



ICON: An Ontology for Comprehensive Artistic Interpretations

BRUNO SARTINI and SOFIA BARONCINI, University of Bologna
MARIEKE VAN ERP, KNAW Humanities Cluster DHLab
FRANCESCA TOMASI and ALDO GANGEMI, University of Bologna

In this work, we introduce ICON, an ontology that models artistic interpretations of artworks' subject matter (i.e., iconographies) and meanings (i.e., symbols, iconological aspects). Developed by conceptualizing authoritative knowledge and notions taken from Panofsky's levels of interpretation theory, ICON ontology focuses on the granularity of interpretations. It can be used to describe an interpretation of an artwork from the pre-iconographical, iconographical, and iconological levels. Its main classes have been aligned to ontologies that come from the domains of cultural descriptions (ArCo, CIDOC-CRM, VIR), semiotics (DOLCE), bibliometrics (CITO), and symbolism (Simulation Ontology), to grant a robust schema that can be extendable using additional classes and properties coming from these ontologies. The ontology was evaluated through competency questions that range from simple recognition on a specific level of interpretation to complex scenarios. Data written using this model was compared to state-of-the-art ontologies and schemas to both highlight the current lack of a domain-specific ontology on art interpretation and show how our work fills some of the current gaps. The ontology is openly available and compliant with FAIR principles. With our ontology, we hope to encourage digital art historians working for cultural institutions in making more detailed linked open data about the content of their artifacts, to exploit the full potential of Semantic Web in linking artworks through not only subjects and common metadata but also specific symbolic interpretations, intrinsic meanings, and the motifs through which their subjects are represented. Additionally, by basing our work on theories made by different art history scholars in the last century, we make sure that their knowledge and studies will not be lost in the transition to the digital, linked open data era.

CCS Concepts: • **Information systems** → **Semantic web description languages**; • **Computing methodologies** → **Ontology engineering**;

Additional Key Words and Phrases: Iconology, iconography, art interpretation, ontology, cultural heritage, Semantic Web

ACM Reference format:

Bruno Sartini, Sofia Baroncini, Marieke van Erp, Francesca Tomasi, and Aldo Gangemi. 2023. ICON: An Ontology for Comprehensive Artistic Interpretations. *ACM J. Comput. Cult. Herit.* 16, 3, Article 59 (August 2023), 38 pages.
<https://doi.org/10.1145/3594724>

B. Sartini and S. Baroncini contributed equally to this research. S. Baroncini worked on Sections 1, 2, 3.2, 4, 5.5.3, 5.5.4, 5.5.7, 5.5.8, 6, 6.1, 6.2, 6.3. B. Sartini worked on Sections 3, 3.1, 3.3, 3.4, 5, 5.1, 5.2, 5.3, 5.4, 5.5, 5.5.1, 5.5.5, 5.5.6, 5.6, 6.4, 7. A. Gangemi worked on section 5.5.2. All authors contributed to Section 8.

This work was partially funded by the Emilia Romagna Region (grant agreement 462 25/03/2019), the University of Bologna, the SPICE EU H2020 Project 870811 within the program SOCIETAL CHALLENGES—Europe in a Changing World—Inclusive, Innovative and Reflective Societies, and the Odeuropa EU H2020 Project under grant agreement number 101004469.

Authors' addresses: B. Sartini, S. Baroncini, F. Tomasi, and A. Gangemi, University of Bologna, Bologna, Italy; emails: bruno.sartini3@unibo.it, sofia.baroncini4@unibo.it, francesca.tomasi@unibo.it, aldo.gangemi@unibo.it; M. van Erp, KNAW Humanities Cluster DHLab, Amsterdam, The Netherlands; email: marieke.van.erp@dh.huc.knaw.nl.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

© 2023 Copyright held by the owner/author(s). Publication rights licensed to ACM.

1556-4673/2023/08-ART59 \$15.00

<https://doi.org/10.1145/3594724>

1 INTRODUCTION

Distinguishing between what can be considered an artwork from what is not, and reach a precise definition of art itself, can be challenging in a dynamic world in which new forms of art are constantly introduced [1, 47]. For this work, we refer to artwork(s) as a “visual object or experience consciously created through an expression of skill or imagination” [7]. Since an artwork, for its nature, cannot usually be completely understood only from its objective characteristics, it is subjective to observers’ interpretations. In this work, we present a new ontology, ICON, which models art interpretations of the artworks’ subject matter and its possible meanings. In the context of art comprehension, an interpretation is intended as “any kind of assignment of meaning or significance to artworks” [47, p. 113]. Through interpretations, art historians can claim different kinds of explanations about the artwork concerning different aspects such as the artwork’s content, the tendency of the art of a century, or the reception of artworks by the public [31, p. 114]. Among them, the interpretations considering the comprehension of the subject matter of an artwork fall under the domain of iconographic interpretation, which is based only on internal aspects of the artwork (e.g., a child with bows, arrows, and wings depicted in an artwork is recognized as Cupid and as a symbol of love) [31, p. 124]. The interpretation conducted on this basis can be further enriched by other types of interpretations that take as evidence an external source—that is, the cultural context [31, p. 131]. In this sense, artworks are read as symptoms of the contemporary culture [37]. For example, the fact that, during the Middle Ages, the classical deities were represented deprived of their classical form can be read as the incapability of the Medieval artists and society of retaining a classical model with its appearance, since it was too far from their taste and from the new Gothic representational conventions [38].

Therefore, the domain of knowledge of artworks content interpretation is complex and characterized by the subjectivity of the author of each claim. For these characteristics, we believe that Semantic Web technologies, with a focus on ontologies, are a suitable tool to conceptualize this semantic expressivity by means of the high level of granularity and flexibility offered. Nevertheless, the study in our earlier work [4] highlights how current knowledge graphs, in the Semantic Web, do not express iconographic and iconological statements with the correct granularity, suggesting the introduction of domain-specific ontologies that conceptualize these aspects.

Since the domains of iconography and iconology described previously concern the description and comprehension of artworks’ content, we model art interpretations according to them. Therefore, the ontological modeling proposed here aims to answer to the following research questions:

RQ1: To what extent is it possible to model the domain of knowledge of iconology and iconography, to provide art historians and cultural institutions a way for expressing complex art subjects and meanings, the interlinking among them, and claims about their interpretations?

RQ2: Can the newly developed model outperform current work in terms of granularity?

Among the approaches adopted by the scholars in the context of the artworks’ content interpretation, the iconographical-iconological method¹ formalized by Erwin Panofsky had, in the last century, the greatest relevance and influence on contemporary art historians [43]. Panofsky followed the approach first adopted by Warburg [61] in his studies, which was aimed at understanding the artistic subjects and motifs as witnesses of socio-cultural phenomena. Whereas today his studies are fundamental for the approach itself, the prevailing perspective is the one formulated by Panofsky in a three-layered framework of the artwork’s understanding [34]. On that basis, some scholars proposed variations over levels subdivision [3, 8], sometimes including other aspects, such as the artist’s psychology [62] or the iconic language of the image [32]. Although the framework is recognized as a valid method or approach, it is commonly accepted to acknowledge that this kind of interpretation is subjective and intuitive when practically applied [34].

