this question, we calculate the results of SQ15. To define the rare and common attributes, we calculated the percentage of appearance of each attribute on the total of iconography occurrences. Then, we define rare attributes as those that appear less than 35% of the time and common ones that appear more than 65% of the time. These values allow us to include iconographies that occur only three times.

Figure 7.15 shows the results. It is possible to note that 5 iconographies (i.e., Amor Carnalis, Medusas' head, the constellation Perseus, the personification of

Time, and Time the revealer) present only common attributes, showing high scores. Indeed, 6 of 7 attributes appear 100% of the time. The lower score (80%) is shown by the sword of Constellation Perseus. The low frequency of the iconographies, occurring only twice for 4 of them, contributes to the regular registration of common recognizing attributes. However, the results highlight that, despite the low occurrence, there are constant attributes that allow the recognition of these characters.

In contrast, 3 iconographies present only attributes marked as rare, namely Saturn, Mercury, and Venus, with a frequency score of 12, 6, and 3. Whereas Mercury and Venus present only three attributes, Saturn shows six attributes marked as rare, namely a child (1/12), a crutch (1/12), the dragoon biting its tail (3/12), a gloomy attitude (3/12), a spade (1/12), and a veil (3/12).

The great variety of Saturn's attributes marked as recognizing may be a hint of the great variations that this iconography had, reflecting the complexity of the figure through the centuries as highlighted by Panofsky (1972). In fact, if we compare these results with the scores of an iconography with a similar frequency, namely Cupid, this assumption is supported. Despite the similar frequency, Cupid presents two core attributes with a high frequency, namely the wings (12/13) and arrows (10/13). In addition to that, some rare attributes appear, such as a bandage (2/13), the act of shooting an arrow (1/13), and a torch (3/13).

In conclusion, **the visualization of the most rare and common attributes marked as recognizing may give a first hint on which iconographies are more frequently represented in contrast to other ones that present a high degree of variation**. These quantitative results may consequently be an exploratory analysis constituting the premises of further qualitative iconographical studies.

7.2.3.6 SQ17. Which are the most common level 1 subjects not marked as recognizing?

As introduced, level 1 subjects not marked as recognizing may be of interest to art historians willing to have a deep understanding of artworks. This sub-question, performed through the SPARQL query in Listing 7.19 aims to retrieve and highlight secondary variations of the iconographies.

As a result, we obtained 135 iconographies in which none of the level 1 subjects was marked as recognizing. For each iconography, we visually represented the frequency of level 1 subjects occurring in the Images in which there is no recognizing attribute. Figure 7.16(a) shows the results of Saturn iconography as an example.

Furthermore, to better contextualize the non-recognizing attributes frequency,

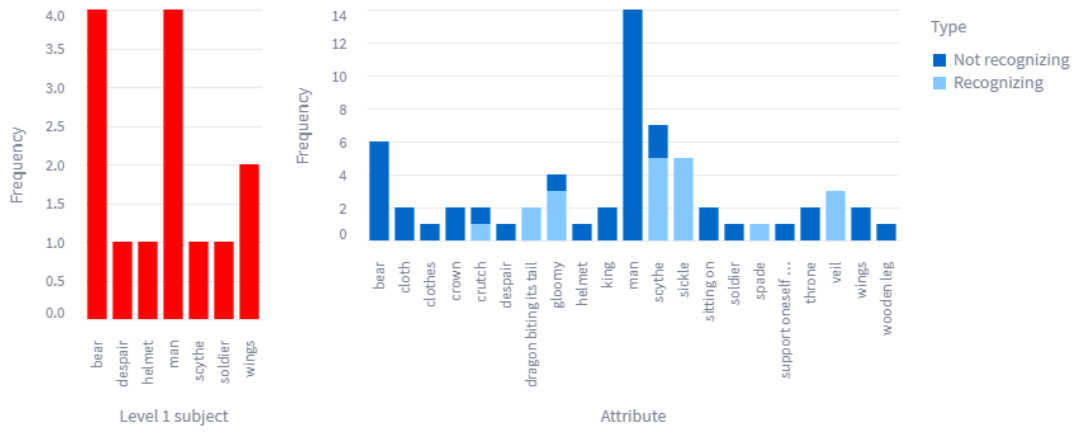

(a) (b)

Figure 7.16: a) Frequency of attributes not marked as recognizing of the iconography of Saturn; b) total frequency of recognizing and non-recognizing attributes of the same iconography

we repeated the analysis on all the iconographies having both level 1 and 2 descriptions, and we grouped the attributes according to the role of recognizing or not recognizing. The example of Saturn is shown in figure 7.16(b). On the one hand, this kind of visualization helps in **understanding how many times the same object constitutes a hint for the subject identification**. For example, whereas the sickle is always underlined as a specific attribute of Saturn, the gloomy attitude, the scythe, and the crutch present at least one occurrence in which it does not have this role. On the other hand, **it helps in understanding which are the recurring iconography features beyond the recognizing attributes**. Hence, in the provided example, Saturn is often represented as a man with a bear.

7.2.3.7 Summary

In summary, in this section, we analyzed various aspects of iconographies with a quantitative approach, such as the variation of iconographies, their change over time, the role of attributes having a symbolical meaning, and the analysis of recognizing attributes and of level 1 subjects not marked as recognizing. As the variation of iconographies requires having a close insight into each iconography, the results of this section particularly underline the role that quantitative analyses can have as a supportive tool for thorough qualitative analyses. This can be made through visualizations of single iconographies according to the criteria explored above.

7.2.4 Evidence

Question 10, which this analysis expresses, refers to those iconological analyses in which a peculiar iconography is interpreted thanks to its strong dependence on a coeval textual description. This fact gives also insights into the narrow relations of the artwork, or extensively, of the artist, with a specific cultural area.

Consequently, this section focuses on the analysis of the sources cited as evidence and tries to provide data about the sources that the artists may have known, according to the available art historians' interpretations. This information may indeed support other historians seeking to interpret other artworks by the same author, by highlighting the sources that the artist may have known, or, at least, seeing which later sources were used by other scholars to interpret the artist's works. Similarly, as the sources known during the same period of time may be of interest for contemporary artwork interpretation, we retrieve also the artworks' date and group the textual sources according to them.

In the following, we try to retrieve the interpretations in which the availability of a textual source led to a deeper interpretation of the artwork.

7.2.4.1 SQ18. What were the known textual sources to which the artwork refers? Is the artwork involved in a cultural phenomenon?

This sub-question fully expresses Q10. To answer it, we first retrieve all the artworks citing a piece of textual evidence and the artists, by performing the query shown in Listing 7.20.

We visualize the results by relating the textual sources to the authors, as shown in figure 7.17[14]. In this way, we obtained 37 artworks out of the 428 described ones (8,6%), 32 cited texts, and 21 artists.

Among the artists, the one having the highest number of cited sources is Piero di Cosimo, to whom 8 texts are related. Excluding the sole historiographical text (i.e., Vasari, *Vite*) and the Medieval one (Boccaccio), the others belong to the classical culture, including text by Vitruvius, Virgil, Servius, Ovid, Homer, and Lucretius. This result shows that the author of the interpretations (in this case, Panofsky) made abundant use of classical texts to interpret Piero's artworks. **This seems to suggest that, according to him, the artist knew a good amount of classical texts**. Besides the numbers seem to suggest that, **it's important to underline that these sources are narrowly dependent on the art historian making the interpretation**. Therefore, more historical studies are required to support the claim that the artist concretely knew and made reference to such

[14]the graphs of this sub-question are available in an interactive format on the online dashboard at https://iconology-dataset.streamlit.app/Analysis_of_evidence > Textual sources.

sources in his own work.

The second artist for the number of cited sources is Sandro Botticelli, for whom 6 texts are cited. Also in this case, the majority of texts belong to the classical culture, made an exception for the contemporary *Stanze* by Poliziano, cited as a source for two artworks. The same trend can be seen also in Lorenzo Lotto, who is associated with 4 classical sources and only a contemporary one. Another artist with several sources is Michelangelo. Contrarily to the previous artists, his sources are medieval (Petrarch, Landino), contemporary to the artist (Vasari), or of a later date (Hyeronimus Tetius). It can be assumed that texts belonging to a later date than the artist's life are used by the art historian as a historiographical source. As Vasari's *Vite* is a historiography of art, describing both earlier and contemporary artists, it cannot be considered an artist's primary source even when it is related to a contemporary artist, as is in this case. For its nature, it is likely that it is used as a bibliographical source by the art historian making the claim, as it provides more descriptions of the artwork under examination.

Titian presents four sources as well, of which two belong to the classical culture, one (Ripa, *Iconologia*) has a slightly later date, and the remaining one has an immediately earlier date (unspecified writings by Marsilio Ficino).

We repeated the analysis by retrieving the artwork's date of creation, if available. As a result, we obtained 43 artworks out of the 428 described (10%) and 39 texts (see figure 7.18).

The first clear result is that **the major part of artworks having a date and citing a textual source belongs to the XV and XVI centuries**. If we compare these results with the most represented periods shown in figure 7.1, we see that, despite being a well-represented period in the data, only one of the artworks of the XIII Century, to which a consistent number of artworks belong, cite a textual source, namely the *Commentaries on Ovid*. Similarly, only 6 artworks of the XIV Century have a reference to a source, which is mainly a contemporary one. This can raise several different hypotheses. The lack of cited sources may be due to Panofsky's method or his different levels of knowledge of the textual sources circulating during the centuries under consideration. Alternatively, the results can highlight the tendency of the Renaissance to a better and more scientific knowledge and interest in the classics which emerges in the visual arts in the form of precise references to the classical culture. Another hypothesis is that the tendency of the Renaissance period to refer to the classical culture is the assumption that led the art historian to search for the source of some iconographies in classical text.

This result can consequently support further studies over the causes of an apparent more abundant practice of using classical texts as a

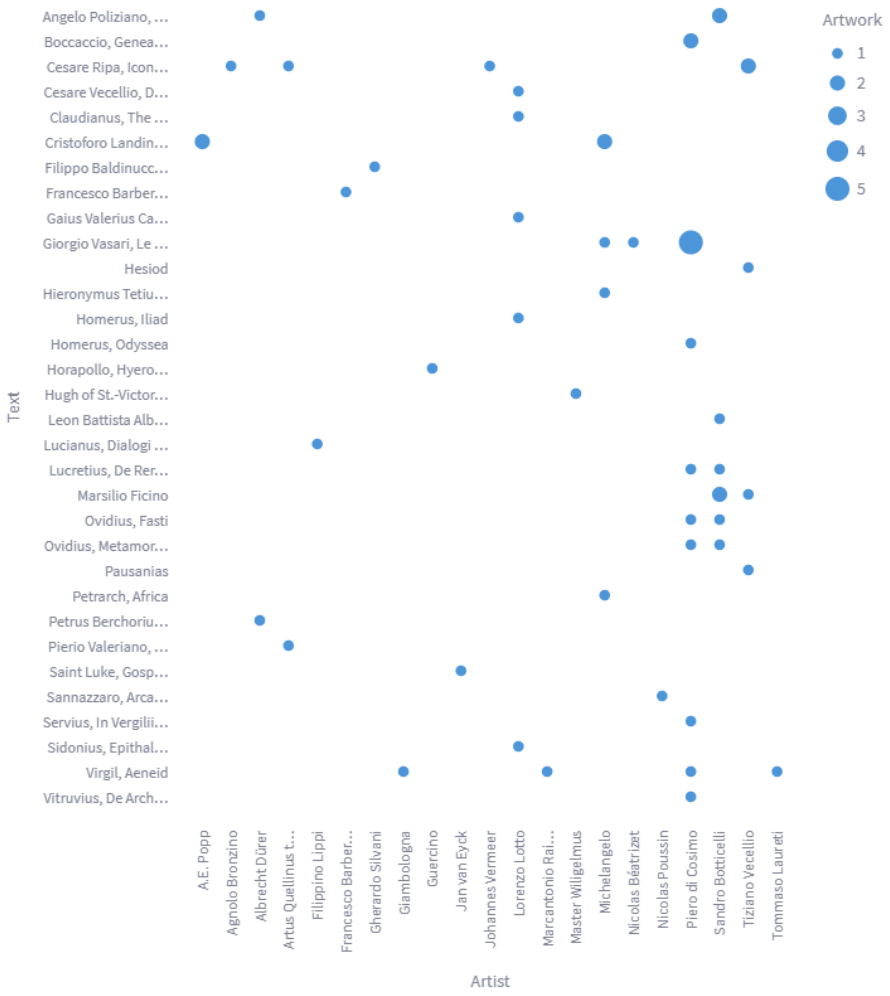

Figure 7.17: Textual sources grouped according to the artwork's author

source of the interpretation of XV-XVI Centuries art.

Then, we retrieved the artworks having a textual source and a third-level meaning associated, obtaining a total of 63 artworks having these characteristics. For each textual source, we show to which deeper meaning it relates, in order to highlight possible relations.

7.2.4.2 Summary

As the presence of pieces of evidence in relation to each type was treated in the previous section, in this section we focused on their type. Of particular interest is the retrieval of the textual sources that the art historians relate to the artwork, either as a later bibliographical source (e.g., Vasari's *Vite*) or as a source that, according to the interpreter, the artist may have directly known.

Gathering this information according to the period can provide a useful overview for art historians of the sources that, according to previous scholars, are needed to interpret an artwork of a certain period. Although we can not provide a distinction

Figure 7.18: Textual sources grouped according to the artwork's date

from data of the role that the source had in the scholar's interpretation, the source's overview can provide hints for making this distinction and helping the reconstruction of the artists' background, furnishing bibliographical hints of the presumably known texts. Furthermore, the analysis provided linking the textual sources to a particular date shows the tendency of recurring in a significant way to textual sources while interpreting Renaissance and Late Renaissance art, compared to other periods that are well-represented in the dataset, such as the Medieval one. This result suggests that such quantitative analysis can lead to further qualitative research either about visual citation practices or the art historians' assumptions and practices when interpreting art of a certain period.

7.2.5 Visual Citation

As introduced in Section 5.1.2, detecting recurrent visual patterns among diverse artworks, sometimes belonging to different cultures and historical periods, is one of the key types of research of iconographical and iconological studies. Furthermore, this citation can be a hint for deeper cultural phenomena taking place, as shown by results of SQ5.

To this end, we perform two SPARQL queries. Whereas the first one aims at retrieving all the artworks involved in a visual citation, the second one aims at detecting the artworks in which such citation is the evidence of a cultural phenomenon, by retrieving all the artworks having both a visual citation and a cultural phenomenon.

7.2.5.1 SQ19. What artworks cite the visual pattern of other ones?

This sub-question aims at retrieving all the artworks involved in a visual citation, with the role of citing or cited ones. Hence, the query presented in Listing 7.21 requests the citing (?art, ?artLabel) and cited (?art2, ?art2Label) artworks.

In total, 50 artworks cite another one. Table 7.8 shows an example of the results, presenting five artworks by Albrecht Dürer citing patterns from 7 ones. Among them, are other artworks by the same artist. **Chains of citations can be seen**, as two artworks visually cited, one of which (ART1516, ART1517), in turn, cites another artwork. The interpretation provided by Panofsky (1955) according to which Dürer knew Classical art during his trip to Italy, and this contact occurred especially with the mediation of contemporary Renaissance Italian artists, emerges from the presence of copied artwork both from classical art and from the Renaissance painter Mantegna.

In order to give a clearer view of the visual pattern cited, we provided an interface for comparing the two artworks, if a picture is available, as shown in

figure 7.19.[15] The example shows the visual citation by Dürer of an artwork by the Italian painter Pollaioulo of the pose of Hercules portrayed in the act of shooting an arrow. Furthermore, we compare the level 1 subjects involved in the visual citation for each artwork and the second-level subject, if any, to which they are linked.

Artwork ID	**Artwork title**	**Copied artwork ID**	**Copied Artwork Title**
ART1517	'Albrecht Dürer, Apollo-Sol and Diana (L.233), 1501-1504, London, British Museum'	ART1456	'Andrea Mantegna, Bacchanal with the Vat (B. 19), 1475'
ART1517	'Albrecht Dürer, Apollo-Sol and Diana (L.233), 1501-1504, London, British Museum'	ART1516	'Albrecht Dürer, Apollo Medicus or Aesculapius (L. 181), 1500, Berlin, Staatliche Museen'
ART1517	'Albrecht Dürer, Apollo-Sol and Diana (L.233), 1501-1504, London, British Museum'	ART1518	Apollo Belvedere
ART1517	'Albrecht Dürer, Apollo-Sol and Diana (L.233), 1501-1504, London, British Museum'	ART1455	Helios Pantokrator
ART1516	'Albrecht Dürer, Apollo Medicus or Aesculapius (L. 181), 1500, Berlin, Staatliche Museen'	ART1456	'Andrea Mantegna, Bacchanal with the Vat (B. 19), 1475'
ART1516	'Albrecht Dürer, Apollo Medicus or Aesculapius (L. 181), 1500, Berlin, Staatliche Museen'	ART1518	Apollo Belvedere
ART1515	'Albrecht Dürer (workshop), Apollo and Dafne, 1502'	ART1448	'Albrecht Dürer, Rape of the Sabine Woman (L. 347), 1495, Bayonne, Musée Bonnat'
ART1514	'Albrecht Dürer, Hercules at the Crossroads (B. 73), 1498'	ART1448	'Albrecht Dürer, Rape of the Sabine Woman (L. 347), 1495, Bayonne, Musée Bonnat'
ART1457	'Albrecht Dürer, The Fall of Man (B. 1), 1504'	ART1517	'Albrecht Dürer, Apollo-Sol and Diana (L.233), 1501-1504, London, British Museum'

Table 7.8: Example of copying and copied artworks

7.2.5.2 SQ20. In which cases do those artworks which are involved in a visual citation have a cultural phenomenon?

In Warburg, Wittkower, and Panofsky's studies the visual citation is often interpreted as evidence of a specific attitude. Usually, the citation of another artwork witnesses the presence of a cultural phenomenon. To this end, we retrieve all the artworks visually citing other ones and having a cultural phenomenon, and we visualize them as a network to highlight shared phenomena, if any.

As a result, we obtained a network composed of 92 artworks and 110 phenomena,

[15]The interface is available at https://iconology-dataset.streamlit.app/Analysis_of_citations > Visual pattern citations.

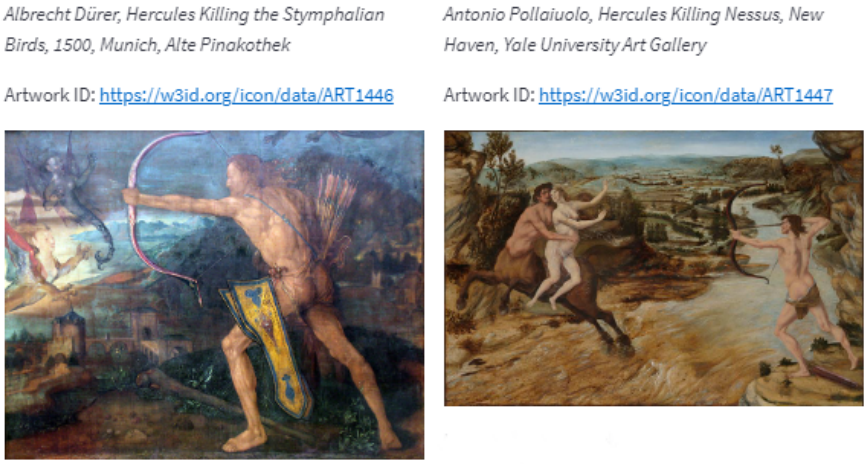

Figure 7.19: Interface for comparing the artworks involved in a visual pattern citation. In this example, Durer's *Hercules Killing the Stymphalian Birds* (on the left) visually cites Pollaiolo's *Hercules Killing Nessus* (on the right). All the artwork images are in the public domain

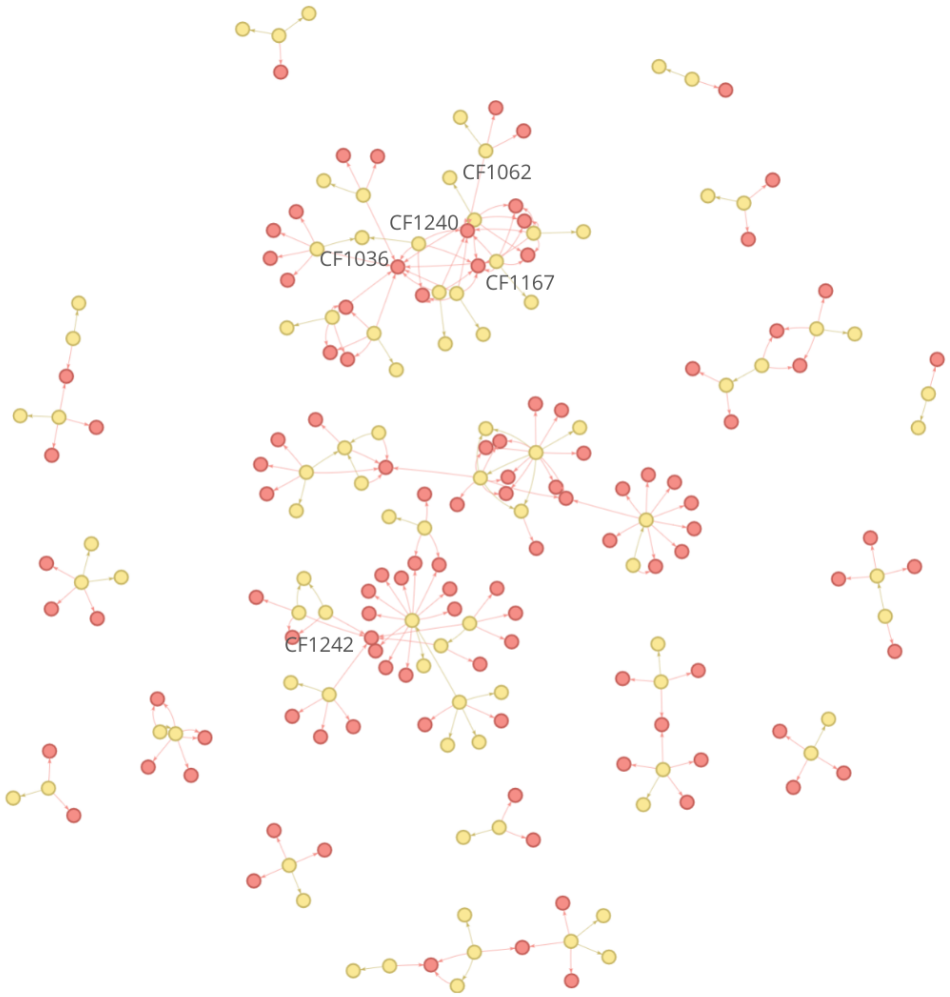

Figure 7.20: Network of artworks involved in a visual citation (yellow) and their cultural phenomena (red)

grouped into 19 clusters (Fig. 7.20)[16]. As we aim to detect phenomena shared among artworks, we examine phenomena having the highest number of related artworks and the type of art they cite. The most common phenomenon is *classical motif recovery* (CF1036), shared by 9 artworks, all from the Medieval period. On the contrary, all the artworks they visually cite, except one, belong to the classical one, including, for example, a bust of Antoninus Pius, ancient sarcophagi, a relief, and a vase. One cited artwork, namely Hercules, represents an affirmed iconographical type rather than a specific artwork. The second well-represented phenomenon is the one expressing the *Renascenses of classical antiquity before the Renaissance* (CF12040), counting 7 artworks, six of which are shared with the previous phenomenon, to which is added the citation of an entablature by an entablature of St. Gilles Church. Similarly, the phenomenon *Intrinsic classicism was created by the meeting of Gothic and proto-renaissance tendencies* (CF1167) is linked to 5 artworks, shared with the previous phenomenon. Hence, **the results obtained by the query, consisting of artworks visually citing mostly classical artworks, confirm that this phenomenon can be examined through this inquiry.**

In another cluster, the phenomenon *Representation of classical content with contemporary formal motifs* (CF1242) is related to 5 artworks as well, belonging either to the Medieval or Renaissance period. The cited artworks consist of two iconographical types, namely the Virgin lamenting over the body of Christ and a strigilated sarcophagus (Roman period).