¹Which inspired the name of this ontology.

For its complete formalization and historical relevance, we will refer to Panofsky’s approach as a representative method of the discipline,² yet considering the enrichment given by other scholars mentioned in Section 2. In detail, Panofsky subdivides the act of interpretation in three levels, the first two of which fall under the traditional domain of iconography—that is, the identification of iconographies (level 2) attributes and variants with which they can be represented (level 1), whereas the last level concerns the socio-cultural interpretation of artworks content, closer to Warburg’s iconological approach [59].

We consider Panofsky’s theoretical approach as suitable for modeling because it is, from the authors’ point of view, the most complete attempt in the literature to formalize the discipline in detail. Indeed, it not only defines the mechanisms of interpretation and its components but also gives precise indications on (i) the types of subjects, and (ii) how the subjects, their components, and meanings are related.

In this work, we consider the concept of interpretation according to the presented theory—that is, an observer interpreting what is represented by one or more artworks, their possible iconography, and meaning. Therefore, other types of interpretations, such as the results of the observation of the physical object (e.g., measurement) and its metadata definitions (e.g., dating, author, and title attribution) are not considered in the scope of this work. The concept of subjectivity is limited to the described situation. As the content and meaning interpretation always depend on the viewer perception and background knowledge, multiple, incompatible interpretations may derive from different observations of the same artwork. For the sake of clarity, we will use the term *interpretation* to refer to the overall claim made by an observer in which the artwork’s subject matter and meaning are understood (e.g., “this painting depicting Venus expresses a wish of good marriage”), whereas the term *recognition* will refer to the identification of subjects taking place at the different levels (e.g., the recognition that a child is Cupid). Therefore, an interpretation is composed of a set of recognitions contributing to the overall understanding of the artwork.

The article is structured as follows. Section 2 introduces the theoretical background of art interpretation, on which the ontology is based. In Section 3, we analyze the state of the art of ontologies and general domain schemas that deal with interpretations, relevant concepts like symbolism, or description of cultural heritage objects. Then, Section 4 explains the requirements that were used to design the ontology, along with potential users and lexical usage. Section 5 describes our design process of the ontology, describing in detail all the design iteration that have been undertaken to model different parts of our work, together with the axiomatization details, the alignments, and reuse of existing ontologies. In Section 6, the evaluation process of the ontology is explained, with both automatic evaluations provided by relevant tools and quality-based evaluation over the granularity potential of the model. Section 7 shortly deals with the release of the ontology and the publication of its documentation. Finally, Section 8 concludes the article with a final discussion about the impact of the ontology, its current limitations, and future work.

2 THEORETICAL CONTRIBUTION

In this section, we illustrate Panofsky’s theory of art interpretation [36, 38] introducing the theoretical aspects that were fundamental for the ontology design phase. During the act of interpretation of an artwork, the formal aspects, such as forms, colors, and compositions, are perceived. When these formal aspects are interpreted as precise objects, the sphere of meanings is considered. According to Panofsky, there are different types of meaning that can be interpreted in an artwork, subdivided in three layers. The depth to which the artwork can be understood depends on the background knowledge of the observer: the more he has knowledge about the artist, stylistic conventions, and cultural context of him/her/them, the more the interpretation at each level is correct, including more profound insights on cultural meanings.

²It has to be specified that albeit the relevance of the theory is affirmed, it was mainly developed to be used in Panofsky’s area of interest (i.e., the Western Middle Ages and Renaissance Art).

Table 1. Levels of Interpretation According to Panofsky [36]

Level	Type	Subject Identified	Recognized Elements	Example: Reading of Leonardo's <i>Ultima Cena</i>	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	13 people, table, food, dishes (all factual meanings), act of talking (expressional meaning)	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	The last Supper, Jesus, Apostles	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	Manifestation of Leonardo's and Renaissance particular attitude	Familiarity with cultural phenomena, tendencies, attitudes

The first layer, namely the pre-iconographical description, requires the knowledge of the representational conventions to allow a correct recognition of factual (e.g., objects, people, actions) and emotional meanings, namely primary or natural subjects. In detail, this description is achieved by the recognition of pure forms (i.e., combinations of forms and colors) as carriers of primary subjects. Pure forms such recognized are called *artistic motifs*, and their combinations are *compositions*. An enumeration of the recognition of artistic motifs constitutes a pre-iconographical interpretation of the artwork [36, p. 28].

If the observer is familiar with the literary sources known by the artist, then the subjects already identified at level 1, namely 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). The artistic motifs such recognized are called *images* or *invenzioni*, namely 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 [36, p. 29, note 1].

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

The scholar highlights that the first two levels are a description of facts and are under the domain of iconography, whereas the last level is in the domain of iconology, which is a synthetic intuition rather than a description. Table 1 resumes the synoptic table from the work of Panofsky [36, pp. 40–41], integrating it with further explanation of concepts implemented in the ontology modeling and by adding a practical example.

Although some following scholars made some variations of the model,³ the subdivision of the interpretation in levels is generally accepted. In detail, we highlight that some scholars put the attention on relevant aspects that we considered during the modeling. Van Straten [59] highlights the difference between intentional and unintentional meanings by dividing the third level in two layers. In this way, he recognizes that some more profound meanings are voluntarily expressed by the artist (e.g., the concept of “good wishes” that the artist wants to express in an artwork made for a wedding occasion) and more unconscious, cultural meanings. Another relevant addition is made by Imdahl [32]. He underlines that the iconic sense of the image should not be ignored, since it is the primary means through which visual arts communicate. For example, the disposition of figures in the space can provide insights on their relationships, their actions, or in expressional meanings.

³We refer to previous work [3] for a further comparison between the major theories.

Furthermore, the preliminary studies conducted in previous work [3], which considered approximately 50 articles of the major scholars of iconography and iconology, collected in several other works [29, 32, 37, 38, 59, 61, 62], highlight important features that may be involved in an iconographical-iconological interpretation that should not be ignored. Indeed, from the bottom-up analysis emerged that the following aspects may be relevant for the supporting of the third-level meaning, namely (i) the direct citation of visual patterns from other artworks, (ii) the dependency of certain iconographies from specific sources, (iii) the role of style, (iv) the fact that a cultural meaning generally involves more than one artwork, and (v) the fact that scholars often extend claims by other scholars.

3 STATE OF THE ART

Interpretations in the Semantic Web are a widely discussed topic [3, 16]. In this section, we will analyze (i) related work that covers specific aspects of cultural heritage interpretations and possibly iconographical-iconological content in the form of ontologies,⁴ (ii) how these and related ontologies model the concept of interpretation, (iii) iconographic and iconological elements contained in general domain schemas,⁵ and (iv) existing controlled vocabularies and taxonomies designed to classify iconographical and iconological elements, intended as authoritative sources of knowledge that provide permanent URIs for potential subjects, art styles, and other relevant information in the context of art interpretation. Following the reusability principles of the Semantic Web [52], we reuse parts of the models listed in the following in our ontology by making alignments between our classes/properties and theirs, or by directly reusing parts of their schema. The alignments are described in Section 5.5.

3.1 Ontologies Related to Iconography and Iconology

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. CIDOC-CRM [5] is a widely used ontology in the context of cultural heritage. It has an event-based structure and covers fundamental aspects of the lifecycle of a cultural heritage object. Carboni and de Luca [8] extended it with the VIR ontology, which explores the concept of visual representations in artworks and associates the portion (called the *iconographical atom*) of the cultural heritage object to the recognized subject. We use SKOS alignments to refer to parts of CIDOC and VIR in our ontology, and compare the coverage of our ontology and VIR in Section 6.3. Compared to our ontology, VIR focuses only on subjects of level 2, considering iconographies and their attributes, consequently lacking of a clear distinction between levels. The preliminary study conducted in previous work [3] further extends VIR by the addition of an iconological interpretation class linking to the artwork concepts and external cultural phenomena. It is evaluated over 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 over iconographical and iconological theories, along with the real-base evaluation, it is used as a source for ontology development proposed here, 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 2 subjects and the integration of multiple interpretations by different art historians. Gartner [27] proposes an ontology to facilitate and automate the identification of subjects (level 2) in works of arts through logical inferences. No alignments were possible to this ontology because it has not been released. ArCo's ontology [10] was developed to model Italian cultural heritage artifacts by converting information contained in traditional catalogue sheets into linked open data. Among the possible aspects

⁴We consider here only those ontologies specifically designed to deal with cultural heritage.

⁵By general schemas, we mean data models expressed through an ontology that were not designed with the scope of describing only cultural heritage but still contain relevant aspects of our work. The distinction between these and the previous ones lies only in the purpose of the ontologies in question.

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`,⁶ not exploiting the full potential of the Semantic Web [49]. As we mention in Sections 4 and 5, in the development of ICON we designed specific classes and properties to express this information with the necessary granularity. We reuse some parts of ArCo to refer to the concept of artwork and subject. In an earlier work [50], symbolic meanings in cultural heritage in the Simulation Ontology were modeled. Compared to ICON, this ontology does not consider the hermeneutic act of interpretation of associating the symbolic meanings to artworks. Nevertheless, 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 takes place fits well with our modeling of symbols. Therefore, we reuse the Simulation Ontology classes and properties to express the symbolic meanings in our work, inserting it in the context of an interpretation of an artwork.

3.2 Ontological Modeling of 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.⁷ 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 formalizing 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 that is done directly or through measurements. It represents the “transition between reality and propositions” [20]. 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 addressed also by the history domain to represent the frequent case of disagreeing historians’ interpretations of the same events. As reported by Frank [22], several ontologies afford the theme by modeling different views of the same observed events, such as SEM ontology⁸ MIDM [58] and the **Ontology Design Pattern (ODP) Event-Model-F**,⁹ expanding DOLCE+DnS Ultralite (DUL), in which a distinction between facts and interpretations is already stated [24]. HiCO [18]¹⁰ goes further by adding contextual information to the interpretation, such as interpretation type and criterion.

The concept of interpretation used in our ontology reflects the VIR perspective, narrowing down the field to those interpretations that an observer may have about the visual representation of an artwork, excluding scientific observations about its physical features. 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. [58] relatively to the archaeological field, stating that the views produced by scholars are “the result of an interpretive reasoning

⁶See <https://dati.beniculturali.it/lodview-arco/resource/HistoricOrArtisticProperty/0500653281.html> (`dc:description` value).

⁷<https://dati.beniculturali.it/lodview-arco-onto/ontology/context-description/Interpretation.html>.

⁸<https://semanticweb.cs.vu.nl/2009/11/sem/>.

⁹http://ontologydesignpatterns.org/wiki/Ontology:Event_Model_F.

¹⁰<http://purl.org/emmedi/hico>.

that includes the subjectivity of the author,” due to the uncertainty and incompleteness that often characterizes archaeological data [58].

While modeling an interpretation, it is important to define also what the concept of *meaning* is. Considering semiotics, there are several aspects of meaning that can be modeled [40]. In the context of this study, *meaning* is defined according to the semiotic theory [21] in which a signifier (i.e., an icon) signifies the carried signification (i.e., a meaning) [54, pp. 93–94]. Therefore, following Panofsky’s modeling, all the subjects identified at each level of interpretation are considered meanings [36]. This definition corresponds to the notion of meaning as a social object introduced in the work of Picca et al. [40].

3.3 General Schemas Containing Iconographical and Iconological Elements

Although iconographical and iconological descriptions are not their main focus, some general domain schemas contain information related to art interpretation. Often these schemas rely on the sole subject property (e.g., `dc:subject`, `schema:about`) to describe any information about the iconographical and iconological content [2, 15, 33, 44]. One of the outliers in this characteristic is Wikidata and its Wikidata schema [60]. The Wikidata schema contains different properties that link an artwork to its content, such as `wdt:P921` (main subject) or `wdt:P180` (depicts). Moreover, Wikidata allows adding qualifiers to the statements made with the property `wdt:P180` to address specific aspects of elements depicted in the artworks (i.e., their symbolic meaning, qualities).¹¹ Compared to ICON, Wikidata does not explicitly distinguish between levels of interpretations, linking elements of the first and second level to the artwork with the same property. This limits the possibility of art historian driven research questions to be answered with the Wikidata model. Finally, compared to our work, Wikidata does not include potential intrinsic meaning of artworks provided by the third level of interpretation. A more in-depth analysis of the qualifiers and comparison between Wikidata and ICON is presented in Section 6.3.

The scope of iconography has some overlaps with the domain of narratology in the description of the plot for what concerns the represented characters and their actions. A good wealth of studies concerns the semantic modeling of the topic¹² [12, 56], among which some focus on the narrative representation in visual images [14, 63]. Even though they do not describe the iconographical subject with the necessary granularity required by an iconographical study, they provide a solution for organizing the common archetypical knowledge of stories, events, and characters participating in them [14].

3.4 Controlled Vocabularies for Iconography

Controlled vocabularies and taxonomies of art and culture, despite not being ontologies, are essential for standardizing the reference to elements that belong to the first, second, or third level of interpretation. Iconclass [13] is a classification system that mostly deals with iconographical subjects. The Getty **Art & Architecture Thesaurus (AAT)** [30] provides permanent identifiers for people, concepts, and places that might be contained in artworks. The two aforementioned taxonomies cover a wide amount of information,¹³ but other exists that were created ad hoc for museums or cultural institutions, such as the Rijksmuseum’s thesaurus [19].

These controlled vocabularies and taxonomies can foster the interoperability between different knowledge graphs that use them, but they do not provide statements regarding the type of element that is depicted in a work of art (whether it belongs to the first, second, or third level of interpretation). For this reason, it is essential for these characteristics to be modeled with a specific ontology.

¹¹<https://www.wikidata.org/wiki/Property:P4878>.

¹²For the modeling of fictional entities from a philosophical perspective, see <https://plato.stanford.edu/entries/fictional-entities/>.

¹³Although Iconclass admittedly is based on Eurocentric subjects.