The last phenomenon with a high number of related artworks is CF1062 (*Durer assimilated Classical art through the intermediate of Italian Renaissance art, comprehending it better than them*). It is linked to 5 artworks by Durer, which either cites other of his own artworks (which are drawings or references to Italian art), the Italian painter Pollaiolo, or the classical Apollo Belvedere.

7.2.5.3 Summary

This analysis focused on the artworks registered as part of a visual citation. SQ19 demonstrated that it is possible to retrieve them and build an interface on top of it to compare the artworks involved, both at a metadata level (including artwork details and level description) and at the visual one, if an image is available. Furthermore, the results of SQ20 show that, by retrieving both cultural phenomena and artworks involved in a visual citation and visualizing them as a network, it is possible to easily detect the groups of artworks having similar characteristics and sharing a common cultural phenomenon, becoming possible helpful support

[16]The interactive version of the network is available at https://iconology-dataset.streamlit.app/Analysis_of_citations > Visual pattern citations > network

for qualitative analyses of such phenomena.

7.3 New inquiries into the ICON dataset

In the previous section, we quantitatively addressed various key aspects of the research conducted in iconological studies. Although this is a relevant aspect, the current study aims to go further in the research to seek advantages and new research possibilities that could be opened thanks to the use of semantic data and quantitative analyses, which would be challenging to conduct with traditional methods.

Therefore, we focus on two key aspects. As the Iconology dataset reflects art historians' (specifically Erwin Panofsky's) interpretation, we will verify if his own method can be quantitatively characterized (RQ4.1a). Next, as the iconological approach does not present a unique theoretical framework shared by historians, we will inquire if patterns supporting the recognition of a phenomenon emerge from data (RQ4.1b).

7.3.1 Quantitative art historiography: characterizing Panofsky's approach

RQ4.1a asks whether it is possible to quantitatively characterize Panofsky's approach from data that express his own interpretations.[17] The following sections will illustrate the results obtained by focusing on specific aspects of the investigation, namely the role of textual sources (a1), the description of the artwork in levels of interpretation (a2), and the complexity of iconological claims (a3).

7.3.1.1 The usage of textual sources

RQ4.1a1 concerns the role of textual sources in the interpretation process. In contrast to what was expected, only 27% of the overall interpretations are supported by evidence (Fig. 7.21(a)). Considering recognitions that are part of interpretations (Fig. 7.21(b)), those recognitions citing more often a piece of evidence are the iconological recognitions (third level). Although the evidence consisting of a textual source (Information Object, Fig. 7.21(c)) is potentially relevant at all levels, it is less cited when proposing an iconological recognition. On the contrary, the majority of recognitions cite other artworks as supportive proof. Figure 7.21(d) shows that artworks are the most frequent type of evidence (72 items), compared to texts (41 items).

[17] Section 7.3.1 is based on the version published in Baroncini, Daquino, et al. (2023a)

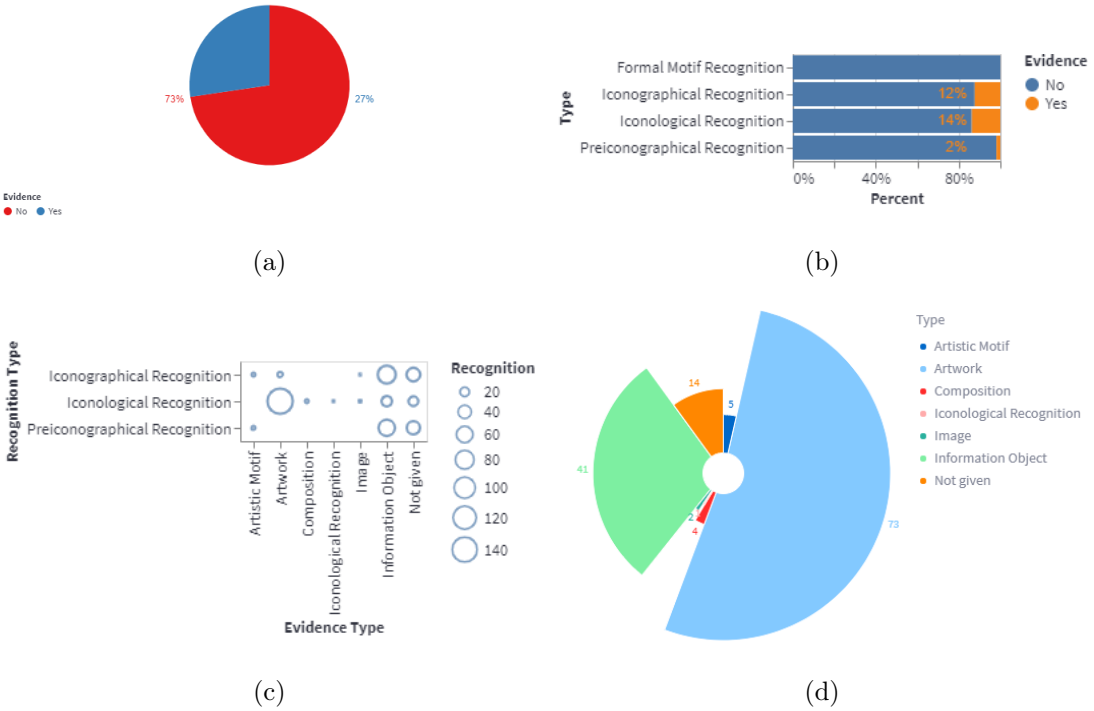

Figure 7.21: a) Percentage of interpretations citing evidence; b) percentages of recognitions citing evidence; c) correlation between recognition and evidence type; d) frequency of piece of evidence per type.

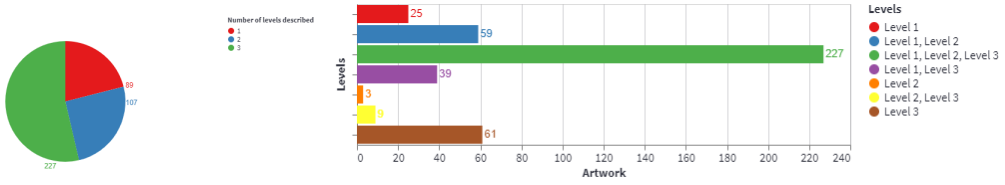

Figure 7.22: a) Number of artworks having a description of one, two, or three levels; b) levels described in each artwork.

7.3.1.2 Levels of description

RQ4.1a2 concerns the extent to which Panofsky adopts, in his interpretations, his own three-level interpretation theory. As he is the most represented person responsible for recognitions in the dataset, the analysis was conducted on all the data. Surprisingly, only nearly half of the artworks (53%) have at least one recognition at all three levels of interpretation (Fig. 4.a). Among the artworks having only one level described, the majority have only the third level (68%, fig. 4.b), whereas, among the ones having two levels described, the most common are artworks having a description at levels 1 and 2 (13%).

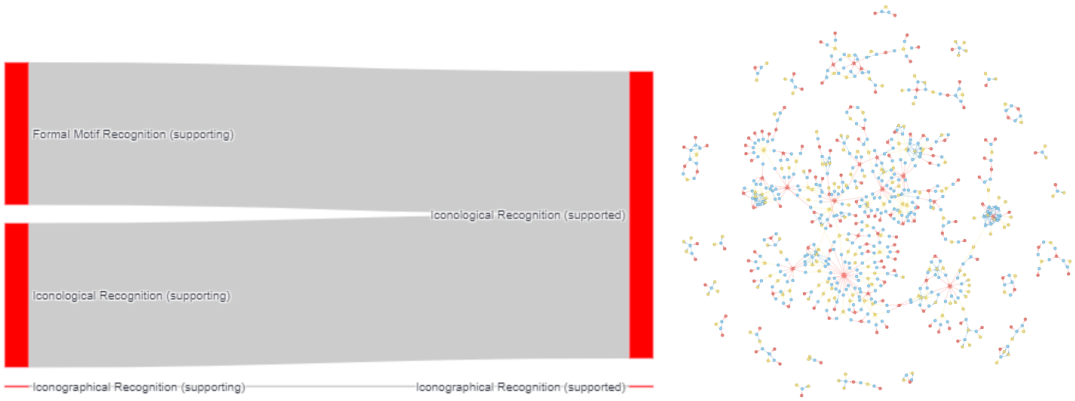

(a) (b)

Figure 7.23: a) Recognitions giving support to other ones divided per type; b) the network of iconological recognitions.

7.3.1.3 Iconological interpretations

The last sub-question (RQ4.1a3) concerns the extent to which the adopted modelling, based on Panofsky's theory, can represent his own sophisticated network of iconological interpretations. Figure 7.23(a)[18] shows how the distribution of supporting citations occurs according to the type of recognition. Formal Motifs Recognitions (level 1) and Iconological Recognitions equally support Iconological Recognitions. The network of cultural phenomena (Fig. 7.23(b)), including supporting citations and artworks cited as evidence, is highly interconnected, with an average of 1,29 edges per node. Nearly half of the artworks (200 out of 423) are included in the network, and only 10 clusters have a sole cultural phenomenon associated. On the contrary, a consistent part of the nodes is interlinked. In the bigger cluster, several important phenomena described by the art historian are connected, which are mainly devoted to the representation of classical content during the Middle Ages and Renaissance (see Table 3). Despite the artworks in this cluster being described by the art historian in different books, a connection between them is present in the data. Therefore, thanks to network visualisation, we can establish links between interpretations that are not explicitly stated in the source text. Nevertheless, some thematically related phenomena (e.g., CF1233, "reintegration of the classical astrological type with scientific and mythological antiquity"), that we would expect to be connected to the cultural phenomena cited above, do not present any connection to the cluster. In addition, although the network is highly interconnected, it is difficult to reverse it to the interpretation process made by Panofsky in the text, i.e., retrieving the exact corpus of artworks

[18]The interactive version of the graph is available at https://iconology-dataset.streamlit.app/Analysis_of_interpretations > The iconological interpretation

that he considered while making an interpretation (e.g., all the artworks supporting the interpretation of the iconography of Father Time and its evolution).

Table 7.9: The most connected cultural phenomena in the iconological network

Phenomenon	**Number of incoming edges**
Representation of classical content with contemporary formal motifs (CF1242)	26
Iconographic evolution of classical subjects (CF1135)	13
Reintegration of classical style and content in the Italian Renaissance (CF1231)	13
Iconographic evolution of the representation of the three-headed companion of Serapis (CF1144)	12

7.3.2 Inquiries into the iconological interpretation

As results from analyses performed with regard to cultural phenomena, it is challenging to provide a methodological definition of the causes leading the art historian to perform an iconological recognition. By examining correlations, we did not obtain relevant results between phenomena types and the objects at other levels, making exceptions for the Natural Element type, which correlates with different objects at all levels (SQ4). Furthermore, by analyzing the citation of textual evidence, the third level is not well supported (Section 7.3.1.1). Moreover, third-level recognitions are those registering the highest number of artworks described only at one level (Section 7.3.1.2).

Consequently, it is difficult to establish which are the recurrent characteristics proper of an iconological interpretation. The only assumption that seems supported by data is that the iconological interpretation is narrowly bound to the recognition of phenomena in other artworks or the recognition of details in other artworks and the artwork itself. This is supported by the analysis of the recognitions citing as support other ones, which highlight the fact that nearly the total of supported recognitions is of third-level type, as shown by the analysis in 7.3.1.3. Furthermore, 16% of the recognitions of a phenomenon explicitly refer to a specific portion of the artwork (see results of SQ5). Another analysis showing promising results is the network including iconological recognitions citing as support other ones, artworks cited as evidence, and the common phenomena that they may share (see Section 7.3.1.3).

From these results, we assume that the relations among artworks and phenomena have an important role in the understanding of artworks' deeper meanings. However, the question of what elements are needed to identify a cultural phenomenon, in Panofsky's practice, remains open, as only a few recognitions make use of sources,

do not apparently share the same types of subjects, and a good amount of artworks are described only at the third level.

That being so, in this analysis, we inquire into which characteristics emerge from the data that can give us hints on the practically applied iconological method (RQ4.1.2).

The first research question we address is *does the recognition of a deeper phenomenon need a thorough description at the previous levels? (RQ4.1b1)*. Our hypothesis is that, despite registering the highest number of artworks described at only one level, such artworks with a solely third-level description are involved in a cultural phenomenon that includes other artworks described at the other levels.

The second research question is *are the artworks discussed in different books treated with the same level of detail? (RQ b2)*. Our hypothesis is that the presence of a more thorough description may be a consequence of Panofsky's writing style.

To this end, we retrieve all the artworks that have at least a cultural phenomenon associated, the cultural phenomena, and a description at other levels, if present. We then group artworks in two groups, namely artworks described 1) only at the third level or 2) at multiple levels. Then, we calculate how many artworks of each group the cultural phenomena have.

A total of 374 phenomena and 327 artworks were retrieved. Most artworks (79%) are described at multiple levels, in contrast to those having a solely third-level description (see Table 7.11).

Table 7.10: The most connected cultural phenomena in the iconological network

Object	**Total**
Cultural phenomena	379
Artworks	327
Artworks described at multiple levels	268
Artworks described only at the third level	59

Table 7.11: Total of artworks having at least a cultural phenomenon, distinguished for the degree of description (multiple levels or only third one)

To have a better insight into the results, we divided the phenomena into five groups according to the degree of description of the artworks in which the phenomenon is recognized.[19] As shown in figure 7.24(a), the majority of phenomena is linked to artworks described at the sole third level (79%). The ones involving an artwork not described further are the rest 21%. To better understand how the remaining phenomena are distributed, we represent them in a separate pie chart, shown in figure 7.24(b). Half of such phenomena (40, corresponding to 50%)

[19] An interactive version of the analysis is available at https://iconology-dataset.streamlit.app/Analysis_of_interpretations > 'artworks per level' > 'artworks described with only cultural phenomena'.

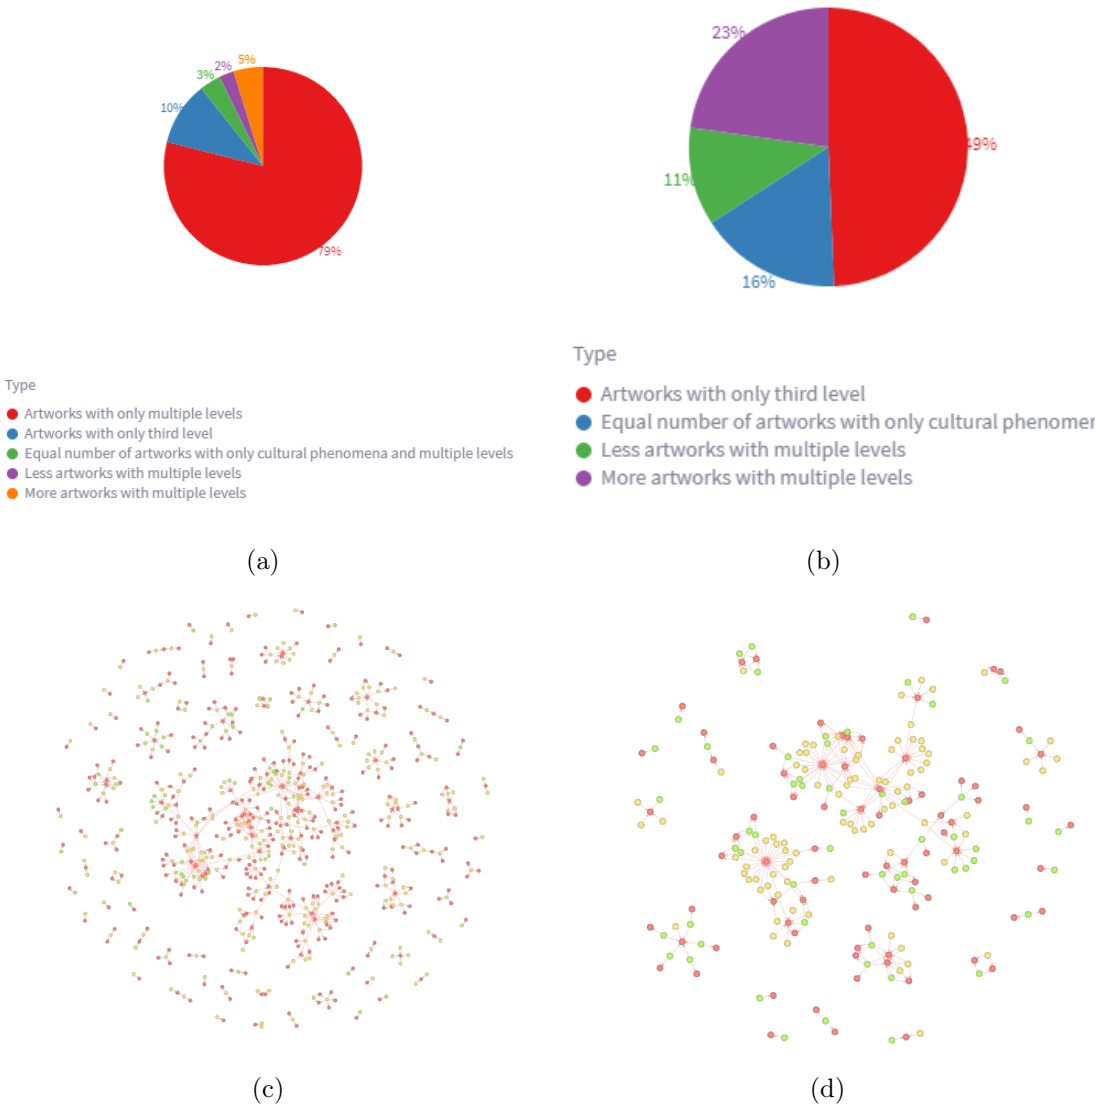

Figure 7.24: Overview of the level of description of artworks involved in a common cultural phenomenon

present only artworks described at the sole third level, for a total of 50 artworks. Most of the remaining ones (23%) are involved in cultural phenomena that present more thoroughly described artworks, followed by an equal number of artworks from the two groups (16%).

We further visualize the results in the network in order to see if such poorly described artworks are indirectly related to other phenomena (see figures 7.24(c) and 7.24(d)). From this network, it is possible to see that only 10 out of the 49 artworks are not included in a cluster in which there is at least one artwork described at multiple levels. Similarly, only 12 out of the 40 phenomena remain isolated, drastically reducing the number of phenomena related to artworks having only an iconological description. By examining the average of artworks per phenomenon, it results that the ones of the typology under examination (i.e., the ones having

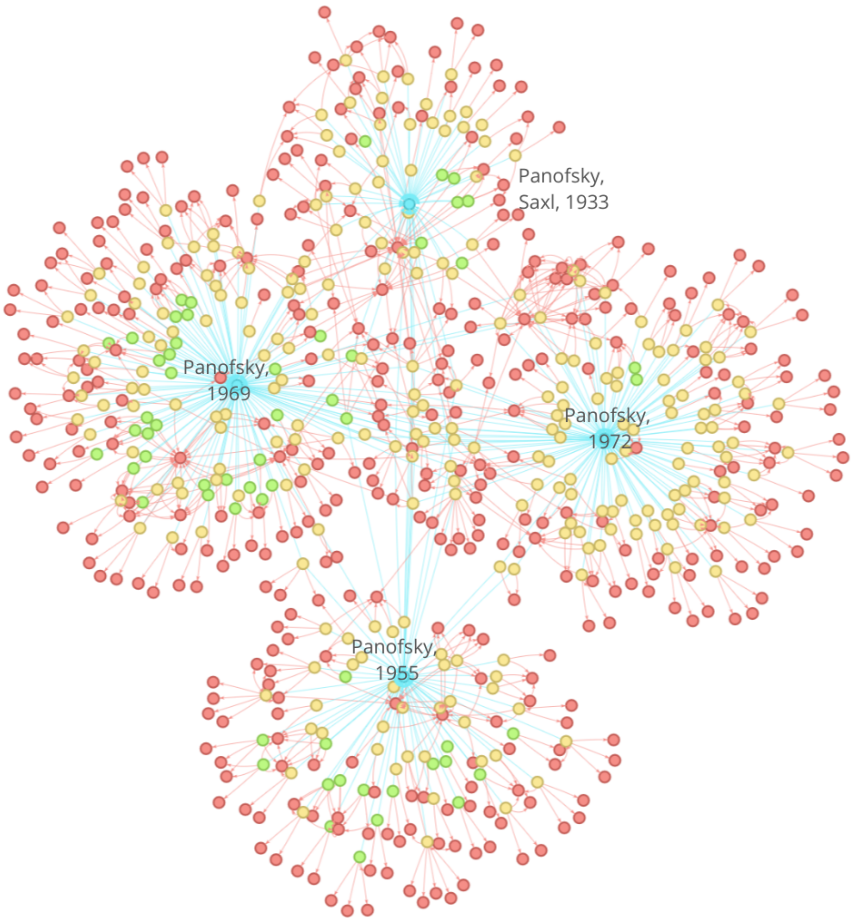

Figure 7.25: Distribution of artworks having only a cultural phenomena description (green) according to the bibliographical source (RQ4.1)

only third-level artworks) have an average of 1.2 artworks per phenomenon. This may reflect that such phenomena tend to be the ones more thoroughly describing the phenomenon specified. Hence, if we qualitatively look at some of the clusters having only one yellow artwork (i.e., described at multiple levels), we see that all the clusters refer to the same phenomenon subdivided into more specific ones. For example, a cluster of artworks and phenomena is thematically related to the wider phenomenon according to which the theory of proportions is a symbolic form of the respective culture, but it is split into more detailed phenomena.

Then, we examined if the distribution of only the third-level described artworks varies according to the bibliographic reference from which the description is extracted. The network shown in Figure 7.25 shows that the book *Renaissance and Renascenses in Western Art* (blue node in the bottom left) tends to present a more consistent quantity of such artworks (represented by green nodes), compared to the remaining three sources.

In conclusion, if we integrate the artworks having only a third-level phenomenon description into the network of common phenomena, the initial result of 58 artworks is reduced to 10 artworks remaining isolated, constituting only 3% of artworks having a cultural phenomenon. Despite we could not quantitatively characterize the reason motivating a deeper understanding of such a small group of artworks, it may be affirmed that the initial hypothesis according to which the phenomenon recognition is made on a similarity basis, is partially supported by data.

Furthermore, the tendency to provide a lighter description could be influenced by a stylistic choice of the art historian himself, as one out of the four bibliographical references contains a greater number of artworks having only cultural phenomena.

Listing 7.1: SPARQL query of SQ1

```
SELECT (count(distinct ?artwork) as ?tot) WHERE {

?iconol icon:aboutWorkOfArt ?artwork;
    icon:recognizedIntrinsicMeaning / icon:
        recognizedCulturalPhenomenon ?lev3.
}
```

Listing 7.2: SPARQL query of SQ2 (Level 1)

```
SELECT DISTINCT ?CFLabel ?lev1Label (count(?artwork) as ?tot) WHERE {

    ?artwork icon:iconologicallyRepresents ?meaning; icon:
        preiconographicallyDepicts ?lev1.
    ?meaning a icon:CulturalPhenomenon;        rdfs:label ?CFLabel.
    ?lev1 rdfs:label ?lev1Label.
} GROUP BY ?CFLabel ?lev1Label
ORDER BY desc(?tot)
```

Listing 7.3: SPARQL query of SQ2 (Level 2)

```
SELECT DISTINCT ?CFLabel ?lev2Label (count(?artwork) as ?tot) WHERE {

    ?artwork icon:iconologicallyRepresents ?meaning; icon:
        iconographicallyDepicts ?lev2.
    ?meaning a icon:CulturalPhenomenon;        rdfs:label ?CFLabel.
    ?lev2 rdfs:label ?lev2Label.
} GROUP BY ?CFLabel ?lev2Label
ORDER BY desc(?tot)
```

Listing 7.4: SPARQL query of SQ3

```
SELECT DISTINCT ?rec ?evid ?typeLabel WHERE {

        ?rec icon:aboutWorkOfArt ?artwork;
            cito:citesAsEvidence ?evid; icon:recognizedIntrinsicMeaning /
                icon:recognizedCulturalPhenomenon ?cf.
        OPTIONAL {?evid a ?type. ?type rdfs:label ?typeLabel.}
        FILTER NOT EXISTS {?evid2 crm:P106_is_composed_of ?evid}
    }
    }
```

Listing 7.5: SPARQL query to retrieve the total number of cultural phenomena recognitions citing an evidence (SQ3)

```
    SELECT DISTINCT (count(distinct ?rec) as ?tot) WHERE {

        ?rec icon:aboutWorkOfArt ?artwork;
            cito:citesAsEvidence ?evid;
            icon:recognizedIntrinsicMeaning / icon:
                recognizedCulturalPhenomenon ?cf.
    }
```

Listing 7.6: SPARQL query of SQ4

```
    SELECT DISTINCT ?artwork ?typeLabel ?CFBroaderLabel WHERE {
        ?artwork icon:preiconographicallyDepicts | icon:
            iconographicallyDepicts ?subj.
        ?subj a ?type.
        ?type rdfs:label ?typeLabel.
        ?artwork icon:iconologicallyRepresents ?cf.
        ?cf a icon:CulturalPhenomenon;
          crm:P2_has_type ?CFBroader.
        ?CFBroader rdfs:label ?CFBroaderLabel.
    }
```

Listing 7.7: SPARQL query of SQ5

```
    SELECT distinct ?rec ?typeLabel ?type2Label WHERE {
        ?rec icon:aboutWorkOfArt ?aboutArt;
            cito:givesSupportTo ?citedRec; a ?type.
        ?citedRec icon:recognizedIntrinsicMeaning / icon:
            recognizedCulturalPhenomenon ?cf2; a ?type2.
        ?type rdfs:label ?typeLabel.
        ?type2 rdfs:label ?type2Label.
    }
```

Listing 7.8: SPARQL query of SQ5

```
SELECT DISTINCT ?rec ?typeLabel ?rel ?obj WHERE {
    VALUES ?rel {icon:hasImage icon:hasComposition icon:
        hasArtisticMotif}
    ?rec icon:recognizedIntrinsicMeaning ?meaning; a ?type.
    ?meaning  icon:recognizedCulturalPhenomenon ?cf;
        ?rel ?obj.
    ?type rdfs:label ?typeLabel.
}
```

Listing 7.9: SPARQL query of SQ6

```
SELECT DISTINCT ?phenomenon ?artwork ?styleLabel ?artwork2 ?
    style2Label WHERE {
    ?style crm:P2_has_type <http://vocab.getty.edu/aat/300015646>.
    ?artwork a icon:Artwork;
        crm:P2_has_type ?style;
        icon:iconologicallyRepresents ?phenomenon.
    ?style rdfs:label ?styleLabel.
    ?phenomenon a icon:CulturalPhenomenon.

    ?style2 crm:P2_has_type <http://vocab.getty.edu/aat/300015646>.
    ?artwork2 a icon:Artwork;
        crm:P2_has_type ?style2;
        icon:iconologicallyRepresents ?phenomenon.
    ?style2 rdfs:label ?style2Label.
    ?phenomenon rdfs:label ?cfLabel.

    FILTER (?artwork != ?artwork2)
    FILTER (?style != ?style2)
}
```

Listing 7.10: SPARQL query of SQ7 and 8

```
SELECT DISTINCT ?simulacrumL ?realityL WHERE {
    VALUES ?rel {icon:hasCharacter icon:hasEvent icon:hasNamedObject
        icon:hasPlace icon:hasPersonification crm:P138_represents}

    ?img icon:hasSymbol ?simulation.
    ?simulation sim:hasSimulacrum ?simulacrum; sim:
        hasRealityCounterpart ?reality.
    ?simulacrum rdfs:label ?simulacrumL.
    ?reality rdfs:label ?realityL.
```

```
}
```

Listing 7.11: SPARQL query of SQ9

```
SELECT DISTINCT ?sim ?contextL ?periodLabel ?startDate ?endDate WHERE
    {
    ?rec icon:aboutWorkOfArt ?artwork; icon:recognizedImage / icon:
        hasSymbol ?simulation.
    ?simulation rdfs:label ?sim.


    ?obj crm:P65_shows_visual_item ?artwork;
        ^crm:P108_has_produced / crm:P4_has_time_span ?time .
    ?time crm:P82a_begin_of_the_begin ?startDate;
        crm:P82b_end_of_the_end ?endDate.

    OPTIONAL {?simulation sim:hasSimulationContext ?context. ?context
        rdfs:label ?contextL}
    OPTIONAL {?production crm:P108_has_produced ?obj; crm:
        P10_falls_within ?period. ?period rdfs:label ?periodLabel}
}
```

Listing 7.12: SPARQL query of SQ10

```
SELECT DISTINCT ?simulacrumL ?realityL ?evid WHERE {
    ?rec icon:aboutWorkOfArt ?artwork; icon:recognizedImage / icon:
        hasSymbol ?simulation.
    ?simulation sim:hasSimulacrum ?simulacrum; sim:
        hasRealityCounterpart ?reality.
    ?simulacrum rdfs:label ?simulacrumL.
    ?reality rdfs:label ?realityL.


    OPTIONAL {?simulation cito:citesAsEvidence ?evid.}
    OPTIONAL {?rec cito:citesAsEvidence ?evid.}
}
```

Listing 7.13: SPARQL query of SQ11

```
SELECT distinct ?rec WHERE {
?rec icon:aboutWorkOfArt ?aboutArt;
    icon:recognizedIntrinsicMeaning / icon:recognizedConceptualObject
        ?concept;
    cito:citesAsEvidence ?evid.
```

```
    ?evid a crm:E73_Information_Object.
    }
```

Listing 7.14: SPARQL query of SQ12, retrieving the level 2 subjects and the related Artistic Motif of Iconography

```
SELECT DISTINCT ?artwork ?img ?lev2Label ?amOrComp WHERE {
    VALUES ?rel {icon:hasCharacter icon:hasEvent icon:hasNamedObject
        icon:hasPlace icon:hasPersonification icon:hasSymbol crm:
        P138_represents}
    ?icrec icon:recognizedImage ?img; icon:refersToArtisticMotif ?
        amOrComp; icon:aboutWorkOfArt ?artwork.
    ?img a icon:Image;
        ?rel ?lev2.
    ?lev2 rdfs:label ?lev2Label.
}
```

Listing 7.15: SPARQL query of SQ12, retrieving the respective level 1 subjects

```
SELECT DISTINCT ?lev1Label WHERE {
    {?rec icon:recognizedArtisticMotif <'''+amOrComp+'''>.
    <'''+unique+'''> icon:hasFactualMeaning | icon:
        hasExpressionalMeaning | crm:P138_represents ?lev1.
    }

    UNION
    {?img icon:recognizedComposition <'''+amOrComp+'''>. <'''+amOrComp
        +'''> icon:hasPart ?am.
    ?am icon:hasFactualMeaning | icon:hasExpressionalMeaning | crm:
        P138_represents ?lev1. }
    ?lev1 rdfs:label ?lev1Label
}
```

Listing 7.16: SPARQL query of SQ13

```
SELECT DISTINCT ?artwork ?lev2Label ?lev1Label ?simLabel ?evidence ?
    evidenceLabel WHERE {
    VALUES ?rel {icon:hasCharacter icon:hasEvent icon:hasNamedObject
        icon:hasPlace icon:hasPersonification crm:P138_represents}
    ?icrec icon:recognizedImage ?img;
        icon:aboutWorkOfArt ?artwork;
        icon:refersToArtisticMotif ?comp.
    ?icrec2 icon:recognizedImage ?img2;
        icon:aboutWorkOfArt ?artwork;
```

```
        icon:refersToArtisticMotif ?am.
    ?img ?rel ?lev2.
    ?lev2 rdfs:label ?lev2Label.

    ?img2 icon:hasSymbol ?simulation.
    ?simulation rdfs:label ?simLabel.
    ?comp icon:hasPart ?am.
    ?am a icon:ArticsticMotif;
        icon:hasFactualMeaning | icon:hasExpressionalMeaning | crm:
            P138_represents ?lev1 .

    ?lev1 rdfs:label ?lev1Label.

    FILTER(?img != ?img2)
    FILTER(?icrec != ?icrec2)
}
```

Listing 7.17: SPARQL query of SQ14, retrieving the respective level 1 subjects

```
SELECT DISTINCT ?subj2Label ?subj1Label ?artwork ?startDate ?endDate
    WHERE {
    VALUES ?rel {icon:hasCharacter icon:hasEvent icon:hasNamedObject
        icon:hasPlace icon:hasPersonification icon:hasSymbol crm:
        P138_represents}
    ?icrec icon:recognizedImage ?img;
        icon:aboutWorkOfArt ?artwork;
        icon:refersToArtisticMotif ?l1.
    ?img ?rel ?subj2.
    ?subj2 rdfs:label ?subj2Label.

    {?l1 icon:hasPart ?am.
    ?am a icon:ArticsticMotif;
        icon:hasFactualMeaning | icon:hasExpressionalMeaning | crm:
            P138_represents ?subj1 }

    UNION

    {?l1 a icon:ArticsticMotif;
        icon:hasFactualMeaning | icon:hasExpressionalMeaning | crm:
            P138_represents ?subj1 }

    ?subj1 rdfs:label ?subj1Label.
```

```
    ?obj crm:P65_shows_visual_item ?artwork;
        ^crm:P108_has_produced / crm:P4_has_time_span ?time .
    ?time crm:P82a_begin_of_the_begin ?startDate;
        crm:P82b_end_of_the_end ?endDate.

}
```

Listing 7.18: SPARQL query of SQ15

```
SELECT DISTINCT ?img ?lev2Label ?lev1Label WHERE {
    VALUES ?rel {icon:hasCharacter icon:hasEvent icon:hasNamedObject
        icon:hasPlace icon:hasPersonification icon:hasSymbol crm:
        P138_represents}
    # ?icrec icon:aboutWorkOfArt ?artwork; icon:recognizedImage ?img.
    ?img a icon:Image;
        ?rel ?lev2.
    ?lev2 rdfs:label ?lev2Label.
    {?img icon:hasRecAttribute ?am} UNION {?img icon:hasRecAttribute ?
        comp. ?comp icon:hasPart ?am}
    ?am icon:hasFactualMeaning | icon:hasExpressionalMeaning | crm:
        P138_represents ?lev1.
    ?lev1 rdfs:label ?lev1Label

}
```

Listing 7.19: SPARQL query of SQ17

```
SELECT DISTINCT ?img ?lev2Label ?lev1Label WHERE {
    VALUES ?rel {icon:hasCharacter icon:hasEvent icon:hasNamedObject
        icon:hasPlace icon:hasPersonification icon:hasSymbol crm:
        P138_represents}
    ?icrec icon:recognizedImage ?img.
    ?img a icon:Image;
        ?rel ?lev2.
    ?lev2 rdfs:label ?lev2Label.
    {?icrec icon:refersToArtisticMotif ?am} UNION {?icrec icon:
        refersToArtisticMotif ?comp. ?comp icon:hasPart ?am}
    ?am icon:hasFactualMeaning | icon:hasExpressionalMeaning | crm:
        P138_represents ?lev1.
    ?lev1 rdfs:label ?lev1Label.
    FILTER NOT EXISTS {?img icon:hasRecAttribute ?am}
}
```

Listing 7.20: SPARQL query of SQ18

```
SELECT DISTINCT ?artwork ?artLabel ?artist ?artistLabel ?textLabel
    WHERE {
    {?rec icon:aboutWorkOfArt ?artwork;
        cito:citesAsEvidence ?evid.} UNION
    {?rec icon:aboutWorkOfArt ?artwork;
        icon:recognizedImage / icon:hasSymbol ?simulation.
    ?simulation cito:citesAsEvidence ?evid. }

    ?obj crm:P65_shows_visual_item ?artwork;
        rdfs:label ?artLabel;
        ^crm:P108_has_produced / crm:P14_carried_out_by ?artist .
    ?artist rdfs:label ?artistLabel.
    ?evid a crm:E73_Information_Object; rdfs:label ?textLabel.
}
```

Listing 7.21: SPARQL query of SQ19

```
SELECT DISTINCT ?art ?artLabel ?art2 ?art2Label where {

    ?rec icon:aboutWorkOfArt ?art;
        a icon:FormalMotifRecognition.
    ?rec icon:hasPrototypicalMotif ?am; icon:hasCopiedMotif ?am2.
    ?preicrec icon:aboutWorkOfArt ?art2; icon:recognizedArtisticMotif
        | icon:recognizedComposition ?am .
    ?obj crm:P65_shows_visual_item ?art; rdfs:label ?artLabel.
    ?obj2 crm:P65_shows_visual_item ?art2; rdfs:label ?art2Label.

    FILTER (?artLabel != '')
    FILTER (?art2Label != '')
}
```

Listing 7.22: SPARQL query of SQ19

```
SELECT DISTINCT ?art ?artLabel ?cfArtLabel ?fromArt ?fromArtLabel ?
    cfFromArtLabel where {

    ?rec icon:aboutWorkOfArt ?art;
        a icon:FormalMotifRecognition.
    ?rec icon:hasPrototypicalMotif ?am; icon:hasCopiedMotif ?am2.
    ?rec2 icon:aboutWorkOfArt ?fromArt; icon:recognizedArtisticMotif |
         icon:recognizedComposition ?am .
    ?obj crm:P65_shows_visual_item ?art; rdfs:label ?artLabel.
```

```
    ?fromObj crm:P65_shows_visual_item ?fromArt; rdfs:label ?
        fromArtLabel.

    # respective cf

    ?iconol2 icon:aboutWorkOfArt ?fromArt; icon:
        recognizedIntrinsicMeaning / icon:recognizedCulturalPhenomenon
         ?cfFromArt.
    ?cfFromArt rdfs:label ?cfFromArtLabel.

    ?iconol icon:aboutWorkOfArt ?art; icon:recognizedIntrinsicMeaning
        / icon:recognizedCulturalPhenomenon ?cfArt.
    ?cfArt rdfs:label ?cfArtLabel.

    FILTER (?artLabel != '')
    FILTER (?fromArtLabel != '')
}
```

CHAPTER 8

Evaluation

This chapter illustrates how the resources created in the doctoral project, i.e., the Iconology Dataset and the ICON Ontology, were evaluated.

The data set is evaluated with respect to its data quality, its ability to answer domain research questions, and its compliance with the FAIR principles. The ontology is tested through 1) competency questions performed on a subset of the Iconology dataset, 2) a comparison with other ontologies, and 3) an automatic evaluation of its syntax correctness, logical consistency, and compliance with FAIR principles.

8.1 Dataset evaluation

In the vision of the semantic web, to guarantee the correct interoperability of data, a crucial aspect is the quality assessment of published datasets. This led to consistent data validation and verification techniques available in the literature (Zaveri et al., 2016), as illustrated in Section 3.6.1. For the quality assessment of the Iconology Dataset, we performed two types of evaluation. First, we evaluated the dataset regarding the intrinsic, contextual, representational, and accessibility aspects by performing a selection and adaptation of the metrics proposed by (Färber et al., 2018). This ensures that the data set is acceptably complete for what concerns the content, error-free from a formal point of view, consistent, accessible, and interlinked. In particular, the evaluation of data consistency against the constraints of the used ontologies was a crucial point, for which we decided to perform a more thorough evaluation through SHACL shapes.

The second step of the evaluation involves assessing the dataset's adherence to the FAIR principles through online self-assessment surveys.

8.1.1 Metrics

For the quality assessment on the Iconology Dataset, we adopt the applicable metrics of the study by Färber et al. (2018), as, compared with the state of the art presented in Section 3.6.1, it provides a clear and complete methodology for the evaluation of data quality based on the available literature. They propose a total of 34 criteria, divided into 11 dimensions, grouped into four categories. 23 of the 34 criteria were adopted, as they were compatible with the dataset. All criteria score 0 to 1, where 1 corresponds to a high-quality grade concerning the selected metric. We sometimes specialized the metrics to suit the dataset-specific features better. All re-used metrics are listed in Table 8.1. For each, the adopted metric and the results are provided[1]. In the following sections, we better describe the metrics for each category and their eventual refinement, providing motivations for each score. The overall result of all the adopted metrics is 0.84, showing that the Iconology dataset has a good quality grade in several aspects.

8.1.1.1 Intrinsic category

Intrinsic data quality is the first category taken into account. It denotes the inherent quality of the data itself, namely its context-independent characteristics. The intrinsic category includes three dimensions: accuracy, trustworthiness, and consistency.

Accuracy aims at assessing the formal correctness of the data, considering the aspects of syntax and semantics. Indeed, it verifies that the data respect RDF syntactic rules and that the literals adhere to standards. In our case, we verified whether the dates, expressed as literals, comply with the ISO8601 standard. The semantic validity metric aims to assess whether the content expressed by literals is semantically sound. As the information provided in the Iconology dataset is extracted from Panofsky's bibliographic references, we adapted the metric to the specific case. In detail, we extracted the information concerning a sample of five randomly selected artworks, and we verified if the content of triples, including literals, holds true by comparing them with the information provided in the bibliographic reference. We focussed on 1) triples that include the artwork label to verify the correctness of artwork metadata and 2) triples describing the subjects recognized in the artworks to verify if the actual content was correct. All the metrics of this dimension performed well, with an average of 1.

The second dimension, namely **trustworthiness**, aims at assessing if the infor-

[1]A notebook with the analysis performed is available in the project repository

mation provided in the data set is correct, true, real, and credible. Färber et al. (2018) formalize this dimension into two aspects: trustworthiness on the Knowledge Graph level and the statement level. Although a further metric dedicated to the representation of unknown and empty values is provided, we do not apply it to the current dataset evaluation, as this type of data is not included in it. The first metric focuses on the adopted data curation and insertion methods. Specifically, it considers whether 1) data were manually inserted by experts, community contributors, or automatically extracted, and 2) whether the operation took place in an open or closed system. As the Iconology Dataset was manually created and curated in a closed system, the score 1 is assigned. The second metric is then used to assess the presence of the provenance assertion. As iconographic and iconological statements are documented through the ICON ontology relations, this metric gets 1 as a score.

The **consistency** dimension evaluates whether data in a dataset conflict with each other. As introduced in Chapter 4, due to this dimension's relevance in the literature, we dedicate a specific evaluation of it through SHACL to verify if instances are compliant with the logical constraints established at the ontology level. The SHACL shapes were extracted from the ontologies used (e.g., ICON, CIDOC-CRM, RDFS, DOLCE, CiTO, PROV-O, etc.) through the free online tool provided by Astrea[2](Cimmino et al., 2020). Among reused ontologies, FOAF, XML Schema, and DCAT-AP_IT were excluded from SHACL shapes generation, as their prefix was not regularly resolved by the tool, that is, no RDF was retrieved when the ontology prefix was passed[3].

The dataset was validated against the obtained definition through the tool ELI Validator[4], created in the context of the European Legislation Identifier project. Validation was performed iteratively on the data, which were consequently fixed according to the results.

In the last iteration, some errors remained unsolved. An error is that instances of Allegory and Story do not have the type of superclass `Invenzione` declared, a specification which, according to the author, would make data less readable. Second, the ICON ontology requires that the range of the property `icon:recognizedConcept` is a `dul:SocialObject`. As, in the context of the current modeling framework, we extensively reuse CIDOC-CRM for the description of fundamental aspects of the included objects (i.e., artworks, books, people), we decided to use the closest entity of CIDOC, namely crm:E28_Conceptual_Object, rather than the DOLCE's one.

[2] https://astrea.linkeddata.es/
[3] the generated SHACL shapes are available in the project's repository
[4] https://webgate.ec.europa.eu/eli-validator/validate

Another error is that the dataset does not declare the inverse property of relation `pro:hasRoleInTime`. To avoid lower performance in data retrieval, we included the inverse properties only of ICON, which is the ontology that describes the core aspects of iconography and iconology.

The overall score performed by the dataset for the intrinsic category is 1.

8.1.1.2 Contextual category

Contextual data quality underscores the need to evaluate data quality within the specific context in which data are used. According to Färber et al. (2018), it includes the dimensions of relevance, completeness, and timeliness. For the evaluation of the current dataset, we adopt the only dimension of completeness. As data mainly reflect the art interpretations expressed by a single art historian in a determined range of time, the dimensions of relevancy, which evaluates the presence of a ranking system of divergent assertions, and timeliness, which evaluates if the data are updated for a successful usage for the task at hand, are not applied in this evaluation.

Completeness expresses the degree to which the data is comprehensive enough to achieve the desired goal required by the specific context in which they are used. Färber et al. (2018, p. 8) identifies three metrics for this dimension, namely schema completeness, column completeness, and population completeness. As the metric of population completeness verifies whether the considered KG covers the basic population, we consider it not applicable to this case, in which a selection of Panofsky's claims was chosen for the dataset population as a case study.

Schema completeness determines whether the adopted schema can fully describe the domain of interest by comparing its classes and attributes to a golden standard. As the data set focuses on a detailed description of iconographical and iconological interpretations, we adopt as a gold standard the golden example schema formulated in Section 3.6.3.2 (see Figure 3.1), which is already used to evaluate the completeness of the schema of currently available KGs in relation to the *icon* statements. Similarly, we adopt the same criterion, consisting of 1) the alignment of the properties used in the KG to the gold schema and 2) the schema completeness as the division between the number of properties of the example that have been aligned and the total number of properties in the example. As all the properties were aligned, the score obtained is 1.

Column completeness determines whether the attributes of classes defined at the ontology level are present in the instance level of the Knowledge Graph. We specified this metric for the core classes of the dataset, namely the Human-made Objects (E22), Artworks, and the four types of Recognition. As the relations

describing them intend to provide means for a more thorough description, when necessary, we narrow the scope solely to the relations considered necessary or highly recommended. The final score is calculated by dividing the number of instances of each specific class that occur with the property considered by the total number of instances of that class. The dataset performed 0.85 as a score for this metric.

The overall score performed by the dataset for the intrinsic category is 0.83.

8.1.1.3 Representational category

Representational data quality encompasses elements tied to data format and its meaning. It comprises two key dimensions, viz. the ease of understanding by a human interpreter and interoperability, focusing on machine-readability.

Ease of Understanding indicates how easily data can be understood by a human data consumer without any confusion or ambiguity. To this end, the criteria of 1) description of resources, 2) labels in multiple languages, 3) provision of an understandable RDF serialization and 4) the presence of self-describing URIs are used to determine whether the data comply with this dimension (Färber et al., 2018, p. 10).

The first criterion evaluates whether the described resources present a textual description specified through the relations `rdfs:comment`, `rdfs:label` or `schema:description`, and it is calculated by dividing the number of resources that have such description by the total number of resources.

The second criterion aims at verifying if textual descriptions are provided in multiple languages, an aspect that would improve the ease of understanding of the dataset. As we did not provide labels in other languages, this metric performs 0.