Table 2. Description of Requirements 1 to 5 According to the 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 artworks interpretations online and integrate them with existing data so as to enhance the query potentiality of the cultural institutions' data Use 2. Conduct specific and detailed quantitative analysis to answer research questions in the domain research field Use 3. Provide a semantic structure for knowledge extraction

4 REQUIREMENTS

Based on the iconographical and iconological literature analysis described in Section 2, we formulated the requirements using the SEEMP framework [55]. The terminology was mainly selected from Panofsky's theory [36]. The output document is described in Tables 2 and 3.

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 at being 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 requirement¹⁴ is the reuse and alignment to the standards shared across the community to allow of reusability. Furthermore, the **Competency Questions (CQs)** formulated for the functional requirements aim at expressing the various aspects of the iconographical-iconological approach described in Section 2. We summarize the main themes that can be extracted from the requirements listed in Table 3 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, namely wings and arrows [38]) must be described.
- (3) The symbolic and cultural meanings attributed to each subject must be included.

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

- (4) The attribution needs to be subjective.
- (5) The sources used by the scholar to state its claim need to be present.

¹⁴See slot 6a of the SEEMP ORSD in the work of Suárez-Figueroa et al. [55].

- (6) The clear distinction between the subjects described at a general level (i.e., the background knowledge necessary for iconographical descriptions cited in Table 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 to gather 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), which 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,” and “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 style should be included.

As Panofsky’s theory is considered a representative formalization of the iconological approach, we take the majority of the ontology’s 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 3, point 7, by extracting the terms that are answering to CQs directly from the definition of his theory. The number following each word indicates its frequency in the selected article,¹⁵ in which Panofsky’s theory is fully illustrated.

5 ONTOLOGY DESIGN

The ICON ontology¹⁶ was designed following the SAMOD [39] and eXtreme Design [42] methodologies. SAMOD is an agile methodology that focuses on the application of small iterative steps to model parts of an ontology. Each step is individually documented and combines motivating scenarios that derive from general domain descriptions with data-centric examples of descriptions formalized with the ontology. We reuse the SAMOD methodology for the main part of the design, as we adopt the iteration-like structure and its outputs. In fact, the design process was divided into four 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 CQs formulated both in natural language and SPARQL (referring to the aligned prototype), and a Jupyter notebook that contains unity tests. All the CQs were tested on real interpretations by Panofsky [37] expressed using the ontology schema. For a more detailed description of the test dataset, see Section 6.1. eXtreme Design is another agile methodology that divides the development of an ontology through iterations but focuses on the reuse of ODPs. 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 reuse of modeling patterns that come from the “solution space,” such as the ODP. We specifically adopted this methodology when dealing with the reuse of ODPs that were specialized in the context of our domain. The following paragraphs describe (i) each SAMOD iteration, (ii) the specialization of ODP to facilitate some modeling issues, and (iii) the refactoring of some classes and property through alignment to relevant ontologies.

¹⁵For this analysis, we referred to the article “Iconography and Iconology: An Introduction to the Study of Renaissance Art” by Panofsky [36], since it is the last published revised version.

¹⁶The ontology is available at <https://w3id.org/icon/ontology/>.

Table 3. Description of Requirements 6 and 7 According to the 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 the characters, events, personifications, named objects, and places recognized at level 2.
	CQ 2.3 In which story or allegory are involved the depicted subjects?
	CQ 2.4 Do the level 2 subjects have a symbolic meaning?
	CQ 2.5 Which is the object that allows the 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 supporting 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 artworks are interpreted on an iconological level but not on an iconographic one?
	CQ 0.5 What are the recognitions supporting another one? Of which 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; Pre-iconographical Description 9; Iconographical Analysis, 8; Composition, 7; Expressional, 7; Artistic Motifs, 5; Factual, 5; Iconological Interpretation, 4; Iconology, 4; Invenzioni, 1

5.1 First Design Iteration: Recognitions

As explained in Section 2, works of art can be analyzed through different layers of interpretations that depend on recognitions. A recognition, 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.¹⁷ Coherent recognitions on the same artwork are collected and

¹⁷More observation on the matter can be found in Section 5.5.

documented by interpretation descriptions (requirement 8, Section 4).¹⁸ 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 [23]) links the recognition to the agent who performed it (requirement 4, Section 4). The class `InterpretationDescription` is linked to (one or many) `Recognition` 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,¹⁹ `isCompliantWithIconographicalRecognition` for iconographical recognitions, and `isCompliantWithIconologicalRecognition` for iconological recognitions. The CiTO [51] 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 4). 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 Pisa cathedral is represented with the same appearance of the nude classical iconography of Venus Pudica (formal motif recognition).²⁰ This interpretation provides support to the third-level recognition of the characteristics of the Proto-Renaissance movement in the cultural context of the Medieval Tuscany [37, p. 157]. To express this using our ontology, the property `cito:givesSupportTo` can link the supporting recognition to another one (requirement 7, Section 4). These elements are also the object of interest of the general CQs (see Table 3, Q0.1–Q0.5).

Depending on the level of interpretation presented in Table 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, and (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 [36]).

These classes and their specific usage will be further described in the following sections. Figure 1 shows a rendering of the classes and properties of this iteration.

5.2 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 (tree of life, Saint Joseph, Excalibur) that are

¹⁸The 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 description.

¹⁹Both 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.

²⁰Figure available at https://www.wga.hu/art/p/pisano/giovanni/pulpit_p/3column5a1.jpg.

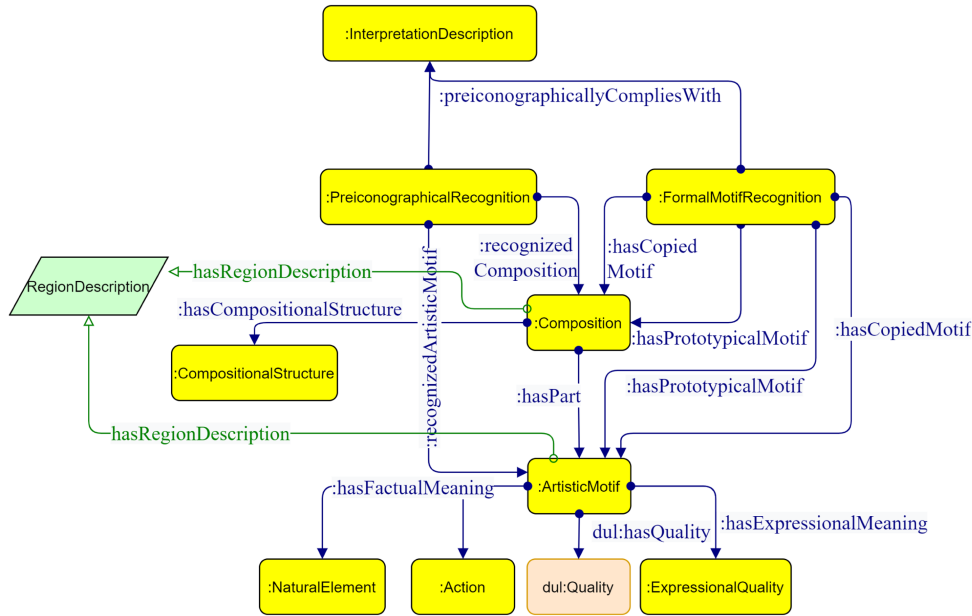


Fig. 2. ICON ontology classes and properties describing the pre-iconographical level of interpretation (level 1).

classes as the domain. As mentioned earlier, the use of IIRF 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 2 shows a graphical rendering of the classes and properties used in this interpretation level.

5.3 Third Design Iteration: Iconographical Recognitions (Level 2)

In this third iteration, we focus on the Panofsky’s second level of art interpretation: the iconographical interpretation. In this level, the interpreter recognizes images and invenzioni²³ 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-level subjects, which are characters, places, events, named objects,²⁴ symbols, and 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.²⁵ For example, in an artwork, you might recognize three images: the first refers to the general subject of Mary, the second refers to the general subject of Angel Gabriel, and the third refers

²³Invenzione is an Italian word used by Panofsky as an umbrella term for allegories and stories [36, 59].

²⁴A named object is a non-living unique element that is often used as an attribute for the recognition of specific characters (Thor’s hammer).

²⁵This definition slightly differs from the Panofsky’s one: whereas 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” [36]. 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

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 the possibility to express symbols as just symbolic meanings recognized, or, for a more thorough description, as *Simulations* (see Section 5.5). 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* (i.e., 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 [12, 14]. 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 3 shows the classes and properties relative to this level of recognition.

5.4 Fourth Design Iteration: Iconological Recognitions (Level 3)

Iconological interpretations (third level) focus on the recognitions of intrinsic meanings.²⁶ An intrinsic meaning links the whole artwork or some parts of it to a cultural phenomenon and a concept that defines it. The *IconologicalRecognition* class is linked to the *IntrinsicMeaning* class²⁷ through the property *recognizedIntrinsicMeaning*. 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*, and *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.²⁸ 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 four subclasses, which specify the type of cultural phenomenon, namely *Attitude*, *Belief*, *CulturalValue*, and *Tendency*. These terms are taken from Panofsky’s vocabulary in 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.

²⁶Even 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.

²⁷Compared 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).

²⁸The concepts, ideas, and abstract elements that are linked to intrinsic meanings on an iconological level are very broad [37]. Therefore, we decided to reuse this DOLCE class (*Social Object*), which conceptualizes a broad set of possible entities [23].

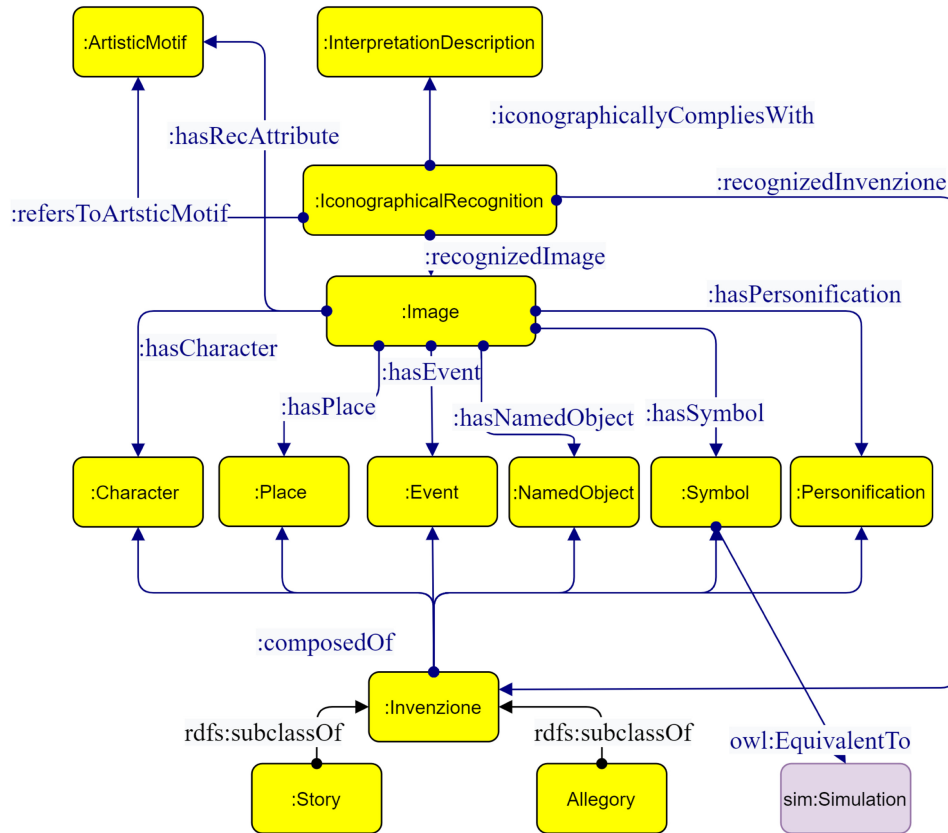


Fig. 3. ICON ontology classes and properties describing the iconographical level of interpretation (level 2).

description of the third level of artistic interpretation.²⁹ 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 4.

5.5 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 [9, 35]. Our ontology selection for reuse and alignment was guided by different principles: (i) standardization for CIDOC-CRM [5] and FRBRoo [46] because they are considered standard frameworks in the domain, (ii) cognitive and formal analysis for the choice of DOLCE foundational ontology [6, 23] in its OWL version (DOLCE Zero), Simulation Ontology [50], VIR [8], HiCO [18], and CiTO [51], as all of them offer design solutions to the CQs defined from the requirements in Section 4.

²⁹Although these subclasses could be formally associated with mental entities just as recognitions, they differ in their function. Recognition are modeled 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 will be discussed in Section 8, further work will be dedicated to a more thorough description of cultural phenomena and their subclasses.