According to Färber et al. (2018), some RDF serializations are more comprehensible to a human than the RDF/XML standard. Therefore, the third metric assesses whether data are provided in other serialization formats. Since the data set is also released in Turtle format, this metric gives the maximum score.

The last metric of the ease of understanding dimension concerns the presence of self-describing URIs. In fact, their use radically increases the understandability of data. Although extensive use of self-describing URIs was made in the creation of the Iconology Dataset (see Section 6.3.4.1), the characteristics and needs of some resources prevented their use in all cases. The unique identification of Cultural Objects and artwork was a fundamental aspect, which would not have been guaranteed by the use of self-describing URIs as artworks may have the same descriptive metadata (e.g., title, artist). Second, the length of the description provided for Cultural Phenomena prevented the creation of self-describing URIs based on their label, making necessary the creation of an alphanumeric string to

identify them. For these aspects, the result for this metric is 0.5.

The overall score obtained by this dimension is 0.58.

Interoperability focuses on evaluating the representation of data from a technical perspective. Therefore, it considers whether the data are concise and whether they can be easily integrated with other data. These aspects are evaluated through 1) the presence of blank nodes or reification classes, 2) the provisioning of several serialization formats, 3) the use of external vocabulary, and 4) the interoperability of proprietary vocabulary. As the internal vocabulary is not defined, we consider the last criterion to be not applicable in our case.

Blank nodes and reification classes are necessary when further information is provided about the assertion. However, they may prevent straightforward data operations, such as data integration and data processing on the client side. As both CIDOC-CRM and ICON extensively use reification to provide provenance information about the assertion, but the use of blank nodes is avoided, this metric scores 0.5.

Secondly, since we followed best practices for data publishing, we provided data in two formats, Turtle and RDF/XML (score: 1).

Since no classes and relations were internally defined in the dataset, but all the vocabulary used belongs to standards or newly-created ontologies that are externally defined and aligned with standards, all the terms used for resource description are taken from external vocabularies. Therefore, the external vocabulary usage metric performs 1 as a score.

The interoperability dimension scored an overall score of 0.83.

8.1.1.4 Accessibility category

The final element of the assessment is accessibility, which looks at how data can be obtained. It is divided into the dimensions of accessibility, licensing, and interlinking.

The dimension of **accessibility** measures the degree to which data can be accessed or obtained quickly and easily. In the following paragraphs, we elaborate more on the adopted criteria.

The dereferencing possibility of resources examines whether the URIs of RDF resources are accessible and return RDF information. To this end, we extracted a sample of 20 URIs, equally divided into internally created and external to the data set, and checked whether RDF information was returned. As a result, all the URIs of the sample were deferenceable. The availability of the KG verifies the Graph availability over time. Since it has just been released, we can not evaluate this

criterion by performing the metric proposed by Färber et al. (2018). Instead, we consider whether the dataset is hosted on a server that guarantees its availability over time. As the DH.arc University research center hosts the dataset, we consider this metric satisfied. The criterion *provisioning of public SPARQL endpoints* checks whether a public SPARQL endpoint is provided. Since we provide a SPARQL endpoint,[5] this metric scored 1. Similarly, since the data set is provided in two formats, the criterion *provisioning of an RDF export* also obtains the maximum score. The *support of content negotiation* metric verifies whether the endpoint supports different data formats during the content negotiation request and whether the server provides the requested content type. The ICON endpoint supports several data formats (e.g., JSON, CSV, XMLRDF, TURTLE), but it does not return results accordingly, so this metric scores 0.5. Finally, as structured metadata are provided,[6] the criterion *provisioning of metadata about a KG* obtained the maximum score. The overall accessibility score is 0.9.

The **licensing** dimension aims at assessing whether an appropriate license for reuse is provided. Since we published it under a Creative Common license, specified in the data through the relation `dct:license`, this metric obtained 1 as a score.

The **interlinking** dimension aims to assess whether entities expressing the same concept are linked to external sources to allow the exploration of further knowledge. It includes two criteria, viz. 1) the interlinking through `owl:sameAs` and 2) the validity of external URIs. The first checks the amount of resources linked to external ones through the specified property. To calculate it, we selected only actual resources (i.e., we filtered out the reification classes). We then calculated the total of resources that have at least an `owl:sameAs` property and divided it by the total of resources. As we performed a semi-automatic alignment and every entity was manually checked before being added to the dataset, we do not perform further metrics of precision and recall of the alignment algorithm. Although resource alignment is a central part of the development of the data set, few resources were aligned with this property (score: 0.13). The main reason for this is the lack of respective sources in other KGs because of the dataset's high granularity and domain specificity.

Linking to external resources is a pillar of the semantic web. However, such links can become obsolete and not be available anymore. Hence, the last criterion in the accessibility section verifies whether the external links included in the data set are active. To this end, we selected a sample of 25 URLs representative of

[5] As the SPARQL endpoint is hosted on a university server, it is considered stable. Since the dataset has small dimensions, no timeout issues have been experienced.

[6] https://w3id.org/icon/data

the aligned main source. In fact, the sample consisted of 10 URLs of pictures (`foaf:picture`), 5 from Iconclass, and 5 from Getty Vocabularies. This criterion performed well (score: 1).

Table 8.1: Performed evaluation metrics

Name	Metrics	Result
Intrinsic category		
Accuracy	Data accuracy is defined as the degree to which the data set is exact, reliable and free of errors (Färber et al., 2018, p. 4)	
Syntactic Validity of RDF Documents	Validation through TurtleValidator[7]	1
Syntactic Validity of Literals	Check if dates comply with the ISO8601 standard through regular expressions	1
Semantic Validity of Triples	Adapted metric: select triples of a sample of 5 artworks that have literals and verify if they are true by comparing them with the textual source from which the information was extracted or with the artwork picture	0.99
Accuracy average		**1**
Trustworthiness	The degree to which the information is accepted as correct, true, real, and credible. Trustworthy applications need to be able to verify the origin of data.	
Trustworthiness on KG level	Metrics: 1: manual data curation, manual data insertion in a closed system, 0.75: manual data curation and insertion, both by a community, 0.5: manual data curation, data insertion by user or automated knowledge extraction from structured data sources, 0.25: automated data curation, data insertion by automated knowledge extraction from structured data sources, 0: automated data curation, data insertion by automated knowledge extraction from unstructured data sources	1
Trustworthiness on statement level	Metrics: 1: provenance (through the use of properties) on statement level is used, 0.5: provenance on resource level is used, 0: otherwise	1
Trustworthiness average		**1**
Contextual category		
Completeness	The extent to which data are of sufficient breadth, depth, and scope for the task at hand	

[7] https://github.com/IDLabResearch/TurtleValidator

Schema Completeness	Definition of a gold standard for iconographical and iconological description from a golden example. The properties used in the KG are aligned to the gold standard. The schema completeness metric is obtained by dividing the number of properties of the example that have been aligned on the total number of properties in the example.	1
Column Completeness	Ratio of the number of instances that have both the class k and a value for the relation r, to the total number of instances that have class k. We apply this metric to the core classes, namely: Human-Made Objects, Artworks, Recognitions, and to their necessary or highly recommended relations	0.85
Completeness average		**0.92**
	Representational category	
Ease of Understanding	Dimension defining the grade of understandability of a data source by a human data consumer	
Description of resources	The metric is computed as the number of described resources divided per the number of all considered resources. We apply the computation only actual instances, excluding reification classes.	0.94
Labels in multiple languages	Presence of labels in other languages than English	0
Understandable RDF serialization	Data are provided in RDF in formats more human-understandable than RDF/XML (e.g., N3, N-Triple, and Turtle). **Metric:** Verify if the serialization formats provided include other formats that RDF/XML.	1
Self-describing URIs	Metrics: 1 self-describing URIs always used 0.5 self-describing URIs partly used 0 otherwise	0.5
Ease of understanding average		**0.61**
Interoperability	interpretability focuses on the representation from a technical perspective	
Avoiding blank nodes and RDF reification	Metrics: 1: no blank nodes and no reification; 0.5: either no blank nodes or no reification; 0: otherwise	0.5
Provisioning of several serialization formats	This criterion verifies the presence of multiple serialization formats	1

Using external vocabulary (ontologies)	Presence of external vocabulary (i.e., classes and relations) for resource description.	1
Interoperability average		**0.83**
	Accessibility category	
Accessibility	The extent to which data are available or easily and quickly retrievable	
The dereferencing possibility of resources	Over a sample of 20 URIs, we check if the URIs are accessible via HTTP requests and if they return useful RDF data by analyzing the HTTP status codes.	1
Availability of the KG	**Adapted metric:** whether access to the dataset is provided by a server that guarantees its continuity	1
Provision of public SPARQL endpoints	Whether an official SPARQL endpoint is publicly available.	1
Provisioning of an RDF export		1
Support of content negotiation	Metrics: 1 content negotiation is supported, correct content type is returned; 0.5: content negotiation is supported, but results are not in the desired content type; 0: otherwise	0.5
Provisioning of metadata about a KG	Presence of structured metadata on KG	1
Accessibility average		**0.92**
License	Provision of machine-readable licensing information	1
Interlinking	The degree to which entities representing the same concept are connected, within or between multiple data sources, is referred to as interlinking.	
Interlinking via owl:sameAs	Proportion of resources that have a connection to an external knowledge graph via `owl:sameAs`. We exclude reification classes from this metric	0.13
Validity of external URIs	We assess the accuracy of external URIs by analyzing a sample set of URIs to determine if there is a timeout, a client error (HTTP response 4xx), or a server error (HTTP response 5xx). It is tested over a random sample of 25 URIS, composed of 5 external URLs of pictures, 5 from Iconclass, 5 from Getty Vocabularies	1
Interlinking score		**0.55**

Overall score	**0.84**

8.2 Evaluation: domain-specific research questions

To evaluate to which extent traditional iconological issues can be quantitatively addressed through the modeling and data created, we assessed which of the 20 sub-questions presented in Chapter 7 (see Table 7.1) could be answered. All the sub-questions were translated into SPARQL queries and performed against our dataset in our analysis.

As shown in Table 8.2, about 85% of the queries could be answered, while 15% were not fully answered due to missing or limited data that hamper a reliable analysis of the phenomenon. For example, symbols rarely appear multiple times (i.e., the dataset includes very diverse subjects), hindering the observation of their evolution over time.

Table 8.2: Percentage of addressed and partially addressed sub-questions expressing the RQs of the domain

Theme	**Fully addressed**		**Partially addressed**	
	Questions	Percentage	Questions	Percentage
Cultural phenomena	SQ1, SQ2, SQ3, SQ5, SQ6	83,34%	SQ4	16,67%
Symbols	SQ7, SQ8, SQ10, SQ11	80%	SQ9	20%
Iconographies	SQ12, SQ13, SQ15, SQ16, SQ17	83,33%	SQ14	16,67%
Evidence	SQ18	100%		0
Visual citations	SQ19, SQ20	100%		0
Total		85%		15%

8.2.1 FAIR evaluation

As the FAIRness of the data is a crucial aspect of Linked Data (Wilkinson et al., 2016), we conducted a self-evaluation using the DANS self-assessment questionnaire[8]. The data set obtained an overall score of 76% for compliance with

[8]Available at https://satifyd.dans.knaw.nl/. A copy of the Iconology Dataset Evaluation results is available in the project repository

the FAIR principles, divided into Findable (77%), Accessible (75%), Interoperable (92%) and Reusable (61%).

8.3 Ontology evaluation

The ICON ontology was evaluated in i) its extraction potential by testing the Competency Questions listed in Table 5.6 of Section 5.2 (Uschold & Grüninger, 1996) on a real-world interpretations dataset, ii) to measure its granularity potential by comparing the data of the interpretations written using the model against the same interpretations encoded with other ontologies, iii) logical consistency, FAIRness and syntax using selected tools and services[9].

The following subsections deal with the creation of the evaluation data set and the different evaluation methods.

8.3.1 Evaluation dataset

We tested ICON on the section of the Iconology Dataset containing interpretations extracted from Panofsky's *Studies in Iconology* (Panofsky, 1972). This text was chosen for its historical importance and authoritativeness in the domain at hand. As a result, the dataset contains a total of 28,864 triples about 152 artworks, 1,980 interpretations, and 928 subjects. Additional statistics on the dataset can be found in Table 8.3.

Table 8.3: Test dataset overview of a number of triples, subjects, and interpretations

	Interpretations	**Subjects**	**Triples**	**Artworks**
Level 1	1,662	491		
Level 2	544	297		
Level 3	274	140		
Total	1,980	928	28,864	152

8.3.2 Competency questions evaluation

The ontology was evaluated in Panofsky's dataset using the competency questions listed in Table 5.6. Each iteration of the SAMOD methodology corresponds to a group of questions at levels 0, 1, 2, or 3, and the expected results for each CQ are described. To test the correctness of the single classes, the CQs were

[9]This section was published in Sartini et al. (2023). In detail, Sections 8.3.1, 8.3.2, and 8.3.3 are based on the Sections 6.1, 6.2, 6.3 of the article, for which S. Baroncini is responsible. Section 8.3.4 is based on Section 6.4, for which B. Sartini is responsible.

further subdivided into more detailed ones closer to the ontology structure. All unit tests that query the test dataset are available on GitHub in the form of Jupyter notebooks.[10] For each level, we describe one or two CQs. Table 8.4 contains an overview of the metadata about the artworks included in the queries.

Table 8.4: Information about artworks cited in the queries

ID	Description
ART1195	Piero di Cosimo, *The Finding of Vulcan*, 1485-1490, Hartford, Wadsworth Atheneum
ART1266	Vatican City, Biblioteca Apostolica Vaticana, *L'Ovide moralisé*, XIV Century, Cod. Reg. 1480, folio 5r. Anonymous, *Saturn*
ART1267	*Saturn*, first third of XV Century, Dresden, Kupferstichkabinett
ART1268	Jacopo Caraglio after Rosso Fiorentino, *Saturn*, Engraving B24, 1526
ART1269	*Saturn and his Children*, in Cim. 10, Middle XV Century, Berlin, Kupferstichkabinett
ART1270	Vatican City, Biblioteca Apostolica Vaticana, Cod. Pal. lat. 1368, folio 1v, XVI Century. Anonymous, *Saturn*
ART1284	Giovanni Rost after Agnolo Bronzino, *Flora*, Florence, Galleria degli Arazzi
ART1285	Albrecht Dürer, *The abduction of Proserpine*, 1516
ART1289	Nicolas Poussin, *Phaethon before Helios*, Berlin, Kaiser Friedrich Museum
ART1346	Michelangelo, *Pen drawing*, Fr. 103, 1504-1505, London, British Museum
ART1534	Rubens, *Saturn devouring a Son*, 1636-1638, Madrid, Prado

8.3.2.1 CQ Level 1

The query presented here is part of the CQ 1.2, aimed at retrieving all first-level meanings of the artworks considered, distinguishing between *Natural Elements*, Actions, and *Expressional Qualities*. The expected results, corresponding to the ones obtained, are shown in Table 8.5.

CQ 1.2.1: *Retrieve all the natural, expressional meanings, and actions recognized in the artistic motifs of ART1195.*

```
SELECT DISTINCT ?natural ?expressional ?action WHERE {
?icrec icon:aboutWorkOfArt d:ART1195;
```

[10] https://w3id.org/icon/development

```
    {?icrec icon:recognizedArtisticMotif ?am}
UNION {?icrec icon:recognizedComposition ?comp. ?comp icon:hasPart ?am
    }
{?am icon:hasExpressionalMeaning ?expressional}
UNION {?am icon:hasFactualMeaning ?natural. ?natural a icon:
   NaturalElement}
UNION {?am icon:hasFactualMeaning ?action. ?action an icon:Action}
}
```

Table 8.5: CQ 1.2.1 results

natural	**expressional**	**action**
man		
	dazed	
woman		
	charitable	
		helping
group of women		
natural landscape		
dog		
		gathering flowers
	surprise	
	amusement	
	pity	
	protectiveness	
	kindliness	
	hospitality	

The second level 1 query listed below is part of CQ 1.3, aimed at retrieving all level 1 subjects that are formally derived or copied from other artworks level 1 subjects. CQ 1.3.1 applies this question to ART1284 and ART1285. Results in Table 8.6 show how this structure can allow a detailed and qualitative comparison of the phenomenon of visual motifs copy and migration since the relation between the single portions interested can be made explicit.

CQ 1.3.1: *What are the level 1 subjects (i.e., copied subjects) copied by ART1284 from ART1285, including the ones identified by a composition? What are the corresponding original subjects in ART1285 (i.e., subjects)?*

```
SELECT DISTINCT ?subject ?copiedSubject WHERE {
?rec a icon:FormalMotifRecognition;
    icon:aboutWorkOfArt d:ART1284, d:ART1285.
{?rec icon:hasPrototypicalMotif ?am. ?am a icon:ArtisticMotif}
```

```
    UNION {?rec icon:hasPrototypicalMotif ?comp. ?comp icon:hasPart ?am}
    ?am icon:hasFactualMeaning | icon:hasExpressionalQuality ?subject.
    {?rec icon:hasCopiedMotif ?copied. ?copied a icon:ArtisticMotif}
    UNION {?rec icon:hasCopiedlMotif ?comp. ?comp icon:hasPart ?copied}
    ?copied icon:hasFactualMeaning | icon:hasExpressionalQuality ?
        copiedSubject.
}
```

Table 8.6: CQ 1.3.1 results

subject	**copiedSubject**
woman	woman
riding-on	riding-on
unicorn	ram

8.3.2.2 CQ Level 2

The correspondence of level 1 subjects with level 2 ones offers the chance to explore the variations in the subjects' representation, a fundamental research aspect for the domain of iconography and iconology. The query below represents CQ 2.6, aimed at retrieving the representative variations of a level 2 subject, to the Character 'blindfold Cupid'. This type of research question can be further explored by retrieving the artwork's date and place of production to obtain a detailed representation of the subject variations over place and time. Consequently, it can be useful for art historians to integrate qualitative iconographic analysis with a quantitative overview of the phenomenon.

CQ 2.6.1 *What are the variants of the subject 'blindfold Cupid'? Retrieve all the level 1 subjects corresponding to this subject along with how many times they appear.*

```
    SELECT DISTINCT ?lev1 (count(?lev1) as ?tot) WHERE {
    VALUES ?rel {icon:hasFactualMeaning icon:hasExpressionalMeaning}
    ?rec icon:recognizedImage ?img;
        icon:aboutWorkOfArt ?art;
        icon:refersToArtisticMotif ?am;
        icon:recognizedImage ?img.
    ?img icon:hasCharacter d:blindfold-cupid.
    {?am a icon:ArtisticMotif; ?rel ?lev1} UNION
    {?am icon:hasPart ?a. ?a ?rel ?lev1}

    } GROUP BY ?lev1
```

```
ORDER BY DESC(?tot)
```

Table 8.7: CQ 2.6.1 results

level 1	tot	level 1	tot	level 1	tot
wings	8	hearts	1	sleeping	1
bandage	7	string-of-hearts	1	natural-landscape	1
bow	6	throne	1	standing-on	1
arrows	6	arrow	1	sphere	1
boy	4	band	1	putto	1
child	4	spear	1	snuggling-in-her-lap	1
griffon-claws	3	standing-on-a-horse	1	talons	1
crown-of-roses	2	horse	1	running	1

8.3.2.3 CQ Level 3

Concerning level 3, we present a query retrieving all the artworks linked to the same cultural phenomenon, focusing on the phenomenon 'evolution of the iconography of Saturn'. This approach is useful to group all the artworks that are involved in the same cultural phenomenon as a starting point of further analysis, considered fundamental for the researcher. For example, it could be interesting to explore the second-level subjects involved in it, their variations at level 1 according to time and space, and the literary sources involved.

CQ 3.4.1 *retrieve the artworks where an intrinsic meaning is associated with the cultural phenomenon CF1087 'Evolution of the iconography of Saturn'*

```
SELECT DISTINCT ?artwork WHERE {
?rec icon:aboutWorkOfArt ?artwork;
    a icon:IconologicalRecognition;
    icon:recognizedIntrinsicMeaning ?intrinsic.
?intrinsic icon:recognizedCulturalPhenomenon d:CF1087.
}
```

Results: ART1269, ART1270, ART1266, ART1267, ART1268, ART1534, ART1535, ART1289.

8.3.2.4 General CQs

For the general-level CQs, we present two competency questions. The first (CQ 0.2.1) is presented to show how different interpretations can be represented. It retrieves all the interpretation descriptions of an artwork and the types of recognition included in it. It is performed over ART1195, which is the object of

Table 8.8: CQ 0.2.1 results

Person	**Description**	**RecognitionType**
Erwin Panofsky	ART1195-DESC1	PreiconographicalRecognition
Erwin Panofsky	ART1195-DESC1	IconographicalRecognition
Erwin Panofsky	ART1195-DESC1	IconologicalRecognition
A. E. Austin	ART1195-DESC2	IconographicalRecognition
R. van Marle	ART1195-DESC3	IconographicalRecognition
L. Venturi	ART1195-DESC4	IconographicalRecognition

contrasting interpretations. While describing it (Panofsky, 1972), the art historian Erwin Panofsky states that his position diverges from the usual interpretation according to which the artwork depicts the myth Hylas and the Nymphs. He cites the works of three other scholars as references for this general interpretation (A. E. Austin, R. van Marle, L. Venturi). In contrast, he says that it represents the finding of Vulcan by the inhabitants of the island of Lemnos after he precipitated from Mount Olympus because he was kicked out by his mother. For his interpretation, Panofsky considers various features of the first level (e.g., the general atmosphere of kindness and hospitality, which would be inappropriate to the rape and sexual aggression of the Nymphs to Hylas, described by the myth).

CQ 0.2.1 *What is the person responsible for the recognitions at each level in ART1195? Do they belong to different descriptions?*

```
SELECT DISTINCT ?personLabel ?desc ?type WHERE {
?rec icon:aboutWorkOfArt d:ART1195;
    ?rel ?desc;
     crm:P14_carried_out_by ?person;
     a ?type.
    ?desc a icon:InterpretationDescription.
    ?person rdfs:label ?personLabel

} ORDER BY ?desc
```

Results are shown in Table 8.8. The artwork has 4 different Interpretation Descriptions, among which only DESC1 has recognitions at all the levels, whereas the remaining has only second-level recognitions. To better see their agreement and disagreement, it is possible to retrieve the content of such recognitions or to see if the recognitions are already described with CiTO's relations of agreement or disagreement (`cito:agreesWith` `cito:disagreesWith`).

The second competency question (CQ 0.4) retrieves artworks that are described at both levels 1 and 3, but not at level 2. This request shows how the ontology can be used to assess the level of detail in which the subjects are described in a

dataset according to each level of interpretation.

CQ 0.4 *What artworks are interpreted on an iconological level but not on an iconographical one?*

```
SELECT DISTINCT ?art WHERE {
?rec icon:aboutWorkOfArt ?art;
     a icon:IconologicalRecognition.
?rec1 icon:aboutWorkOfArt ?art;
   a icon:PreiconographicalRecognition.
MINUS {?rec2 icon:aboutWorkOfArt ?art;
     a icon:IconographicalRecognition.
}
}
```

Result: ART1346.

All the CQs confirmed the expected results in their output. Therefore, it is possible to state that this ontology allows a meaningful representation of iconological research questions through a quantitative approach.

8.3.3 Comparison with existing ontologies

To comply with RQ2, we propose a qualitative comparison between ICON and existing ontologies. We do so by describing an authoritative example from Panofsky's bibliography, in which there are iconographical and iconological interpretations. For this purpose, we selected the most complete ontologies for iconographical descriptions, namely: Visual Representation Ontology and Wikidata.

We chose the frontispiece of François Perrier's 'Segmenta nobilium signorum statuarum' (1638 Edition)[11] depicting the iconography of Father Time, a subject that emerged from the Renaissance onward but that, despite its origins from classical sources, was never visually represented in previous times. Panofsky reconstructed its genesis, claiming that it originated by the fusion of a medieval French iconography of Time represented with wings (Temps) and the sinister characteristics of Saturn (e.g., old age, scythe, the act of devouring his children) (Panofsky, 1972). Late antique writers had already enriched the figure of the god Saturn with attributes referring to time (i.e., a dragon or snake biting its tail) or through a re-interpretation of the traditional Saturn's attributes, such as the sickle, associated with the times recurring,[12] or the act of devouring the children, reinterpreted as Time devouring 'whatever he has created' (Panofsky, 1972, p. 74).

[11] Figure available at https://www.britishmuseum.org/collection/object/P_1895-1031-28

[12] 'tempora quae sicut falx in se recurrunt', trad. by the authors: 'times which, like the sickle, recur'(Sabbadini et al., 1888, pp. II, 406)

Table 8.9: 'Time the destroyer' description according to the three levels of interpretation

1	man, nude and old, with a scythe and wings gnawing away a statue snake biting its tail fragments of classical buildings and statuary
2	Father Time as a Destroyer, symbolically devouring the past by devouring the classical Torso del Belvedere
3	Evolution of the Iconography of Time: 1) Renaissance art 'produced an image of Time the Destroyer by fusing a personification of Temps with the frightening figure of Saturn, and thereby endowed the type of Father Time with a variety of new meanings' (Panofsky, 1972) 2) Pseudomorphosis: 'certain Renaissance figures became invested with a meaning which, for all their classicizing appearance, had not been present in their classical prototypes, though it had frequently been foreshadowed in classical literature' (Panofsky, 1972)

Panofsky claims that such an evolution of the iconography of time is evidence of the phenomenon of *pseudomorphosis*, according to which figures with a classical appearance did not exist in the classical visual arts, although they were described in the classical literature (level 3). In addition to that, the example here described faithfully shows Cesare Ripa's description of time the destroyer as a demon with iron teeth standing among ruins, a symbol of the fact that time ruins everything without any effort.[13] Table 8.9 resumes the understanding of this artwork at each different level, while figures from 8.1 to 8.4 show how this example can be modeled with the ontologies considered.

Table 8.10 gives an overview of the comparison results. It is apparent from these schemas that neither VIR nor Wikidata include properties or classes that represent a third-level meaning. Nevertheless, they do express important aspects of the domain. On the first level of description, VIR offers only the limited expressivity given by the class Iconographical Atom, which is intended to describe the physical portion of the artwork to which a subject is bound (Carboni & de Luca, 2019). In contrast, Wikidata offers several level 1 specifications of the subjects through the qualifiers of the property wdt:P180 (e.g., "nudity" and 'old' referred to the subject 'man', through the qualifier wdt:P1354 'shown with features'). Nevertheless, none of the ontologies distinguishes the types of objects described, failing to express immaterial items properly, such as actions or emotions (VIR), and the connections existing

[13] La Ruina, e la Bocca aperta, e i Denti di ferro mostrano, che il Tempo strugge, guasta, consuma, e manda per terra tutte le cose senza spesa, e senza fatica.' trad. by the authors: 'The Ruina, and the Open Mouth, and the Iron Teeth show, that Time presses, spoils, consumes, and sends all things to the ground without expense, and any effort'. (Ripa, 1764)

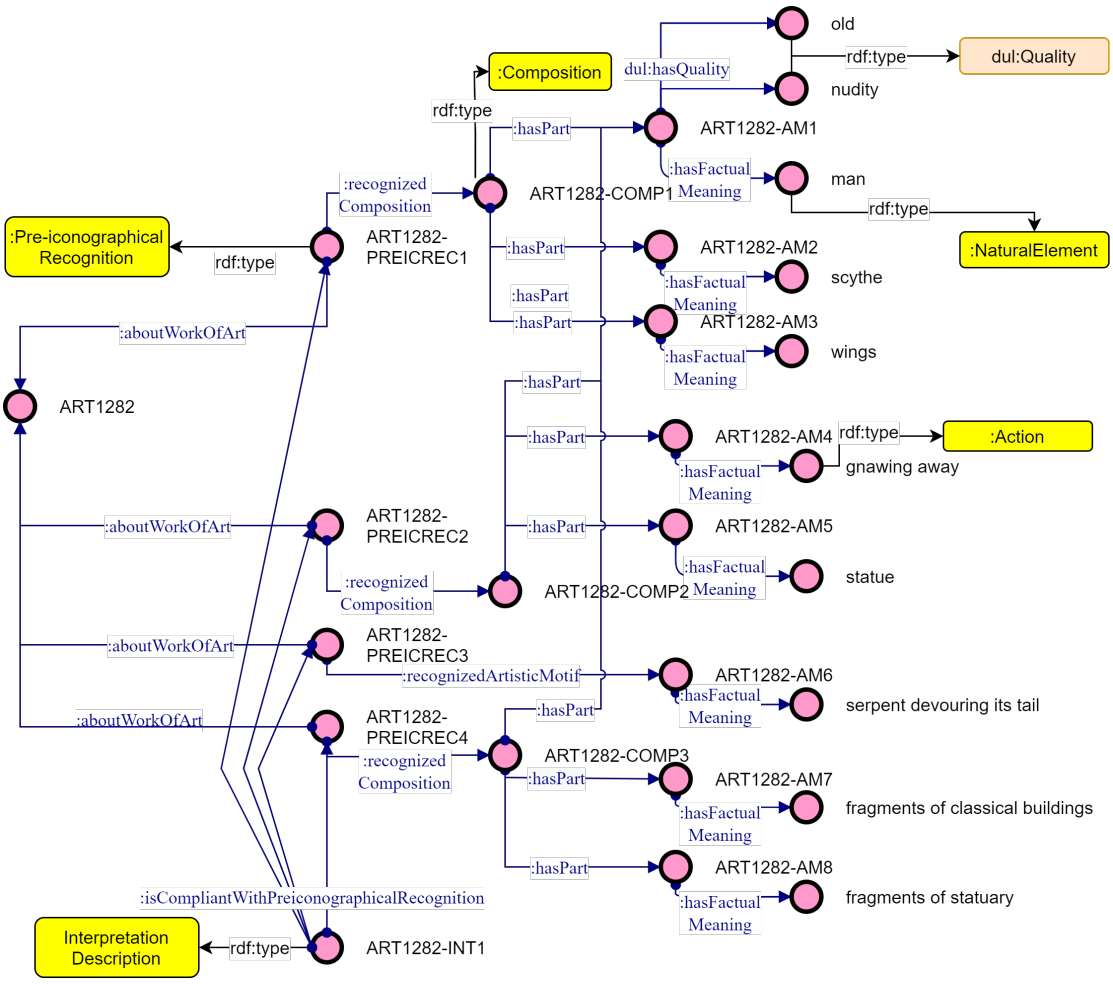

Figure 8.1: Description of the example using ICON (level 1)

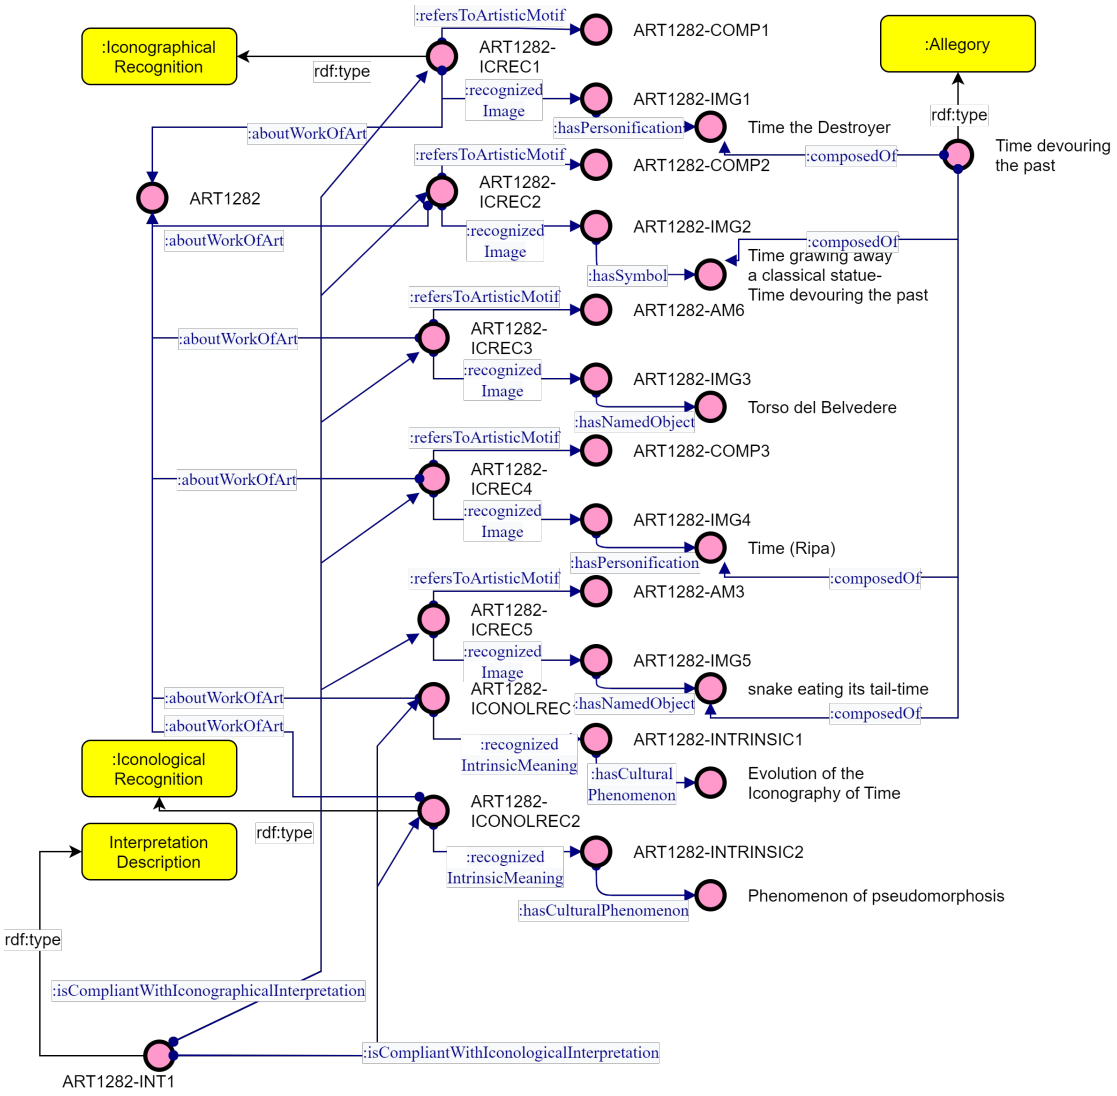

Figure 8.2: Description of the example using ICON (level 2 and 3)

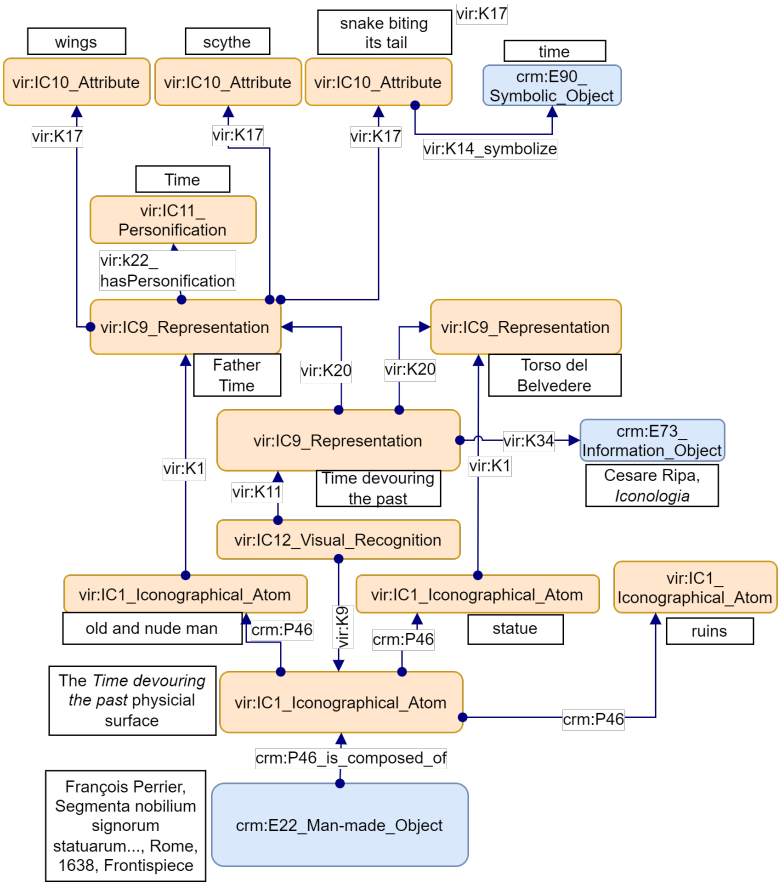

Figure 8.3: Description of the example using VIR

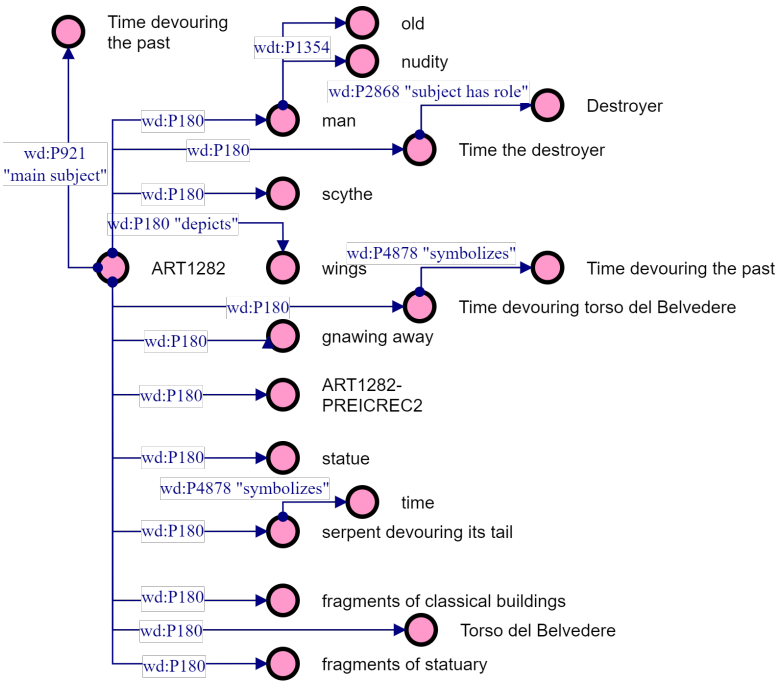

Figure 8.4: Description of the example using Wikidata

between these immaterial aspects and the subjects doing them. Therefore, ICON aims at solving this issue by introducing Panofsky's concept of the Composition of Artistic Motifs. As a result, level 1 objects as depicted in the specific artwork can not only be described in detail but also gathered in meaningful groups. This structure allows specifying the actors involved in the action, as shown in Figure 8.1 (the subjects 'man', 'gnawing away' and 'statue' are part of the same Composition ART1282-COMP2), or the actors feeling emotions.

At the second level of interpretation, with VIR it is possible to express important characteristics of the representation, such as attributes, personifications, symbols, places, and characters (Figure 8.3), but not events. Wikidata, on the contrary, tends not to specify the type of object depicted, even if a symbolical meaning can always be expressed through the qualifier `wdt:P4878` 'symbolizes' (fig. 8.4). The distinction between stories and allegories is not included in both of them and, for different reasons, the contextual appearance of the subjects cannot be carefully described. This is because in VIR it is not possible to properly describe the subjects at the first level, whereas in Wikidata, the depicted subjects cannot be related to each other. Therefore, as shown by the example, the act of a person gnawing away at a statue cannot be related to the allegory of Time devouring the past in VIR, and cannot be recognized as an allegory in Wikidata. To solve that, the ICON Iconographical Recognition allows relating the second-level subjects to its level 1 representation, and the subjects can be part of Stories or Allegories. In addition, the n-ary class `icon:Image` allows separating the general description of subjects from the contextual one. In this way, subjects can be described carefully, including characteristics that would be inappropriate to include in the vocabulary-level description of the subject considered, highlighting variations in their contextual representation.

Whereas the VIR and Wikidata features described above allow a description of the first two levels of interpretation, none of them represents the domain of knowledge of iconology by considering the third level of interpretation. Therefore, ICON introduces `icon:IconologicalRecognition` relating a third-level meaning (concept or cultural phenomenon) to the whole artwork or its specific parts.

Regarding the attribution of responsibility to the interpretation, both VIR and Wikidata allow one to register the person responsible for the statement in different ways. In VIR, the person responsible and possible sources can be related to the class `IC12 Iconographical Recognition`, on which the whole recognition of the artwork's content without a specification for every subject recognized depends. In Wikidata, every statement can be attributed to a person responsible, but information about sources is not always provided. A consistent difference lies in the fact that whereas VIR expresses the person responsible for the claim, Wikidata

considers the person responsible for the data inserted, specified through statements properties, avoiding a possible interesting comparison of authoritative art historian claims. As a result, ICON introduces an interpretation for each statement, giving subjectivity and authoritativeness for each recognized subject. In this way, more agreeing claims can be expressed in the same recognition description to represent better the realistic case in which an art historian agrees with the claim of others, quoting them and adding further interpretations. Additionally, this structure fosters the interoperability of online sources and data integration.

Table 8.10: Comparison between Wikidata, VIR, and ICON

Ontology	**Lev. 1 subj**	**Lev. 2 subj**	**Lev. 3 subj**	**Distinction contextual/general subject description**	**Interpretation subjectivity**	**Distinction and relation between levels**	**Subjects variations description**
VIR	attributes only	✓		poor	only for the main interpretation		only attributes and representations
Wikidata	✓	✓		✓	only the data author		only level 1 aspects
ICON	✓	✓	✓	✓	the person responsible for the recognition for each statement	✓	All the levels

8.3.4 Automatic Evaluation

Apart from the evaluations made to verify the explicit goals declared in our research questions, we also evaluate some more technical aspects of our ontology using automatic tools and services. We validated our ontology syntax with the W3C RDF Validation Service.[14] No syntax problems were highlighted by this tool. Then we evaluated the logical consistency of our ontology through the OOPS (Poveda-Villalón et al., 2014) tool[15] that provides feedback on the ontology in form of highlighted pitfalls of different levels of importance. Most of the issues raised by this tool do not come from our modeling; instead, they are linked to the reused ontologies that might have missing information (for example, no ranges and domain in properties or inconsistent labeling). The only highlighted flaw that was directly related to the classes developed by the ontology is 'P30: Equivalent classes not explicitly declared.' This issue suggests the possibility that classes such as `icon:Character` and `dul:Quality`, `dul:Role`, and `dul:Reference` should be equivalent. At the same time, it suggests that `icon:Image` should be equivalent to `sim:Simulacrum` and that `icon:Story` should be equivalent to `dul:Narrative`.

[14] https://www.w3.org/RDF/Validator/

[15] We suggest replicating this evaluation by pasting the whole RDF/XML file that contains the ontology content into the evaluation website (https://oops.linkeddata.es/index.jsp

The first equivalence would be fundamentally wrong because the classes themselves do not represent the same concept. The same can be said for Image and Simulacrum, as in the Simulation Ontology, a simulacrum is said to be a general symbol that could have different representations, which is why it is equivalent to our class icon:Symbol. An image is instead considered as a specific representation of a symbol in an artwork; therefore, it is not considered to be the general concept of the symbol itself. Furthermore, `icon:Story` and `dul:Narrative` might have similarities, but the latter class has no description, so we refrained from making ambiguous equivalences. Apart from these three cases, the ontology was evaluated pitfall-free in all the other aspects (considering issues that dealt with created classes only). Finally, we analyzed our ontology with the FOOPS (Garijo et al., 2021) tool, which evaluates how much an ontology complies with the FAIR principles. Our ontology scored 90%. In particular, it received a score of 8.5 out of 9 on reusability, 8 out of 9 on findability, 3 out of 3 on interoperability, and 2 out of 3 on accessibility. The main problem highlighted by this tool is that the ontology is not yet inserted in the linked open data vocabulary (lov).[16] This issue will be addressed in the future. Finally, we verified that the logical axioms of all the external imported classes and properties that were aligned with our ontology did not cause any inconsistency in ICON by running the HermiT Reasoner[17] (Glimm et al., 2014).

[16] https://lov.linkeddata.es/dataset/lov/
[17] Version 1.4.3.456

CHAPTER 9

Discussion and conclusions

This thesis explored the benefits of computational methods for the representation, retrieval, and analysis of iconography and iconology information in the context of the semantic web. Due to the lack of data and representation means for formally describing the topics of the domain, the research aimed primarily at furnishing reusable tools for iconographical and iconological statement description and a domain-specific authoritative dataset. Furthermore, the study aimed to demonstrate the usefulness of conducting quantitative research. These objectives were expressed in two main questions, focusing on 1) the provision of ontological modeling for fostering access, retrieval, and analysis of iconological interpretations and 2) assessing the advantages of browsing and quantitatively analyzing a semantic network of iconographical and iconological interpretations.

The adopted methodology included a preliminary study of the domain on which the modeling of a new ontology was based. A selection of statements from the art historian Panofsky, chosen as a case study, was described accordingly. The data analysis was then conducted based on domain-specific research questions. Further quantitative inquiries into the art historian's method and iconological approach were performed. The ontology and data set were evaluated according to quality metrics and compliance with FAIR principles.

The purpose of this chapter is to discuss the results of the research and draw conclusions. We provide an overview of the contribution of the work to the research objectives. We further discuss the key findings of the analysis conducted. Next, we highlight the impact of the research on the field of Digital Art History. Finally, we illustrate the limitations of the work, conclusions, and future work.

9.1 Research objectives and contributions

The major contribution of the current work is threefold, as it responds to the needs raised by the defined research objectives. The main research contributions to each research objective may be illustrated as follows.

- **RO1: Providing reusable means for iconography and iconology formal description**. The objective was addressed by developing the domain-specific ICON ontology for the description of artistic interpretations, which allows a thorough illustration of artworks' iconographical-iconological subjects and meanings
- **RO2: providing a semantic dataset of domain-specific art interpretations**. The creation of the Iconology dataset contributes to this objective, an RDF dataset of authoritative iconographical-iconological art interpretations of the art historian Erwin Panofsky in a structured format
- **RO3: addressing domain research questions in a quantitative way to prove the usefulness of the approach for iconographical-iconological inquiries**. A data analysis of the Iconology dataset was performed. In detail, domain-specific research questions were quantitatively addressed, and new inquiries into the iconological method were performed

In the following sections, we provide an overview of the characteristics of each contribution and the key findings of the quantitative analysis.

9.1.1 Ontological modeling

In this thesis, we modeled ICON, an ontology dedicated to conceptualizing artistic interpretations designed by formalizing the content of several interpretation theories. In line with the principles of reuse and interoperability of the Semantic Web, the ontology reuses (directly and indirectly) several existing ontologies. It is released alongside documentation that guides potential users in formalizing art interpretations using our model. ICON was evaluated for its extraction potential, syntax, metadata, and FAIRness. Moreover, its granularity was highlighted by comparing it with current ontologies on their respective serialization of the same interpretation. The results show how our work elevates the potential of expression of artistic interpretations in the context of the Semantic Web by providing a granularity level not reached by other ontologies on this topic. Finally, its effectiveness in describing the Iconographical and Iconological complex domain is confirmed by the results of the proposed competency questions, formalized in SPARQL queries, run on a test dataset containing artistic interpretations.