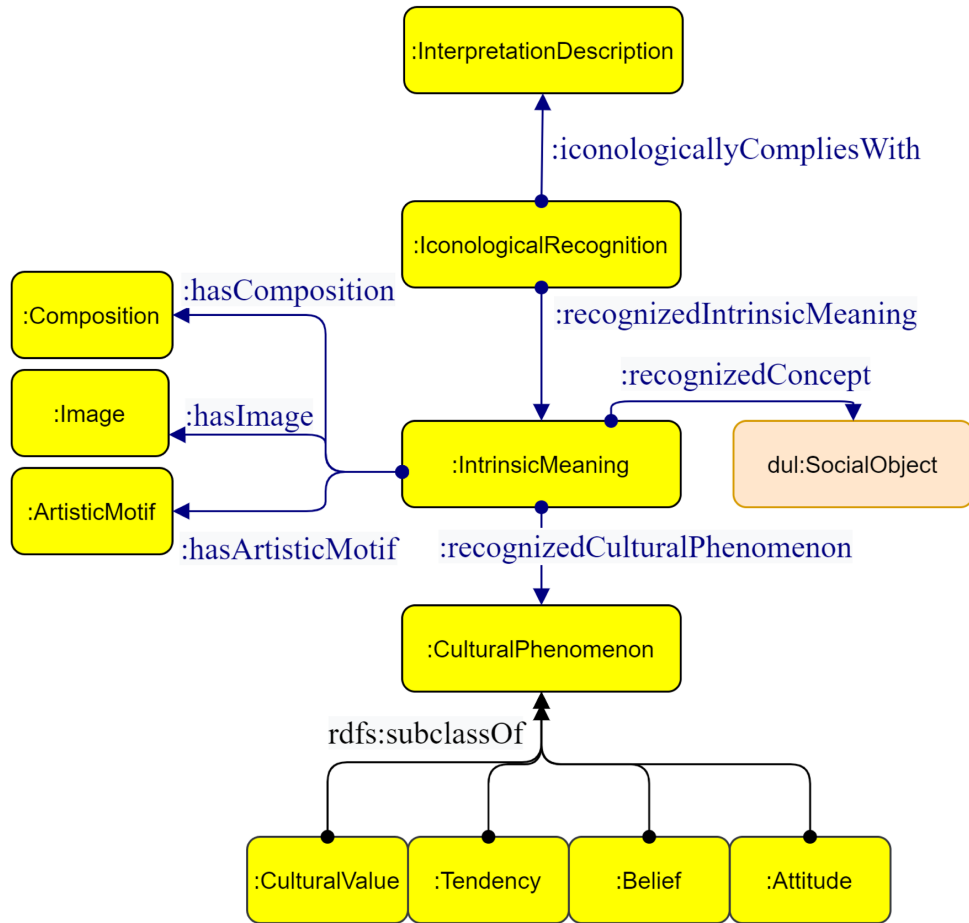


Fig. 4. ICON ontology classes and properties describing the iconological level of interpretation (level 3).

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 [9], which, depending on the specific cases explained in the following, 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 O_1 are logically *aligned* to an external ontology O_2 , which we want to reuse with its full-fledged semantics, because it is compatible, desirable, or necessary. For example, if we extensionally align a O_1 class *Organisation* to a O_2 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 O_3 as purely intensional constructs when we want a limited interoperability, which does not include accepting in O_1 all the semantics provided in O_3 , because it may be partly incompatible. For example, we may intensionally align a O_1 class *Image* to a O_3 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*.

Table 4. Direct Reuse of Classes and Properties in ICON

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

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, and 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 defensible, 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 (see Section 7) for further details on the single alignments. Table 4 shows the direct reuse of external classes and properties in ICON, and Table 5 shows the indirect alignments.

5.5.1 Recognitions as Situations. According to the guidelines of eXtreme Design [42], 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 (e.g., the agent performing it; what is recognized in the form of the first, second, or third level of interpretation entities; the artwork). We have chosen the Situation ODP³⁰ as a solution because it was designed to solve modeling issues regarding multiple contextual information connected to the same class in the form of n-ary relationships. The Situation ODP is reused via the import of DOLCE Ultralite.³¹ The n-ary relationship ODP is specialized by our `icon:Recognition` class, by making it a subclass of `dul:Situation`.

³⁰<http://www.ontologydesignpatterns.org/cp/owl/situation.owl>.

³¹The aforementioned ODP was reused through the DOLCE Ultralite ontology because the Situation class in DOLCE is linked to other classes that are reused in our ontology as well, such as the agent.

Table 5. 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; IntrinsicMeaning	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

5.5.2 Interpretations as Descriptions. The types of recognitions that we have presented are formalized as situations. In the Descriptions and Situations pattern³² that is also formalized in DOLCE Ultralite 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; and in an everyday situation, an observer may interpret it differently according to her perspective, culture, or intention. In the iconographical and iconological domain, as also applied in the ArCo ontology network [10, 11], 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, and so forth [24], which is satisfied by a situation when it involves entities that are classified by one of the parameters, roles, tasks, and so forth 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, and the time period, as well as (potentially) additional iconographical aspects. Hence, we formalize this complex relation in terms of compliance: `InterpretationDescriptionPaintingA isCompliantWithPreiconographicalRecognition`

³²<http://www.ontologydesignpatterns.org/cp/owl/descriptionandsituation.owl>.

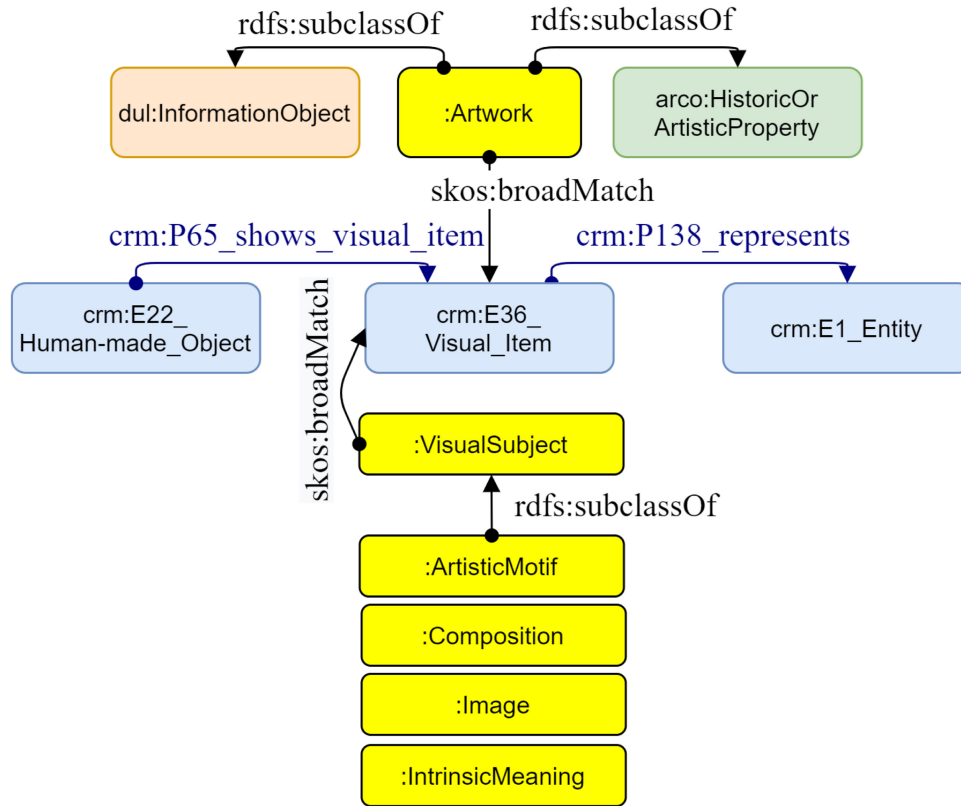


Fig. 5. Alignment with other ontologies of the artwork description.

LionRecognitionInA and isCompliantWithIconographicalRecognition LionCourageRecognitionInA. The property isCompliantWithPreiconographicalRecognition is made a sub-property of dul:isSatisfiedBy that links a dul:Description (our InterpretationDescription is subsumed under description) to one or more dul:Situation (our Recognition is subsumed under situation).

5.5.3 *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, 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. The 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 Information Object that fits with our Artwork definition.

5.5.4 *Interpretation Details.* The class Recognition has been aligned with classes from HiCO, CIDOC-CRM, and DOLCE, as shown in Figure 1. The class hico:InterpretationAct is intended to represent the context

in which a recognition³³ is made—that is, 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`, since `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`, 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`, and respectively, `RecognizedArtisticMotif` or `RecognizedComposition` at level 1, and `RefersToArtisticMotif` and `RecognizedImage` 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 criterion³⁴ `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 of the relief *Allegory of salvation* from the Roman relief depicting *Hercules and the Caledonian Boar* described by Panofsky and Saxl [38, p. 228; Figure 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`.

5.5.5 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, whereas 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 artifacts 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 a subclass of `icon:VisualSubject`, which is disjoint with `arco:Subject` to underline their different nature and role. Table 6 displays the division between subjects and visual subjects according to the different iconographic and iconological levels.

5.5.6 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 [50]. This ontology,

³³In the context of this study, since we align `hico:InterpretationAct` to `Recognition`, we refer to it with the term *recognition* for the sake of clarity.

³⁴In the case of our ontology, interpretation criterion is linked to every single recognition and not with the interpretations.

Table 6. 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

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.

5.5.7 Expression of Style. The expression of style is an important feature related to iconographical and iconological studies (requirement 10, Section 4). Knowing the history of styles is, according to Panofsky [37], a fundamental requirement for the correct interpretation of level 1 objects. Furthermore, as is evident, among the others, from Warburg’s studies on *Pathosformeln* and *Nachleben der Antike*, forms of style are a subject of interest in iconology. Therefore, we reuse CIDOC-CRM to model it according to the solution adopted by linked.art project,³⁵ 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 those of CIDOCs. Therefore, we reuse this existing solution to model requirement 10 of Section 4. In this way, both the artwork itself and every portion of the image identified at each level can have its own style specification declared.

5.5.8 Citations, Sources, Evidence. As shown in Figure 1, 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

³⁵<https://linked.art/model/object/aboutness/#style-classification>.

representation is fundamental to encourage a documented description of the recognition, its reference, and support.

5.6 Logical Constraints in Artistic Interpretations

Artistic interpretations are, by definition, subjective. Nevertheless, the characteristics and relationships that surround recognized elements can be subject to logical constraints. Each recognition is made on exactly one artwork, involves exactly one agent, and different recognitions require adequate elements. A pre-iconographical recognition targets as recognized elements, either artistic motifs or compositions. A formal motif recognition instead deals with a prototypical motif and a copied motif, and both of them can be either an artistic motif or a composition. Then, an iconographical recognition recognizes exactly one image or one *invenzione*. Finally, an iconological recognition refers to an intrinsic meaning. For what it concerns the elements that are recognized, the recognition of an artistic motif, in Panofsky terms (and thus in our ontology) implies that the interpreter associates either a natural or expressional meaning to a portion of the artwork. At the same time, the recognition of an image implies the presence of either a character, event, named object, symbol, personification, or specific place in the artwork. Furthermore, images, artistic motifs, and intrinsic meanings cannot be instantiated without having a recognition that addresses them. To ensure this, we added several restrictions. For example, images must be linked to exactly 1 recognition through the property `isIconographicallyRecognizedBy`. The difference between stories and allegories is that the former generally includes characters, places, events, and named objects, and the latter is more focused on symbols and personifications. Nevertheless, a story might contain symbols and an allegory may contain characters, so the logical restrictions in these cases are not very strict. An intrinsic meaning can refer to either a cultural phenomenon or a conceptual object.

An exemplification of some restrictions on main classes through OWL axioms follows, formalized in Manchester Syntax³⁶:

- *InterpretationDescription*:

```
isCompliantWithIconographicalRecognition some IconographicalRecognition
```

```
isCompliantWithIconologicalRecognition some IconologicalRecognition
```

```
isCompliantWithPreiconographicalRecognition some (PreiconographicalRecognition or
FormalMotifRecognition)
```

- *Recognition*:

```
aboutWorkOfArt exactly 1 Artwork
```

```
includesAgent exactly 1 Agent
```

- *PreiconographicalRecognition*:

```
recognizedArtisticMotif exactly 1 ArtisticMotif or
```

```
recognizedComposition exactly 1 Composition
```

- *FormalMotifRecognition*:

```
hasCopiedMotif exactly 1 ArtisticMotif or
```

```
hasCopiedMotif exactly 1 Composition
```

³⁶<https://www.w3.org/TR/owl2-manchester-syntax/>.

- hasPrototypicalMotif exactly 1 ArtisticMotif or
 - hasPrototypicalMotif exactly 1 Composition
- *IconographicalRecognition*:
 - recognizedImage exactly 1 Image or
 - recognizedInvenzione exactly 1 Invenzione
- *IconologicalRecognition*:
 - recognizedIntrinsicMeaning exactly 1 IntrinsicMeaning
- *IntrinsicMeaning*:
 - ((recognizedConcept exactly 1 SocialObject)
 - and (recognizedCulturalPhenomenon exactly 1 CulturalPhenomenon))
 - or ((recognizedConcept exactly 1 SocialObject)
 - or (recognizedCulturalPhenomenon exactly 1 CulturalPhenomenon))
 - isIntrinsicMeaningOf exactly 1 IconologicalRecognition
- *ArtisticMotif*:
 - (hasExpressionalMeaning min 1 ExpressionalQuality) or
 - (hasFactualMeaning min 1 (Action or NaturalElement))
 - isRecognizedArtisticMotifOf exactly 1 PreiconographicalRecognition
- *Composition*:
 - hasPart min 1 ArtisticMotif
 - isRecognizedCompositionOf exactly 1 PreiconographicalRecognition
- *Story*:
 - composedOf some Character or composedOf some Event
 - or composedOf some NamedObject or composedOf some Place
- *Image*:
 - (hasCharacter min 1 Character)
 - or (hasEvent min 1 Event)
 - or (hasNamedObject min 1 NamedObject)
 - or (hasPersonification min 1 Personification)
 - or (hasPlace min 1 Place)
 - or (hasSymbol min 1 (Symbol or Simulation))
 - isIconographicallyRecognizedBy exactly 1 IconographicalRecognition
- *Allegory*:
 - composedOf some Personification or composedOf some Symbol

Table 7. Test Dataset Overview of the 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

This list contains only the directly created axioms. The restrictions inherited by the alignment to external ontologies are available in the documentation.³⁷

6 EVALUATION

The ICON ontology was evaluated in (i) its extraction potential through testing the CQs [57] on a real-world interpretations dataset to comply with RQ1, (ii) gauging its granularity potential by comparing data of interpretations written using the model against the same interpretations encoded with other ontologies to comply with RQ2, and (iii) logical consistency, FAIRness, and syntax using selected tools and services.

The following sections deal with the creation of the evaluation dataset and the different evaluation methods.