9.1.2 Iconology Dataset

A core part of the doctoral project centered on the collection and the description of a selection of authoritative art interpretations by the art historian Erwin Panofsky. The work resulted in a Linked Open Data dataset containing interpretations mainly by the art historian about ca. 400 artworks, mostly from Middle Ages and Renaissance Western art. The dataset is modeled according to the newly created ICON ontology. Following the good practice of reuse, we used standards and available ontologies to describe cultural objects' metadata, citations, and people's roles. The interpretations are divided into three levels, from a more superficial understanding to a deeper one, as described by the art historian's theory, and inter-level links among identified subjects are provided. The subject types recorded include natural elements, actions, and emotions (level 1), characters, events, places, objects with a specific identity (e.g., the Bible), personifications, symbols, stories, and allegories (level 2), concepts, and cultural phenomena (level 3). For each subject identification, a provenance of the assertion is provided, indicating the author, source, and, eventually, the cited evidence. The dataset was tested according to the intrinsic, contextual, representational, and accessibility quality aspects by reusing available metrics in the literature, performing a high overall score (84%). Further quality evaluations included 1) compliance with ontology constraints through SCHACL shapes and 2) a self-assessment questionnaire to verify the adherence to the FAIR principles. The capability of addressing domain-specific research questions was evaluated by calculating the percentage of the 20 sub-research domain questions formulated in the analysis that were entirely (85%) or partially (15%) addressed.

9.1.3 Data Analysis and key findings

The data analysis over the Iconology Dataset was divided into two parts. First, domain-specific research questions identified during the domain study were specified into 20 sub-questions performed over the dataset. Secondly, new inquiries into Panofsky's approach and the iconological method were conducted.

9.1.3.1 Domain-specific Research questions

It was possible to address all the sub-questions formulated for each research question despite part of them not having enough ground, due to the lack of data in our dataset. This result seems to confirm that the modeling adopted is feasible to address a good amount of domain-dependent research questions quantitatively.

The domain-specific analysis showed some key findings:

- **Recognitions of cultural phenomena are poorly supported by textual evidence and the description at other levels**. There is a poor correlation between types of cultural phenomenon and subjects at previous levels (SQ2, SQ4). Unexpectedly, the recognition of cultural phenomena is the least supported by textual sources (SQ3). On the contrary, they seem more supported by the relation with other artworks (SQ3) and tight relations with visual motif recognitions and other cultural phenomena recognitions (SQ5).

- **When a textual source is cited, the symbol often presents a symbolical meaning different from the one it has when no evidence is provided**. The analysis carried out in SQ10 only interested nine results, of which five correspond to this trend. For the limited number of examined items, the results are not generalizable. Nevertheless, such insight shows that it is possible to detect the (supposed) influence of a specific source on the traditional meaning of a symbol.

- **Scholars tend to explain Renaissance artworks with a more significant number of textual evidence, and in particular Classical sources, compared with the Medieval period**. Although both periods are well represented in the dataset, only about ten textual sources are cited as evidence of artworks before 1400 (SQ18).

Generally, all questions showed the usefulness of the provision of data as an access point to specific information, foreseeing the potential that structured data can have to support thorough qualitative analyses.

9.1.3.2 Quantitative art historiography: characterizing Panofsky's approach

RQ4.1a asks whether we can quantify Panofsky's approach from his own interpretations.[1] The following sections focus on specific aspects of the investigation: textual sources, levels of interpretation of artwork, and iconological complexity.

The usage of text sources. The deep reliance on texts proper of Panofsky's method, often criticized by other scholars, seems not to be confirmed by data, as only 27% of the interpretations cite evidence, showing only 42 texts cited.

As expected, the higher number of recognitions are of iconological nature (the third level). However, unlike Panofsky's official claim, the type of evidence mostly cited to support such recognitions is artworks and not texts. Moreover, the number of textual sources cited to support iconological recognitions is smaller than the

[1]This section is based on the version published in (Baroncini, Daquino, et al., 2023a).

number of those cited to support pre and iconographical recognitions. It should be noted that the high amount of cited artworks is also due to some modeling decisions of complex cultural phenomena. Indeed, our modeling strategy consists of recording which artworks are labeled as being part of the same cultural phenomenon when the recognition of a phenomenon in a certain artwork is also the reason why the same phenomenon is identified in another artwork. To this extent, this result would require further studies. From this first preliminary analysis, data seem to contradict the criticism raised against him by scholars who claimed that he was relying too much on texts when making considerations of iconological nature.

Nevertheless, a comparison of the results with other art historians' practices would be needed to fully understand the frequency of textual citation among the community. Therefore, these results must be interpreted with caution.

Description at multiple levels. Interestingly, unexpected results were obtained by analyzing Panofsky's method. Although Panofsky insists on the subdivision in levels of the interpretation act, it seems he tends not to explicitly make a description at all three levels for all the artworks he analyzed, as only half of the artworks are described at all the levels. In fact, most artworks having a description at only one level (21%) only include the third level of interpretation. This means that the outcome of his iconographical analysis is not presented in his texts, and it constitutes the basis for an iconological recognition without being discussed, justified (nor questioned). A possible explanation for this might be that Panofsky himself states that not all levels may be explicitly stated during an interpretation, as the context makes implicit knowledge clear (Panofsky, 1955).

Iconological interpretations. The network of iconological recognitions supporting each other, citing artworks as evidence, expanded with the network of recognized cultural phenomena, reveals a high interconnection and complexity. Therefore, we can claim it reflects the complexity of an iconological interpretation. Although it is not possible to retrace the single interpretation made by the art historian, the network allows us to connect multiple topics, making explicit previously hidden connections between interpretations. As a consequence, network visualization, as well as supporting statistical graphs, can be a valuable tool for the distant reading (Moretti, 2000) of the overall art historian's understanding of iconological phenomena.

9.1.3.3 Inquiries into the iconological method

Rq 4.1b aimed at 1) evaluating the necessity of detailed artwork description at multiple levels to understand a deeper cultural phenomenon and 2) checking if the

artwork's description grade varies across selected bibliography books due to author writing styles.

Degree of the description of artworks having a cultural phenomenon. Results of the domain-specific analysis previously described seem to suggest that recognitions of cultural phenomena are poorly supported by textual evidence and by the description at other levels. This analysis deepened the extent to which artworks are related to pre and iconographical analysis, as Panosky underlines that they are essential for a correct iconological interpretation. Results show that 10% of the phenomena (40 in total) are recognized in artworks with no further description and no apparent support for recognizing the phenomenon. Nevertheless, this number is drastically reduced by visualizing the results as a network in which artworks having common phenomena are related. Therefore, results seem to suggest that artworks described solely at the third level are part of broader networks of cultural phenomena in which at least one artwork has a multiple-level description. According to Panofsky's indication, a thorough pre- and iconographical artwork description is needed, either in the same artwork or in other ones presumably presenting similar characteristics, to achieve a deeper meaning.

Degree of the description of artworks having a cultural phenomenon in relation to the textual source. The second research question of the inquiry aimed at understanding if the artworks' degree of description varied according to the textual source describing them. To this aim, we visualized a network grouping artworks according to the textual source, distinguishing between artworks with multi-level or only third-level descriptions. As the number of artworks of the last group varies according to the textual source, we can assume that the style of the book or article has a significant impact on the detail with which the artworks are described.

9.2 Impact of the research

Besides the aforementioned contributions, the research conducted in this thesis has potential applications in related contexts.

The study offers a methodology for data-driven iconological inquiries that may be applied to data describing further iconological interpretations. By performing it over representative samples of the art of a period, interesting representative insights may be reached.

Another field of application could be the catalog entries of museum collections, in which a detailed description of the artwork, including the scholar's interpretation,

may be provided in a free text format. Such information may be converted according to ICON ontology to foster the information access and retrieval of complex iconographical and iconological subjects and meanings. Websites storing art bibliographies (e.g., Kubikat[2], OpenBibArt[3]) would similarly benefit from the inclusion in their metadata of such domain-specific descriptions. By having available the references to artworks and their iconography interpreted in the bibliographical records as part of the item metadata, art historians could easily retrieve the bibliographic references in which complex iconographies are described (e.g., searching for all the books in which an artwork depicts Venus with a slate or with a specific deeper meaning associated).

Furthermore, traditional historiographical studies may benefit from the data analysis performed. First, the unexpected results of Panofsky's method, showing a poor description of artworks and a limited reference to textual sources, can foster new qualitative studies on his method. Second, as a unique top-down definition of a theoretical approach to the discipline is challenging, a bottom-up, data-driven approach to finding commonalities among art historians' studies could help define a robust and shared method. To this end, the methodology applied here to Panofsky could be repeated for the theory and studies of other art historians.

Finally, as nearly all the artworks described have a deeper meaning (third level), the dataset could be beneficial for advancements in computer vision and machine learning studies aiming to go further in the research by inferring cultural meanings of artworks.

9.3 Limitations

Despite the positive results, the approach presents some limitations.

The ICON ontology provides a means for describing iconological interpretations granularly. However, despite descriptions of visual aspects can be performed, it is limited to figurative art only. Another pitfall is that, as a consistent use of reification is made, the ontology may not perform good results in terms of efficiency when applied to a large dataset. Nevertheless, complexity is needed when the aim is to provide a thorough description of several aspects with the detail required by domain-specific research.

Furthermore, the modeling focused on the description of works of art. However, during the dataset population phase, we observed that a significant part of the study was often dedicated to the explanation of the context or culture rather than the artworks (e.g., the tradition of a classical text in different contexts and

[2] https://aleph.mpg.de/F?func=file&file_name=find-b&local_base=kub01&con_lng=eng
[3] https://openbibart.fr/cms/?lang=en

the different meanings and textual variations it acquires time after time). The limit of the approach is evident, for example, in the book chapter *Suger abbot of Saint Denis* (Panofsky, 1955). Since the presence of artwork images was the criterion for selection, this chapter could not be included in the current study, as it described passage to a Gothic style perception through the decisions of the abbot himself, with no explicit reference to actual works of art. Additionally, the modeling adopted, which is based on the perspective of the artwork, bound the insight into the evolution of themes of interest of the domain (e.g., iconographies, cultural aspects) to the sole evidence provided by each single artwork. This prevents us from formally describing the chain of evolution reconstructed by the art historian.

Moreover, as underlined by the dataset evaluation, manual data curation guarantees data quality. Nevertheless, given the complexity of Panofsky's writings, the content extracted is subject to the annotator's interpretation.

Finally, since the dataset focuses on some core interests of the art historian Erwin Panofsky, it cannot be considered representative of Medieval and Renaissance Western art, nor should the results of the quantitative analysis be read in this sense. The limited scope also explains why we were unable to fully address some research questions.

9.4 Conclusions and future work

The current work advanced the state of the art of Digital Art History in the domain of iconology and iconography by providing a tool for formally describing art interpretation (i.e., ICON ontology), a dataset of authoritative assertions (the Iconology Dataset), and a quantitative analysis.

The analysis shows the potential that the quantitative approach may have for the research field, both in terms of distant reading and information retrieval. We demonstrated that the expression of Panofsky's theory from an ontological perspective is not only valuable to pursuing quantitative analysis for the domain itself but also for having insights into experts' claims, with unexpected results that can be the basis for further qualitative research, including both iconological and historiographical aspects. Data exploration can constitute a new perspective for deeper insights into the historian's claims. It can be expanded by including iconological interpretations of other art historians to enable comparative analyses and exploration of networks through the competing narratives of experts' interpretations.

This contribution opens up the possibility to link artworks at their content level, allowing content-based research questions in art history to cross into the linked open data realm. The digital space thus becomes the place where multiple

personal views gather, representing in a faithful way the traditional way in which specialistic knowledge is expressed, allowing distant viewing of scholar's thoughts.

Future work is envisioned in multiple directions: ontological modeling, data description, analysis, and retrieval.

For the relevance given to sociocultural aspects in iconological studies, future ontological development includes modeling cultural changes to be pursued in collaboration with experts from anthropology and social sciences. Moreover, more work will need to be done to model the relevant topic of the concept of evolution to describe the iconographical variations and resemantization provided by the scholars' interpretation. A further study could also assess the feasibility of expanding the model to include 1) non-figurative art descriptions and 2) a lighter ontology version.

Furthermore, we envision the creation of a dataset representative of a particular period to extend the investigation and provide the discipline with valuable and representative results.

Finally, to foster access to data by art historians, work to study interfaces to simplify the interaction with data, offering a high level of data collection and information retrieval customization without the need to code, should be addressed.

Part III

IN-PROGRESS APPLICATIONS

CHAPTER 10

Tool for art information integration and retrieval

The work conducted in the thesis project pushed forward the availability of authoritative domain-specific data. However, to encourage their use by art historians and students to concretely achieve the integration of quantitative analyses in the research area, applications that do not require code skills are needed. In the context of an internship at Sony CSL Paris,[1] I collaborated with the European project MUHAI[2] for the development of a user-centered tool aimed at interactively retrieving semantic data about artworks from existing KGs. The central data structure used to reach this objective is the Integrative Narrative Network (INN), a generalization of graphs that semantically represent the process of understanding through a narrative by retrieving and aggregating the knowledge available in existing knowledge bases. In this case, the objective was to exploit the INN structure as a human-centric tool to help humans collect online information and represent it as part of their interpretation. The INN represents the user's progress in understanding, while knowledge about a specific artwork under examination is retrieved from available Knowledge Bases. The research question was addressed through the case study of the representation of a Renaissance painting.

Although the ongoing project is still in its preliminary phases, it can provide insight into possible developments to make domain-specific semantic data more accessible for non-programming art historians.

The following sections illustrate the project. Fist, an introduction to narrative-based understanding and the selected case study is provided in Section 10.1. Then, Section 10.2 introduces the INN data structure. Following, the methodology and results are discussed in Section 10.3. Finally, Section 10.4 draws some conclusions

[1]This research was supported by the EU Pathfinder project MUHAI (EU grant 951846) and by the international exchange program of Bologna University.

[2]https://www.muhai.org/

and an outline of future work.

10.1 Narrative-based Art Interpretation

10.1.1 Understanding as a question-answering

There has been remarkable progress in the past decade in tools to analyze and contextualize artworks based on techniques from computer vision, pattern recognition, ontology engineering, natural language processing, and the semantic web.[3] What is missing, however, are good ways to combine the outcomes of these various tools into one coherent interpretation and to make the results available both to art historians and to viewers.

The human interpretation process strives for narrative-based understanding. When interpreting an artwork the viewer tries to construct a narrative that integrates the visual depictions, memory of past experiences, knowledge of artistic styles, the personal history and prior work of the artist, and general world knowledge, in order to answer a series of questions (e.g., Who is depicted? Why?) and to put an artwork in context (e.g., Who is the painter? What was the function of the work?).

Narrative-based understanding has to be conceived of as a spiraling process.(Gadamer, 1975; Heidegger, 1962). Starting from an initial examination of some input elements with a lot of ambiguity, uncertainty, and indeterminacy, hypotheses of the whole are constructed, which then provide top-down expectations to be tested by a more detailed examination of the same or additional elements, leading to a clearer view of the whole, which then leads back to the examination of additional elements, etc., until the narrative 'makes sense' and resonates with the personal episodic memory of the viewer, reaching a state known as *narrative closure* in literature studies (Carroll, 2007).

When considering partial digital support for interpretation we need not only work on the individual knowledge sources that can answer or raise questions but also on the integration process and this requires as a first, fundamental step a data structure that supports integration.

For this purpose, we present a data structure called **Integrative Narrative Network (INN)**. An INN acts as a kind of blackboard on which different knowledge sources write partial descriptions and through which they can consult the information provided by other knowledge sources to advance understanding. We

[3]This Chapter is based on (Baroncini, Steels, et al., 2023). In detail, Sections 10.1.2 10.3 are based on Sections 1.2, 3.1, 3.2, and 3.4 of the article, for which S. Baroncini is responsible. Sections 10.1.1 and 10.2 are based on the article's Sections 1.1 and 2, for which L. Steels is responsible. Section 10.3.3 is based on the article's section 3.3., for which R. van Trijp is responsible.

operationalize the understanding process in terms of narrative questions and answers, defining questions as open slots that can be filled by knowledge sources by (i) evoking new questions, (ii) introducing answers to questions, (iii) constraining the answers to questions, or (iv) shrinking the set of questions by realizing that the answers to two different questions are the same.

10.1.2 Case study

In this chapter, we explore historical artworks from the late Renaissance period by the Venetian painter Lorenzo Lotto and we focus on integrating results from semantic web resources, specifically ontologies, knowledge graphs, and Wikidata. More concretely, we report on a use case for the painting 'Venus and Cupid' by Lorenzo Lotto at the Metropolitan Museum in New York[4] (see Fig. 10.1). This painting was chosen as a feasible case study for art understanding since the lack of textual documentation and its complex iconography constitute an interpretative challenge for art historians (Christiansen, 1986).

At first glance, the observer's attention is driven by some visual indicators, such as the brightness contrast, to the body of the main characters, i.e. a putto (little angel) and a woman. If the observer is familiar with the art themes depicted during the Italian Renaissance, s/he will recognize immediately these figures respectively as Cupid and his mother Venus, thanks to the identification of their attributes. Nevertheless, the actions performed (e.g., Cupid's act of urinating through a wreath, with a playful expression) and the presence of other objects (e.g., the snake, the cone shell, etc.) may lead the observer to raise new questions: 'What exactly is depicted?', 'Why are these objects represented?', 'What is the scene telling us?'.

These questions may be further enhanced by the observer's past experience with other paintings by Lotto, who often depicted symbols or visual anagrams in his artworks. At this point, the observer, particularly if s/he is an art historian, will start to look up the possible meanings of the iconic objects through visual and textual sources and verify if the newly discovered knowledge about it is coherent with the context.

Narrative closure is reached when answers to all these questions are coherently integrated with each other and grounded in external and past experiences. For example, the art historian Keith Christiansen recognizes that the curious action made by Cupid and the majority of the objects (e.g., the red cloth, the snake, the myrtle wreath, the incense burner) belong to the sphere of marital love, concluding that the artwork expresses a wedding wish[5].

[4]Catalog entry available at https://www.metmuseum.org/en/art/collection/search/436918

[5]A more thorough explanation of this interpretation can be found in the artwork catalog entry of the Metropolitan Museum.

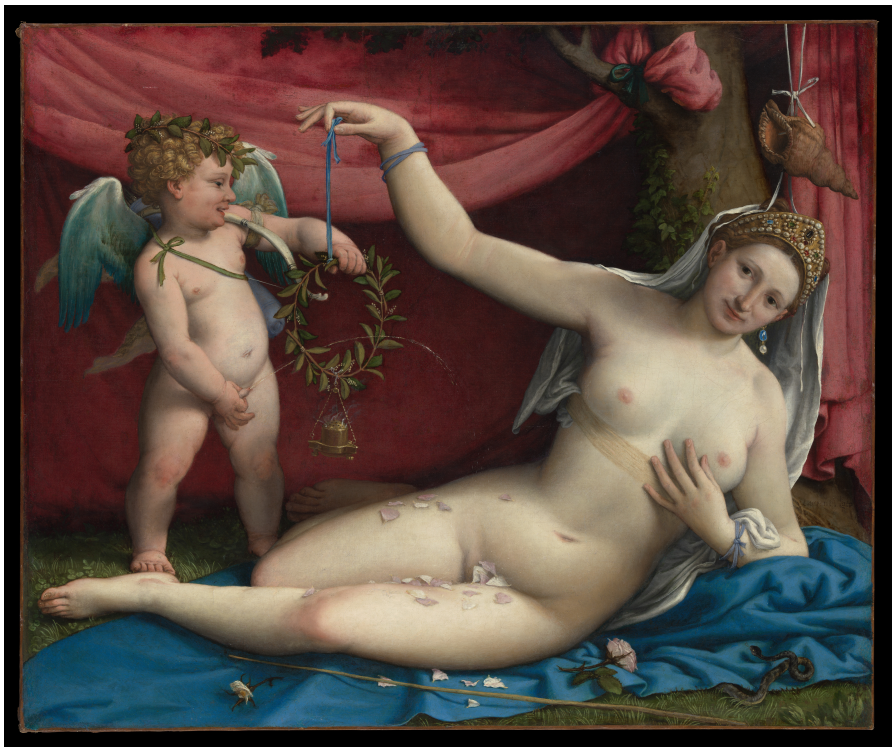

Figure 10.1: Lorenzo Lotto, *Venus and Cupid*, 1520s, Metropolitan Museum of Art, New York

Note that the focus of the interpreter has gone beyond the main focal point of the painting (i.e. Venus and Cupid), to explore secondary details, which are now seen as signs, triggering a deeper interpretation of the scene.

10.2 Integrative Narrative Networks (INN)

As explained, we frame art interpretation as a process whereby the agent raises questions, finds answers to questions, and interlinks questions and answers. We focus in particular on how semantic web resources can help to raise questions and find answers to them.

What are the questions? Questions are computationally operationalized as variables. Following AI tradition, the name of the question is written as a symbol with a question mark in front, as in ?question-name (e.g., ?Where).

What are the answers? The answers to a question are entities in the domain of discourse. Entities are objects, events, or (reified) concepts. They either refer to real-world observational data (for example a physical painting, a region in an image), to virtual entities (which may or may not exist in reality), or to entities in a knowledge graph in which case we use the URI (Universal Resource Identifier) as a unique identifier.

INN makes the relation between a question and an answer computationally

operational in terms of binding between questions (which are technically variables) and identifiers of entities or constants.

Where do the questions come from? We use a frame-based approach that was pioneered in research on knowledge representation and object-oriented programming, starting in the mid-seventies with the proposal by (Minsky, 1975) and technical realizations such as KRL (Bobrow & Winograd, 1977), KRS (Steels, 1985) or CLOS (Gabriel et al., 1991). In the present INN implementation, we use CLOS. In Minsky's original conception, a frame is a bundle of important questions to be asked about a particular type of object. A frame has a set of slots with values, which are in effect the questions that can be asked about a particular entity represented by a frame and respectively the answers to these questions. An entity that is described with a certain frame is called a frame instance of that frame.

Visualization The INN represents all questions that have been posed during the understanding process and all entities that appear as answers to these questions. Questions are visualized with green or red diamonds. The diamond is green if the question can be answered and red if it is still open. The answers to questions are represented with squares and can be frame-instances or constants (e.g., a number or a boolean value).[6]

10.3 Case study for the Lotto painting

This section illustrates the INN and its use in interpretation for a concrete case study on the painting *Venus and Cupid* by the late Renaissance painter Lorenzo Lotto introduced earlier (Section 10.1.2). The case study has been implemented in the sense that all the accesses to diverse knowledge sources are done through queries and then integrated into the INN.

10.3.1 Queries to knowledge sources

The possible interests of users have been adopted from the didactic art literature on how to conduct an interpretation of an artwork [7] (Barrett, 2003). The suggested questions were grouped as sub-questions of five wh-questions, namely *Who? What? Where? When? Why?*. The questions were then expressed as SPARQL queries to be performed over an initial selection of knowledge bases (illustrated in Table

[6] Our paper is supported by an open-source software implementation of integrative narrative networks for Art Interpretation (Apache 2.0 license) as part of the Babel cognitive software suite (Loetzsch et al., 2008; Steels & Loetzsch, 2010) which is available at https://github.com/SonyCSLParis/art-network.

[7] University of Manchester, Art historical research. Url: https://library.wcupa.edu/_art_history

10.1). For example, to answer the question about ?Who (i.e. who is the painter?) the following SPARQL query can be launched to the Wikidata knowledge base[8]:

```
SELECT ?author ?authorLabel
   WHERE {
   wd:Q4009580 wdt:P170 ?author.
   ?author rdfs:label ?authorLabel.
   }
```

The answer from a query (if successful) is then integrated as an answer bound to the question ?Who in the INN. The criteria for selecting knowledge bases is the possible presence of data about artworks, their content, and/or possible meanings, in compliance with the results obtained by the study in (Baroncini, Sartini, et al., 2023). Furthermore, results obtained from computer vision analysis (e.g., some objects identified within a certain portion of the image) were integrated into the Narrative Network through a test on Google API Cloud Vision[9].

More concretely, we searched for the painting ID on the KBs through SPARQL queries looking for the artist's name and title or looking for alignments with the ID previously found (e.g., Wikidata's ID). We then performed the queries expressing the 5W questions on each KB. During the process, results were represented according to the INN data structure, showing a graph that progressively expanded from the initial artwork node. Nodes retrieved from different KBs that were possibly referring to the same concept (e.g., the representation of the character 'Venus' in Wikidata and in Zeri&Lode) were detected through a label similarity fuzzy ratio and Wordnet synonyms, and through the presence of common URI to which they were already aligned.

10.3.2 Results

In the initial phase, we retrieved the artwork ID, its title, and its image, if present, from the selected knowledge bases, starting from Wikidata. The artwork was described only in 3 of them, namely Wikidata, the Iconology Dataset, and Zeri&Lode. The first node, representing the artwork, was added to the graph as an INN `Artwork` node, which constituted the starting point of the information retrieval. The 5W narrative questions and their sub-questions were added[10]. Figure 10.2(a) illustrates how the network visualization appears at this stage. The following

[8]The SPARQL queries performed are available at https://inn-painting-art-interpretation.streamlit.app/Narrative_Questions_and_Queries

[9]https://cloud.google.com/vision

[10]The visualization of the graph at each stage is available at https://inn-painting-art-interpretation.streamlit.app/Network_Evolution

Table 10.1: Knowledge bases included in the case study.

Knowledge Base	**Description**
Wikidata (Vrandečić & Krötzsch, 2014)	Open, collaborative KB hosting structured data from other Wikimedia projects
ArCo (V. A. Carriero et al., 2019)	Italian Cultural Heritage
Zeri&Lode (Daquino et al., 2017)	The photographs about XVI Century art of Federico Zeri's photographic collection
The Iconology Dataset	Iconological interpretations about ca 400 artworks, with a focus on the art historian Erwin Panofsky
HyperReal (Sartini et al., 2021)	Symbols and symbolic meanings in different contexts

sections illustrate in detail the network expansion according to each main question, as information is retrieved from the cited KBs.

10.3.2.1 ?Who

This question concerned the author since no information about other involved people (e.g., the patron) was available on the KBs. It was possible to retrieve the author's date of birth and death, places and art movement. Since this information partially answered other open narrative questions, such as `?When` and `?Art movement` (sub-question of `?What`), a relation between them was added. It was also possible to retrieve other paintings by the same author, which is relevant information for art historical research, e.g., for style comparison, or analyzing if and how the painter depicted similar subjects.

10.3.2.2 ?What

This question considered the retrieval of metadata, such as the title, type, genre, material, and art movement of the painting. All the sub-questions found an answer on Wikidata and Zeri. We then answered the narrative question `?Subject`, sub-question in the `?What` variable, which retrieved a great amount of information. In Wikidata, two ways for registering subjects are available, namely specific objects (`wdt:P180` 'depicts') or the main subject (`wdt:P921` 'main subject'). The graph is enriched with the subjects depicted (`wdt:P180` 'depicts'). Whereas Zeri&Lode included only one entity for the subject ('Venere e Cupido'), related to the Wikidata entity thanks to their common alignment to the Iconclass code 92C454, the Iconology dataset provided a thorough description. The identified subjects are here described according to three layers of understanding, showing how more concrete objects (e.g., 'woman') relate to iconographies (e.g., 'Venus')

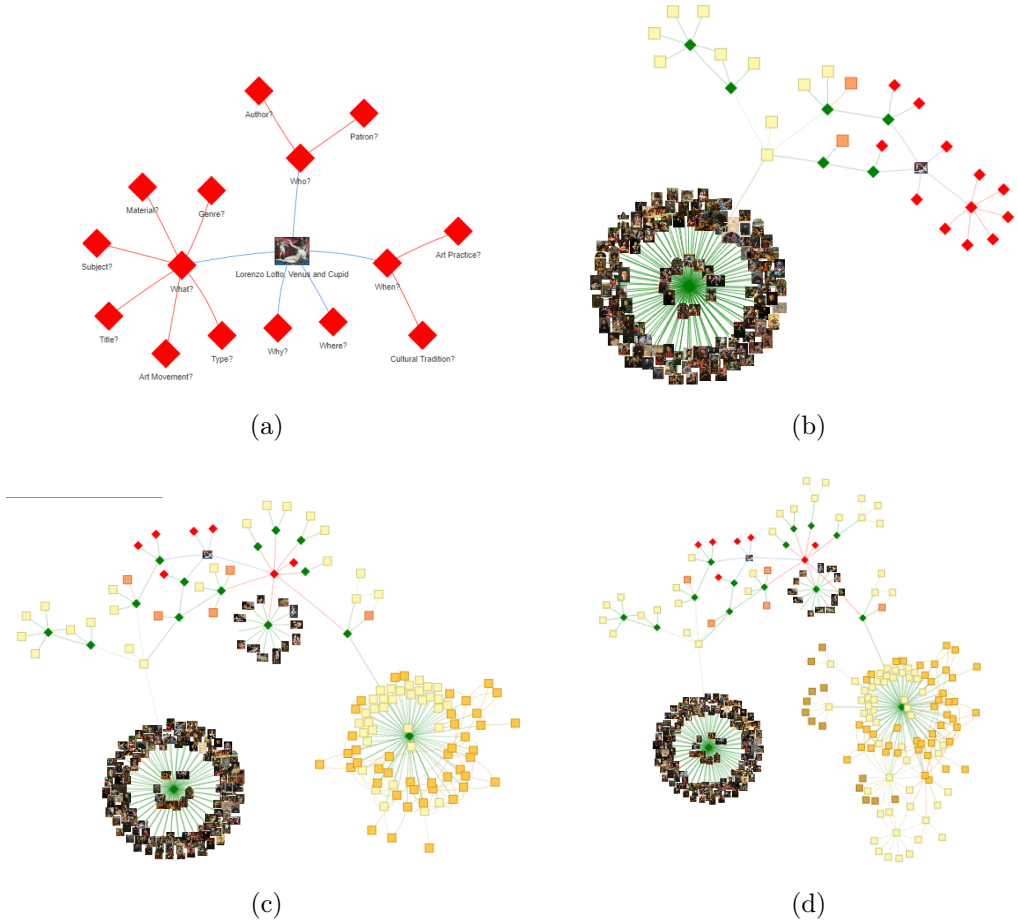

Figure 10.2: How the network appears (a) with the 5WH questions and sub-questions still unbound, (b) after answering the ?Who question, (c) the ?What question, and (d) at the end of the information retrieval.

and to deeper meanings (e.g., 'wedding wish'). For answering this question, also the results obtained by CV analysis were integrated, despite only the concept of 'person' was recognized, and the algorithm erroneously recognized Venus' diadem as a hat. Following, the sources of the depicted theme were retrieved. In detail, the painting is based on an epithalamium, i.e., a classical genre of poems written to celebrate weddings (source: Wikidata), and various objects and their meanings have specific sources, e.g., the incense burner as a proper bridal chamber decoration have as source the epithalamium of the roman poet Sidonius (source: Iconology Dataset).

10.3.2.3 ?When

As anticipated, the retrieval of information about the artist partially answered the question, restricting the range of the possible date of the artwork creation to the one of the author's life. Nevertheless, the information about a more precise date of creation was provided by Wikidata (inception 1530) and Zeri (1520-1556).

10.3.2.4 ?Where

For this section, the main questions raised for a better understanding of the artwork are 1) the location where the artwork was created, and 2) the place where the artwork was intended to be displayed. Whereas the first question can be potentially answered through the modeling adopted by Zeri, the second one can be retrieved on Wikidata. Nevertheless, no further information was available for Lotto's painting. Therefore, the question is partially answered by the information about the author's work location previously retrieved in the `?Who` question. We consequently added a relation between the narrative questions `?Location` and `?Where`.

10.3.2.5 ?Why

The narrative question `?Why` had as a main sub-question the retrieval of the patron, namely the person who commissioned the artwork. Although this sub-question can be formally expressed with Wikidata ontology with the relation `wdt:P88`, such information is not available for the chosen case study. Therefore, the narrative question remains unbound.

10.3.2.6 Second iteration

After the first binding of the main information to the narrative questions, we performed a second iteration to answer potential questions raised by the user. For

example, a user may be interested in continuing the exploration of a certain KB starting from a retrieved node, or retrieving more possible meanings through the exploration of new KBs. For example, we retrieved more information about Venus, Cupid, and their Greek counterparts (Aphrodite and Eros), discovering that they are gods of love and that the latter is the son of the former. Secondly, we explored the potentially embedded symbolism not yet discovered by querying HyperReal. As a result, we included in the INN representation potential symbolical meanings of Venus, the tree, and the cone shell, discovering, for example, that the conch is not only one of the attributes of Venus but is also a symbol of fertility and gestation, concepts supporting the meaning of the artwork as a wedding wish.

10.3.3 Utilisation of the INN by art historians

The INN is a central data structure around which various applications can be developed. One application, which we have operationalized, is intended for art historians. It allows the historian to direct the expansion of the INN by indicating which questions should be preferentially explored and by choosing which knowledge sources should be utilized preferentially. Figure 10.3 shows a snapshot of this interface. The user-centered implementation of the INN data structure is a solution for solving the integration of competing or erroneous information retrieved from different types of sources. The user can indeed remove the erroneous information from his/her own network. Future implementation considers the inclusion of feedback to the knowledge sources about the correctness of the retrieved data, and the chance to integrate the newly added information to collaborative graphs (e.g., Wikidata), to foster the improvement of data quality and the enrichment of KBs with experts' knowledge.

10.3.4 Discussion

By implementing the questions raised by the literature, the major part of the questions could be expressed as SPARQL queries, performed and answered with currently available information on knowledge graphs. Some answers were bound to multiple questions, as they answered both. Although multiple information could be aligned through the strategies of reconciliation with a common vocabulary, word similarity, and synonym detection, challenges in information integration remain open (e.g., when pieces of information slightly differ, as happened with the date of creation). As a solution, we make the INN available as an interactive tool, in which the user can manually select the desired data and perform reconciliations.

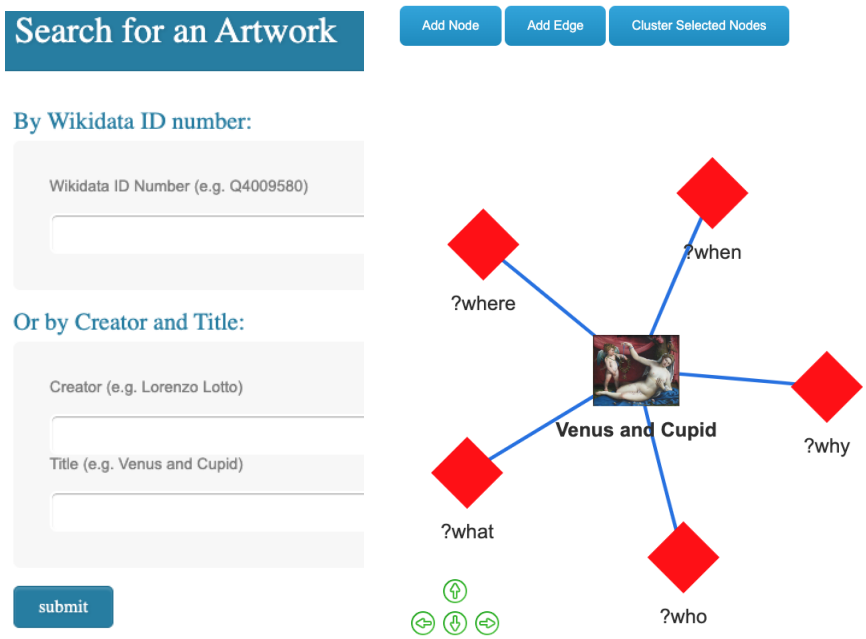

Figure 10.3: This Figure shows an interface for art historians to steer the INN expansion process. The historian can click on certain questions to expand them further and zoom in and out.

10.4 Conclusions and future work

The study presented in this section introduced an application to the art history domain of Integrative Narrative Networks (INN), a data structure for representing progress in interpreting artworks in the form of a graph that connects questions and answers for a specific work. Narrative closure occurs when the main questions of relevance to the human interpreter have been answered. The paper defined the INN and focused on how semantic web resources about art history and interpretation, which are becoming more and more available, can be marshaled to push the interpretation forward. We used a case study of a painting by Lorenzo Lotto to illustrate the proposed methods and techniques. In future work, we plan to integrate additional resources, conduct more case studies, expand the capabilities of the user interface for art historians with the aid of a user study, make advancements toward the generation of KB-independent queries, and develop another interface embedded in an augmented reality device so that viewers seeing an artwork in situ can also interactively explore the semi-autonomous expansion of the Integrative Narrative Network.

Bibliography

Abu-Salih, B. (2021). Domain-specific knowledge graphs: A survey. *Journal of Network and Computer Applications*, *185*. https://doi.org/10.1016/j.jnca.2021.103076

Adams, L. 1. (2010). *The methodologies of art an introduction* (2nd ed). Westview Press.

Andrews, T. (2023). The structured assertion record (STAR) model for event-based representation of historical information. *GrapHNR 2023 : Graphs and Networks in the fourth dimension – time and temporality as categories of connectedness*.

Antal, F. (1937). Some Examples of the Role of the Mænad in Florentine Art of the Later Fifteenth and Early Sixteenth Centuries. *Journal of the Warburg Institute*, *1*(1), 71–73. https://doi.org/10.2307/750076

Auer, S., Bizer, C., Kobilarov, G., Lehmann, J., Cyganiak, R., & Ives, Z. (2007). DBpedia: A nucleus for a web of open data. In *The semantic web* (pp. 722–735). Springer.

Baroncini, S., Daquino, M., & Tomasi, F. (2021). Modelling Art Interpretation and Meaning. A Data Model for Describing Iconology and Iconography. *AIDAinformazioni*, *1-2*, 39–62. https://doi.org/10.48550/arXiv.2106.12967

Baroncini, S., Daquino, M., & Tomasi, F. (2023a). Are domain-specific theoretical approaches valuable for the application of new computational methods? The case study of Erwin Panofsky's artworks interpretations and the semantic web. In *Althage, Melanie / Dröge, Martin / Hiltmann, Torsten / Prinz, Claudia (Hrsg.), Digitale Methoden in der geschichtswissenschaftlichen Praxis: Fachliche Transformationen und ihre epistemologischen Konsequenzen: Konferenzbeiträge der Digital History 2023, Berlin, 23.-26.5.2023, Berlin 2023 (https://doi.org/10.5281/zenodo.8319631)*. Zenodo. https://doi.org/10.5281/zenodo.8321994

Baroncini, S., Daquino, M., & Tomasi, F. (2023b). Exploring iconographical and iconological content in semantic data through art historians' interpretations.

La memoria digitale. Forme del testo e organizzazione della conoscenza. Atti del XII convegno annuale AIUCD, 9–14.

Baroncini, S., Sartini, B., van Erp, M., Tomasi, F., & Gangemi, A. (2023). Is dc:subject enough? a landscape on iconography and iconology statements of knowledge graphs in the semantic web. *Journal of Documentation*, *79*(7), 115–136. https://doi.org/10.1108/JD-09-2022-0207

Baroncini, S., Steels, L., & Trijp, R. v. (2023). Semantic Data Retrieval and Integration for Supporting Artworks Interpretation Through Integrative Narrative Networks. *SWODCH'23: International Workshop on Semantic Web and Ontology Design for Cultural Heritage*.

Barrett, T. (2003). *Interpreting art: Reflecting, wondering, and responding*. McGraw-Hill New York.

Bartalesi, V., Meghini, C., & Metilli, D. (2017). A conceptualisation of narratives and its expression in the CRM. *International Journal of Metadata, Semantics and Ontologies*, *12*(1), 35–46. https://doi.org/10.1504/IJMSO.2017.087692

Behkamal, B., Kahani, M., Bagheri, E., & Jeremic, Z. (2014). A metrics-driven approach for quality assessment of linked open data. *Journal of theoretical and applied electronic commerce research*, *9*(2), 64–79.

Bekiari, C., Bruseker, G., Doerr, M., Ore, C.-E., Stead, S., & Velios, A. (2021). Definition of the cidoc conceptual reference model v7.1.1. https://doi.org/10.26225/FDZH-X261

Bellotto, A. (2020). Medieval manuscript descriptions and the Semantic Web: Analysing the impact of CIDOC CRM on Italian codicological-paleographical data. *Digital Humanities Quarterly*, *014*(1).

Berners-Lee, T., Hendler, J., & Lassila, O. (2001). The semantic web. *Scientific American*, *284*(5), 34–43.

Bikakis, A., Hyvönen, E., Jean, S., Markhoff, B., & Mosca, A. (2021). Editorial: Special issue on semantic web for cultural heritage. *Semant. Web*, *12*(2), 163–167.

Bobasheva, A., Gandon, F., & Precioso, F. (2021). Learning and reasoning for cultural metadata quality: Coupling symbolic ai and machine learning over a semantic web knowledge graph to support museum curators in improving the quality of cultural metadata and information retrieval. *J. Comput. Cult. Herit.* https://doi.org/10.1145/3485844

Bobrow, D. G., & Winograd, T. (1977). An overview of KRL, a knowledge representation language. *Cognitive Science*, *1*(1), 3–46. https://doi.org/10.1016/S0364-0213(77)80003-7

Borgo, S., Ferrario, R., Gangemi, A., Guarino, N., Masolo, C., Porello, D., Sanfilippo, E. M., & Vieu, L. (2022). DOLCE: A Descriptive Ontology for Linguistic and Cognitive Engineering. *Applied Ontology*, *17*(1), 45–69.

Brandhorst, H. (2022). Meaningful patterns. Can a mix of pattern matching and computer-assisted classification produce a paradigm shift in iconography? In *Bibliothek und Wissenschaft 55 (2022): Faszination (Buch-)Handschriften im Jahr 2022* (pp. 287–306). Harrassowitz.

Brouwer, J., & Nijboer, H. (2017). Golden Agents. A web of linked biographical data for the Dutch Golden Age. *Proceedings of the Second Conference on Biographical Data in a Digital World 2017*, 33–38.

Candela, G. (2023). An automatic data quality approach to assess semantic data from cultural heritage institutions. *Journal of the Association for Information Science and Technology*, *74*(7), 866–878. https://doi.org/10.1002/asi.24761

Carboni, N., & de Luca, L. (2019). An ontological approach to the description of visual and iconographical representations. *Heritage*, *2*(2), 1191–1210. https://doi.org/10.3390/heritage2020078

Cardini, F. (2001). *The Chapel of the Magi in Palazzo Medici*. Mandragora.

Carriero, V., Daquino, M., Gangemi, A., Nuzzolese, A., Peroni, S., Presutti, V., & Tomasi, F. (2020). The Landscape of Ontology Reuse Approaches. *Appl. Practices Ontol. Des., Extraction, Reason*, *49*, 21.

Carriero, V. A., Gangemi, A., Mancinelli, M. L., Marinucci, L., Nuzzolese, A. G., Presutti, V., & Veninata, C. (2019). ArCo: The Italian Cultural Heritage Knowledge Graph. In C. Ghidini, O. Hartig, M. Maleshkova, V. Svátek, I. Cruz, A. Hogan, J. Song, M. Lefrançois, & F. Gandon (Eds.), *The Semantic web – ISWC 2019* (pp. 36–52). Springer International Publishing.

Carriero, V. A., Gangemi, A., Mancinelli, M. L., Nuzzolese, A. G., Presutti, V., & Veninata, C. (2021). Pattern-based design applied to cultural heritage knowledge graphs. *Semantic Web Journal*, *12*(2), 313–357.

Carroll, N. (2007). Narrative closure. *Philosophical Studies.*, *135*, 1–15.

Christiansen, K. (1986). Lorenzo lotto and the tradition of epithalamic paintings. *Apollo (London. 1925)*, *124*(295), 166–173.

Cieri Via, C. (1994). *Nei dettagli nascosto: Per una storia del pensiero iconologico*. NIS.

Cimmino, A., Fernández-Izquierdo, A., & García-Castro, R. (2020). Astrea: Automatic Generation of SHACL Shapes from Ontologies. In A. Harth, S. Kirrane, A.-C. Ngonga Ngomo, H. Paulheim, A. Rula, A. L. Gentile, P. Haase, & M. Cochez (Eds.), *The Semantic Web* (pp. 497–513). Springer International Publishing.

Ciotti, F. (2016). Toward a formal ontology for narrative. *MATLIT: Materialidades da Literatura*, *4*(1), 29–44.

Cohen, J. (1960). A coefficient of agreement for nominal scales. *Educational and psychological measurement*, *20*(1), 37–46.

Costa, P. C. G., Laskey, K. B., & Lukasiewicz, T. (2010). Uncertainty representation and reasoning in the semantic web. In *Web Technologies: Concepts, Methodologies, Tools, and Applications* (pp. 1852–1877). IGI Global.

Damiano, R., & Lieto, A. (2013). Ontological Representations of Narratives: A Case Study on Stories and Actions. https://doi.org/10.4230/OASICS.CMN.2013.76

Daquino, M., Mambelli, F., Peroni, S., Tomasi, F., & Vitali, F. (2017). Enhancing Semantic Expressivity in the Cultural Heritage Domain: Exposing the Zeri Photo Archive as Linked Open Data. *J. Comput. Cult. Herit.*, *10*(4). https://doi.org/10.1145/3051487

Daquino, M., & Tomasi, F. (2015). Historical context ontology (hico): A conceptual model for describing context information of cultural heritage objects. *Research Conference on Metadata and Semantics Research*, 424–436.

Davis, E., & Heravi, B. (2021). Linked Data and Cultural Heritage: A Systematic Review of Participation, Collaboration, and Motivation. *J. Comput. Cult. Herit.*, *14*(2). https://doi.org/10.1145/3429458

Didi-Huberman, G. (2017). *The surviving image: Phantoms of time and time of phantoms: Aby Warburg's history of art*. The Pennsylvania State University Press.

Doerr, M., Bekiari, C., Kritsotaki, A., Hiebel, G., & Theodoridou, M. (2014). Modelling scientific activities: Proposal for a global schema for integrating metadata about scientific observation. *Access and understanding–networking in the digital era: The 6th annual conference of CIDOC, the International Committee for Documentation of ICOM, Dresden, Germany.*

Färber, M., Bartscherer, F., Menne, C., & Rettinger, A. (2018). Linked data quality of DBpedia, Freebase, OpenCyc, Wikidata, and YAGO. *Semantic Web*, *9*(1), 77–129.

Fensel, D., Şimşek, U., Angele, K., Huaman, E., Kärle, E., Panasiuk, O., Toma, I., Umbrich, J., & Wahler, A. (2020). Introduction: What is a knowledge graph? In *Knowledge graphs: Methodology, tools and selected use cases* (pp. 1–10). Springer International Publishing. https://doi.org/10.1007/978-3-030-37439-6_1

Filipiak, D., Agt-Rickauer, H., Hentschel, C., Filipowska, A., & Sack, H. (2016). Quantitative Analysis of Art Market Using Ontologies, Named Entity Recog-

nition and Machine Learning: A Case Study. *255*. https://doi.org/10.1007/978-3-319-39426-8_7

Frank, I. (2019). Multi-perspectival representation of historical reality. *JOWO 2019: Proceedings of the Joint Ontology Workshops 2019. Episode G: The Styrian Autumn of Ontology, Graz, Austria, September 23-25*.

Freire, N., & Isaac, A. (2020). Wikidata's linked data for cultural heritage digital resources: An evaluation based on the Europeana Data Model. *International Conference on Dublin Core and Metadata Applications*, 59–68.

Fülöp, C., Kiss, G., Kovács, L., & Micsik, A. (2005). Using a metadata schema registry in the national digital data archive of hungary. In *Research and advanced technology for digital libraries* (pp. 314–322). Springer Berlin Heidelberg.

Gabriel, R. P., White, J. L., & Bobrow, D. G. (1991). CLOS: Integrating object-oriented and functional programming. *Commun. ACM*, *34*, 29–38. https://doi.org/10.1145/114669.114671

Gadamer, H.-G. (1975). Hermeneutics and social science. *Cultural Hermeneutics*, *2(4)*, 207–316.

Gangemi, A., Guarino, N., Masolo, C., Oltramari, A., & Schneider, L. (2002). Sweetening ontologies with DOLCE. *International Conference on Knowledge Engineering and Knowledge Management*, 166–181.

Gangemi, A., & Mika, P. (2003). Understanding the Semantic Web through Descriptions and Situations. In R. Meersman, Z. Tari, & D. C. Schmidt (Eds.), *On the move to meaningful internet systems 2003: CoopIS, DOA, and ODBASE* (pp. 689–706). Springer Berlin Heidelberg.

Garijo, D., Corcho, O., & Poveda-Villalón, M. (2021). FOOPS!: An ontology pitfall scanner for the fair principles. *Proceedings of the ISWC 2021 Posters, Demos and Industry Tracks: From Novel Ideas to Industrial Practice co-located with 20th International Semantic Web Conference (ISWC 2021)*.

Gartner, R. (2020). Towards an ontology-based iconography. *Digital Scholarship in the Humanities*, *35*(1), 43–53. https://doi.org/10.1093/llc/fqz009

Gayo, J. E. L., Prud'hommeaux, E., Boneva, I., & Kontokostas, D. (2018). *Validating RDF Data*. Springer International Publishing. https://doi.org/10.1007/978-3-031-79478-0

Ginzburg, C. (1986). Da A. Warburg a E. H. Gombrich. Note su un problema di metodo. In *Miti, emblemi, spie : Morfologia e storia* (pp. 29–106). Einaudi.

Glimm, B., Horrocks, I., Motik, B., Stoilos, G., & Wang, Z. (2014). HermiT: An OWL 2 Reasoner. *Journal of Automated Reasoning*, *53*(3), 245–269. https://doi.org/10.1007/s10817-014-9305-1

Gombrich, E. H. (1948). Icones symbolicae: The visual image in Neo-platonic thought. *J. Warburg Courtauld Inst.*, *11*(1), 163–192.

Gombrich, E. H. (1970). *Aby Warburg: An intellectual biography*. The Warburg Institute.

Gombrich, E. H. (1972). *Symbolic images*. Phaidon.

Greenwald, D. S. (2021). *Painting by numbers*. Princeton University Press.

Harpring, P., Lanzi, E., McRae, L., et al. (2006). *Cataloging Cultural Objects: A guide to describing cultural works and their images*. American Library Association.

Heidegger, M. (1962). *Being and Time*. SCM Press.

Heist, N., Hertling, S., Ringler, D., & Paulheim, H. (2020). Knowledge Graphs on the Web - an Overview. *Knowledge Graphs for eXplainable Artificial Intelligence*.

Hjorland, B. (2016). Subject (of documents). *ISKO Encyclopedia of Knowledge Organization*.

Houston, K. (2013). *An Introduction to Art Criticism: Histories, strategies, voices*. Pearson.

Hyvönen, E. (2020). Using the Semantic Web in digital humanities: Shift from data publishing to data-analysis and serendipitous knowledge discovery. *Semantic Web*, *11*(1), 187–193. https://doi.org/10.3233/SW-190386

Imdahl, M. (2012). Iconica. l'intuizione delle immagini. *Aisthesis. Pratiche, linguaggi e saperi dell'estetico*.

Impett, L. L. (2020). Analyzing Gesture in Digital Art History. In *The Routledge Companion to Digital Humanities and Art History* (pp. 386–407). Routledge. https://doi.org/10.4324/9780429505188-33

Irving Lavin. (1992). Iconography as a Humanistic Discipline. In *Iconography at the Crossroads. Papers form the colloquium sponsored by the Index of Christian Art. Princeton University 23-24 march 1990*. Princeton University.

Isaac, A., & Haslhofer, B. (2013). Europeana linked open data –data.europeana.eu. *Semant. Web*, *4*(3), 291–297.

Issa, S., Adekunle, O., Hamdi, F., Cherfi, S. S.-S., Dumontier, M., & Zaveri, A. (2021). Knowledge graph completeness: A systematic literature review. *IEEE Access*, *9*, 31322–31339.

Janev, V., Graux, D., Jabeen, H., & Sallinger, E. (2020). *Knowledge graphs and big data processing*. Springer Nature.

Ji, S., Pan, S., Cambria, E., Marttinen, P., & Yu, P. S. (2022). A survey on knowledge graphs: Representation, acquisition, and applications. *IEEE Transactions on Neural Networks and Learning Systems*, *33*(2), 494–514. https://doi.org/10.1109/TNNLS.2021.3070843

Kany, R. (1985). Lo Sguardo Filologico: Aby Warburg E I Dettagli. *Annali della Scuola Normale Superiore di Pisa. Classe di Lettere e Filosofia*, *15*(4), 1265–1283.

Kienle, M. (2017). Between Nodes and Edges: Possibilities and Limits of Network Analysis in Art History. *Visualizing Networks: Approaches to Network Analysis in Art History*, *6*(3), 4–22.

Kim, D., Son, S.-W., & Jeong, H. (2014). Large-Scale Quantitative Analysis of Painting Arts. *Scientific Reports*, *4*(1), 7370. https://doi.org/10.1038/srep07370

Klenczon, W., & Rygiel, P. (2014). Librarian Cornered by Images, or How to Index Visual Resources. *Cataloging & Classification Quarterly*, *52*(1), 42–61. https://doi.org/10.1080/01639374.2013.848123

Kubler, G. (2008). *The Shape of Time: Remarks on the History of Things*. Yale University Press.

Lancioni, T. (2019). Semiotica e iconologia. Nuovi temi per un vecchio confronto. *25*, 63–72.

Liepe, L. E. (2019). *The Locus of Meaning in Medieval Art: Iconography, Iconology, and Interpreting the Visual Imagery of the Middle Ages*. Medieval Institute Publications.

Lindsay, K. C. (1966). Art, Art History, and the Computer. *Computers and the Humanities*, *1*(2), 27–30.

Loetzsch, M., Wellens, P., De Beule, J., Bleys, J., & van Trijp, R. (2008). *The Babel2 Manual* (tech. rep. AI-Memo 01-08). AI-Lab VUB. Brussels.

Madhu, P., Kosti, R., Mührenberg, L., Bell, P., Maier, A., & Christlein, V. (2019). Recognizing Characters in Art History Using Deep Learning. *Proceedings of the 1st Workshop on Structuring and Understanding of Multimedia heritAge Contents*, 15–22. https://doi.org/10.1145/3347317.3357242

Madhu, P., Marquart, T., Kosti, R., Suckow, D., Bell, P., Maier, A., & Christlein, V. (2023). ICC++: Explainable feature learning for art history using image compositions. *Pattern Recognition*, *136*, 109–153. https://doi.org/10.1016/j.patcog.2022.109153

Manovich, L. (2015). Data Science and Digital Art History. *International Journal for Digital Art History*, (1). https://doi.org/10.11588/dah.2015.1.21631

Manovich, L. (2020). *Cultural Analytics*. MIT Press.

Marinescu, M.-C., Reshetnikov, A., & López, J. M. (2020). Improving object detection in paintings based on time contexts. *2020 International Conference on Data Mining Workshops (ICDMW)*, 926–932. https://doi.org/10.1109/ICDMW51313.2020.00133

Matthews, R. J. (1977). Describing and Interpreting a Work of Art. *The Journal of Aesthetics and Art Criticism*, *36*(1), 5–14. https://doi.org/10.2307/430744

McBride, B. (2004). The Resource Description Framework (RDF) and its Vocabulary Description Language RDFS. In *Handbook on ontologies* (pp. 51–65). Springer Berlin Heidelberg. https://doi.org/10.1007/978-3-540-24750-0_3

Milani, F., & Fraternali, P. (2021). A Dataset and a Convolutional Model for Iconography Classification in Paintings. *Journal on Computing and Cultural Heritage*, *14*(4), 1–18. https://doi.org/10.1145/3458885

Minsky, M. (1975). A framework for representing knowledge. In P. Winston (Ed.), *The Psychology of Computer Vision* (pp. 211–277). McGraw-Hill, New York.

Moretti, F. (2000). Conjectures on world literature. *New Left Review*, *1*.

Moxey, K. (1986). Panofsky's Concept of "Iconology" and the Problem of Interpretation in the History of Art. *New Literary History*, *17*(2), 265–274. https://doi.org/10.2307/468893

Müller, M. G. (2014). Iconography and iconology as a visual method and approach. In *The SAGE handbook of visual research methods* (pp. 283–297). SAGE Publications Ltd.

Nauta, G. J. (1993). HYPERICONICS: Hypertext and the social construction of information about the history of artistic notions. *Knowledge Organization*, *20*(1), 35–46. https://doi.org/10.5771/0943-7444-1993-1-35

Osman, I., Yahia, S. B., & Diallo, G. (2021). Ontology integration: Approaches and challenging issues. *Information Fusion*, *71*, 38–63.

Panofsky, E. (1955). *Meaning in the Visual Arts: Papers in and on Art History*. Anchor Books.

Panofsky, E. (1961). The iconography of Correggio's Camera di San Paolo. *Studies of the Warburg Institute*.

Panofsky, E. (1969). *Renaissance and Renascenses in Western art*. Harper & Row.

Panofsky, E. (1972). *Studies in iconology: Humanistic themes in the art of the Renaissance*. Westview Press.

Panofsky, E., & Saxl, F. (1933). Classical Mythology in Mediaeval Art. *Metrop. Mus. Stud.*, *4*(2), 228–280.

Paulheim, H. (2017). Knowledge graph refinement: A survey of approaches and evaluation methods. *Semantic Web*, *8*, 489–508.

Pellegrino, M. A., Scarano, V., & Spagnuolo, C. (2023). Move Cultural Heritage Knowledge Graphs in Everyone's Pocket. *Semantic Web*, *14*(2), 323–359.

Peroni, S. (2016). A simplified agile methodology for ontology development. *Proceedings of the 13th OWL: Experiences and Directions Workshop and 5th OWL reasoner evaluation workshop (OWLED-ORE 2016)*, 55–69. https://doi.org/10.6084/M9.FIGSHARE.3189769.V2

Peroni, S., Shotton, D., & Vitali, F. (2012). Scholarly Publishing and Linked Data: Describing Roles, Statuses, Temporal and Contextual Extents. *Proceedings of the 8th International Conference on Semantic Systems*, 9–16. https://doi.org/10.1145/2362499.2362502

Poveda-Villalón, M., Gómez-Pérez, A., & Suárez-Figueroa, M. C. (2014). OOPS! (OntOlogy Pitfall Scanner!): An on-line tool for ontology evaluation. *International Journal on Semantic Web and Information Systems (IJSWIS)*, *10*(2), 7–34.

Presutti, V., Daga, E., Gangemi, A., & Blomqvist, E. (2009). EXtreme Design with Content Ontology Design Patterns. *Proceedings of the 2009 International Conference on Ontology Patterns - Volume 516*, 83–97.

Previtali, G. (2009). Scritti di Erwin Panofsky. In *Studi di Iconologia* (pp. 349–370). Einaudi.

Preziosi, D. (2009). *The art of art history: A critical anthology*. Oxford University Press.

Rabinowitz, A., Shaw, R., & Golden, P. (2018). 16 making up for lost time: Digital epigraphy, chronology, and the periodo project. In A. D. Santis & I. Rossi (Eds.), *From practice to discipline* (pp. 202–215). De Gruyter Open Poland. https://doi.org/doi:10.1515/9783110607208-017

Radulovic, F., Mihindukulasooriya, N., García-Castro, R., & Gómez-Pérez, A. (2018). A comprehensive quality model for linked data. *Semantic Web*, *9*(1), 3–24.

Rampley, M. (2001). Iconology of the interval: Aby Warburg's legacy. *Word & Image*, *17*(4), 303–324. https://doi.org/10.1080/02666286.2001.10435723

Rantala, V. (1991). The work of art: Identity and interpretation [Num Pages: 271-292]. *Semiotica*, *87*(3-4), 271–292.

Raschka, S. (2018). Mlxtend: Providing machine learning and data science utilities and extensions to python's scientific computing stack. *Journal of Open Source Software*, *3*(24), 638. https://doi.org/10.21105/joss.00638

Rebele, T., Suchanek, F., Hoffart, J., Biega, J., Kuzey, E., & Weikum, G. (2016). YAGO: A multilingual knowledge base from wikipedia, wordnet, and geonames. *International semantic web conference*, 177–185.

Ringler, D., & Paulheim, H. (2017). One Knowledge Graph to Rule Them All? Analyzing the Differences Between DBpedia, YAGO, Wikidata & co. *Joint German/Austrian Conference on Artificial Intelligence (Künstliche Intelligenz)*, 366–372.

Ripa, C. (1764). *Iconologia del Cavaliere Cesare Ripa, Perugino* (Vol. 3). nella stamperia di Piergiovanni Costantini.

Riva, P., & Zumer, M. (2017). FRBRoo, the IFLA Library Reference Model, and now LRMoo: A circle of development.

Robert, S. (1996). *Artworks: Meaning, definition, value.* Penn State University Press.

Rossi Pinelli, O. (2019). *La storia delle storie dell'arte* (3rd ed.). Einaudi.

Rother, L., Koss, M., & Mariani, F. (2022). Taking Care of History: Toward a Politics of Provenance Linked Open Data in Museums.

Rother, L., Mariani, F., & Koss, M. (2023). Hidden Value: Provenance as a Source for Economic and Social History. *Jahrbuch für Wirtschaftsgeschichte / Economic History Yearbook*, *64*(1), 111–142. https://doi.org/10.1515/jbwg-2023-0005

Ruan, T., Li, Y., Wang, H., & Zhao, L. (2016). From queriability to informativity, assessing quality in use' of DBpedia and YAGO. *European Semantic Web Conference*, 52–68.

Sabbadini, R., Thilo, G., & HAGEN, E. (1888). Servii grammatici qui feruntur in vergilii carmina commentarii. *Rivista di Filologia e di Istruzione Classica*, *16*, 322.

Sartini, B., & Baroncini, S. (2023). A Comparative Study of Simple and Complex Art Interpretations in Linked Open Data Using ICON Ontology. *SWODCH'23: International Workshop on Semantic Web and Ontology Design for Cultural Heritage*.

Sartini, B., Baroncini, S., van Erp, M., Tomasi, F., & Gangemi, A. (2023). ICON: An Ontology for Comprehensive Artistic Interpretations. *Journal on Computing and Cultural Heritage*. https://doi.org/https://doi.org/10.1145/3594724

Sartini, B., & Gangemi, A. (2021). Towards the unchaining of symbolism from knowledge graphs: How symbolic relationships can link cultures. *Book of extended abstracts of the 10th national AIUCD conference. AIUCD, Pisa*, 576–580.

Sartini, B., van Erp, M., & Gangemi, A. (2021). Marriage is a peach and a chalice: Modelling cultural symbolism on the semantic web. *Proceedings of the 11th on Knowledge Capture Conference*, 201–208. https://doi.org/10.1145/3460210.3493552

Savnik, I., Nitta, K., Skrekovski, R., & Augsten, N. (2021). Statistics of knowledge graphs based on the conceptual schema. *CoRR*, *abs/2109.09391*.

Schich, M., Huemer, C., Adamczyk, P., Manovich, L., & Liu, Y.-Y. (2017). Network Dimensions in the Getty Provenance Index. https://doi.org/10.48550/arXiv.1706.02804

Schich, M., Song, C., Ahn, Y.-Y., Mirsky, A., Martino, M., Barabási, A.-L., & Helbing, D. (2014). A network framework of cultural history. *Science*, *345*(6196), 558–562. https://doi.org/10.1126/science.1240064

Schneider, P., Jones, J., Hiltmann, T., & Kauppinen, T. (2021). Challenge-derived design practices for a semantic gazetteer for Medieval and Early Modern places. *Semantic Web*, *12*(3), 493–515. https://doi.org/10.3233/SW-200394

Sciolla, G. (1997). Il dibattito metodologico sulla storia artistica 1985-1997. *Rivista storica italiana*, *109*(2), 518–540.

Shatford, S. (1986). Analyzing the subject of a picture: A theoretical approach. *Cataloging & Classification Quarterly*, *6*(3), 39–62. https://doi.org/10.1300/J104v06n03_04

Shenoy, K., Ilievski, F., Garijo, D., Schwabe, D., & Szekely, P. (2022). A study of the quality of Wikidata. *Journal of Web Semantics*, *72*, 100679.

Shotton, D. (2010). Cito, the citation typing ontology. *Journal of biomedical semantics*, *1*(1), 1–18.

Snydman, S., Sanderson, R., & Cramer, T. (2015). The International Image Interoperability Framework (IIIF): A community & technology approach for web-based images. *Archiving Conference*, *2015*(1), 16–21.

Staab, S., & Studer, R. (2010). *Handbook on Ontologies*. Springer Science & Business Media.

Stecker, R. (1994). Art Interpretation. *The Journal of Aesthetics and Art Criticism*, *52*(2), 193–206. https://doi.org/10.2307/431166

Steels, L. (1985). Design requirements for knowledge representation systems. In J. Laubsch (Ed.), *Gwai-84* (pp. 1–19). Springer Berlin Heidelberg.

Steels, L., & Loetzsch, M. (2010). Babel: A tool for running experiments on the evolution of language. In S. Nolfi & M. Mirolli (Eds.), *Evolution of communication and language in embodied agents* (pp. 307–313). Springer.

Stork, D. G., Bourached, A., Cann, G. H., & Griffiths, R.-R. (2021). Computational identification of significant actors in paintings through symbols and attributes. https://doi.org/10.48550/arXiv.2102.02732

Suárez-Figueroa, M. C., Gómez-Pérez, A., & Villazón-Terrazas, B. (2009). How to write and use the ontology requirements specification document. In *On the move to meaningful internet systems: OTM 2009* (pp. 966–982). Springer Berlin Heidelberg.

Szabo, V. (2012). Transforming Art History Research with Database Analytics: Visualizing Art Markets. *Art Documentation: Journal of the Art Libraries Society of North America*, *31*(2), 158–175. https://doi.org/10.1086/668109

Taylor, P. (2008). *Iconography without texts: [colloquium, Warburg Institute, 3-4 june 2005]*. Warburg Institute ; N. Aragno.

Thomson, E. (t.) (2014). Aby M. Warburg by Gertrud Bing (first edition "Rivista storica italiana" LXXII, 1960). *La Rivista di Engramma*, *2014*(116). https://doi.org/10.25432/1826-901X/2014.116.0001

Ufer, N., Simon, M., Lang, S., & Ommer, B. (2021). Large-scale interactive retrieval in art collections using multi-style feature aggregation. *PLOS ONE*, *16*(11), e0259718. https://doi.org/10.1371/journal.pone.0259718

Uschold, M., & Grüninger, M. (1996). Ontologies: Principles, methods and applications. *The Knowledge Engineering Review*, *11*.

Van Den Heuvel, C. (2024). "Beeldleer"- Iconology: Publishing Henri van de Waal's forgotten classification for the history of the arts [under review]. *Oud Holland – Journal for Art of the Low Countries*.

Van den Heuvel, C., & Zamborlini, V. (2021). Modeling and Visualizing Storylines of Historical Interactions. Kubler's Shape of Time and Rembrandt's Night Watch. *Linking Knowledge*, 99–141.

Van Ruymbeke, M., Hallot, P., & Billen, R. (2017). Enhancing CIDOC-CRM and compatible models with the concept of multiple interpretation. *ISPRS Ann. Photogramm. Remote Sens. Spat. Inf. Sci.*, *IV-2/W2*, 287–294.

van Leeuwen, T. (2004). Semiotics and Iconography. In *The Handbook of Visual Analysis* (pp. 92–118). SAGE Publications Ltd. https://doi.org/10.4135/9780857020062.n5

van Straten, R. (2012). *An introduction to iconography* (2nd ed.). Routledge.

Vrandečić, D., & Krötzsch, M. (2014). Wikidata: A free collaborative knowledgebase. *Commun. ACM*, *57*(10), 78–85. https://doi.org/10.1145/2629489

Waal, H. v. d. (1952). *Drie eeuwen vaderlandsche geschied-uitbeelding, 1500-1800: Een iconologische studie*. Nijhoff.

Wang, R. Y., & Strong, D. M. (1996). Beyond accuracy: What data quality means to data consumers. *Journal of management information systems*, *12*(4), 5–33.

Warburg, A. (1999). *The renewal of pagan antiquity: Contributions to the cultural history of the European Renaissance*. Getty Research Institute for the History of Art; the Humanities.

Warburg, A., Ohrt, R., Heil, A., Scherer, B. M., Sherman, B., & Wedepohl, C. (2020). *Aby Warburg: Bilderatlas Mnemosyne : The original*. HKW, Haus der Kulturen der Welt ; The Warburg Institute ; Hatje Cantz.

Wilkinson, M. D., Dumontier, M., Aalbersberg, I. J., Appleton, G., Axton, M., Baak, A., Blomberg, N., Boiten, J.-W., da Silva Santos, L. B., Bourne, P. E., Bouwman, J., Brookes, A. J., Clark, T., Crosas, M., Dillo, I., Dumon, O., Edmunds, S., Evelo, C. T., Finkers, R., … Mons, B. (2016). The

FAIR Guiding Principles for scientific data management and stewardship. *Scientific Data*, *3*, 160018. https://doi.org/10.1038/sdata.2016.18

Wind, E. (1937). Comments on an Observation by Reynolds. *Journal of the Warburg Institute*, *1*(1), 70–71. https://doi.org/10.2307/750075

Wittkower, R. (1977). *Allegory and the migration of symbols*. Thames; Hudson.

Wollheim, R., & Eldridge, R. T. (2015). *Art and its objects: With six supplementary essays* (Second edition, Cambridge Philosophy Classics edition). Cambridge University Press.

Zaveri, A., Rula, A., Maurino, A., Pietrobon, R., Lehmann, J., & Auer, S. (2016). Quality Assessment for Linked Data: A Survey. *Semantic Web*, *7*(1), 63–93.

Zeng, M. L., Žumer, M., & Salaba, A. (2009). Functional Requirements for Subject Authority Data (FRSAD): A Conceptual Model. In M. J. Bates & M. N. Maack (Eds.), *Encyclopedia of Library and Information Sciences, Third Edition* (0th ed., pp. 1–16). CRC Press. https://doi.org/10.1081/E-ELIS3-120049494

Žumer, M., Zeng, M. L., & Salaba, A. (2012). *FRSAD): Conceptual Modeling of Aboutness*. ABC-CLIO.