6.1 Creation of Evaluation Dataset

To the best of our knowledge, existing linked open data hubs do not include detailed descriptions of iconological aspects needed to evaluate the ontology. Therefore, we tested ICON on manually created data extracted from Panofsky's *Studies in Iconology* interpretations³⁸ [37]. This text was chosen for its historical importance and authoritativeness in the domain at hand. 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 has to be considered that the statements obtained depend on the author's subjective comprehension of Panofsky's work. For this reason, the author is indicated as responsible for the graph created, whereas Panofsky's text is always cited as the source of each statement. The data creation was conducted as follows. First, we created a tabular data structure reflecting the ontology structure, including information on each level of interpretation. Second, we interpreted Panofsky's claims and described the features of interest according to the data structure. The quality of the data was assessed through values validation through controlled lists. Finally, the data were converted into RDF³⁹ according to the modeling of the ontology described in Section 5 and reusing CIDOC-CRM for the description of artworks' metadata. As a result, the dataset contains a total of 28,864 triples about 152 artworks, 1,980 interpretations, and 928 subjects. Additional statistics about the dataset can be found in Table 7.

6.2 CQ Evaluation

The ontology was evaluated on Panofsky's dataset through the CQs listed in Table 3. Each SAMOD iteration corresponds to a group of questions at levels 0, 1, 2, or 3, and expected results for each CQ are described. To test the correctness of the single classes, the CQs were further subdivided into more detailed ones closer to the

³⁷<https://w3id.org/icon/docs/>.

³⁸The data is a subset of a dataset containing Panofsky's interpretation available at <https://w3id.org/icon/data/>. The dataset's exploratory data analysis is provided at <https://iconology-dataset.streamlit.app/>.

³⁹The conversion was done using the Python RDFlib library.

Table 8. Information About Artworks Cited in the Queries

ID	Description
ART1195	Piero di Cosimo, <i>The Finding of Vulcan</i> , 1485–1490, Hartford, Wadsworth Atheneum
ART1266	Vatican City, Biblioteca Apostolica Vaticana, <i>L'Ovide moralisé</i> , XIV Century, Cod. Reg. 1480, folio 5r. Anonymous, <i>Saturn</i>
ART1267	<i>Saturn</i> , first third of XV Century, Dresden, Kupferstichkabinett
ART1268	Jacopo Caraglio after Rosso Fiorentino, <i>Saturn</i> , Engraving B24, 1526
ART1269	<i>Saturn and His Children</i> , in Cim. 10, Middle XV Century, Berlin, Kupferstichkabinett
ART1270	Vatican City, Biblioteca Apostolica Vaticana, Cod. Pal. lat. 1368, folio 1v, XVI Century. Anonymous, <i>Saturn</i>
ART1284	Giovanni Rost after Agnolo Bronzino, <i>Flora</i> , Florence, Galleria degli Arazzi
ART1285	Albrecht Dürer, <i>The Abduction of Proserpine</i> , 1516
ART1289	Nicolas Poussin, <i>Phaethon Before Helios</i> , Berlin, Kaiser Friedrich Museum
ART1346	Michelangelo, <i>Pen Drawing</i> , Fr. 103, 1504–1505, London, British Museum
ART1534	Rubens, <i>Saturn Devouring a Son</i> , 1636–1638, Madrid, Prado

ontology structure. All unit tests that query the test dataset are available through GitHub in the form of Jupyter notebooks.⁴⁰ For each level, we describe one or two CQs. Table 8 contains an overview of the metadata about the artworks included in the queries.

6.2.1 CQ Level 1. The query presented here is part of 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 obtained ones, are shown in Table 9.

CQ 1.2.1: *Retrieve all the natural, expressional meanings, and actions recognized in the artistic motifs of ART1195.*

```

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

SELECT DISTINCT ?natural ?expressional ?action WHERE {
?icrec icon:aboutWorkOfArt d:ART1195;
    {?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 a icon:Action}
}
    
```

The second level 1 query listed in the following 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 10 shows 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)?*

```

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

SELECT DISTINCT ?subject ?copiedSubject WHERE {
?rec a icon:FormalMotifRecognition;
    icon:aboutWorkOfArt d:ART1284, d:ART1285.
}
    
```

⁴⁰<https://w3id.org/icon/development>.

Table 9. 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	

Table 10. CQ 1.3.1 Results

Subject	copiedSubject
woman	woman
riding-on	riding-on
unicorn	ram

```

{?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:hasCopiedMotif ?comp. ?comp icon:hasPart ?copied}
?copied icon:hasFactualMeaning | icon:hasExpressionalQuality ?copiedSubject.
}

```

6.2.2 *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 following query 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 date and place of production of the artwork to obtain a detailed representation of the subject variations over

Table 11. CQ 2.6.1 Results

Level 1	Total	Level 1	Total	Level 1	Total
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

place and time. As a consequence, it can be a useful tool for art historians for integrating qualitative iconographical analysis with a quantitative overview of the phenomenon. The results of CQ 2.6.1 are shown in Table 11.

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 do they appear.*

```

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

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)
    
```

6.2.3 *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 to the cultural phenomenon CF1087 “Evolution of the Iconography of Saturn.”*

```

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

```

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.

6.2.4 General CQs. For the general-level CQs, we present two CQs. 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 contrasting interpretations. While describing it [37], the art historian Erwin Panofsky states that his own position diverges from the usual interpretation according to which the artwork depicts the myth Hylas and the Nymphs. He cites the works of other three 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?

```

PREFIX d: <https://w3id.org/icon/data/>
PREFIX icon: <https://w3id.org/icon/ontology/>
PREFIX cito: <http://purl.org/spar/cito/>

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 12. The artwork has four different interpretation descriptions, among which only DESC1 has recognitions at all the levels, whereas the remaining have only level 2 recognitions. To better see the agreement and disagreement of them, 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 CQ (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 details 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?

```

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

```

Table 12. 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

```

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.

6.3 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 selected the frontispiece of François Perrier’s “Segmenta nobilium signorum statuarum...” (1638 edition) depicting the iconography of Father Time, a subject that emerges from the Renaissance onward but, despite it originates 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) [37]. 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 reinterpretation of the traditional Saturn’s attributes, such as the sickle, associated with the times recurring,⁴¹ or the act of devouring the children, reinterpreted as Time devouring “whatever he has created” [37, p. 74]. Panofsky claims that such 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, albeit they were described in 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

⁴¹“Tempora quae sicut falx in se recurrunt”, trad. by the authors: “Times which, like the sickle, recur” [48, II, 406].

Table 13. “Time the Destroyer” Description According to the Three Levels of Interpretation

	man, nude and old, with a scythe and wings gnawing away a statue
1	snake biting its tail fragments of classical buildings and of 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” [37] (2) Pseudomorphosis: “certain Renaissance figures became invested with a meaning which, for all their classicising appearance, had not been present in their classical prototypes, though it had frequently been foreshadowed in classical literature” [37]

Table 14. Comparison Between Wikidata, VIR, and ICON

Ontology	Level 1 Subject	Level 2 Subject	Level 3 Subject	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 levels

iron teeth standing among ruins, a symbol of the fact that time ruins everything without any effort.⁴² Table 13 resumes the understanding of this artwork at each different level, whereas Figures 6 through 9 shows how this example can be modeled with the ontologies considered.

Table 14 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 [8]. 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 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 cannot 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 6 (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,

⁴²“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 without any effort” [45].

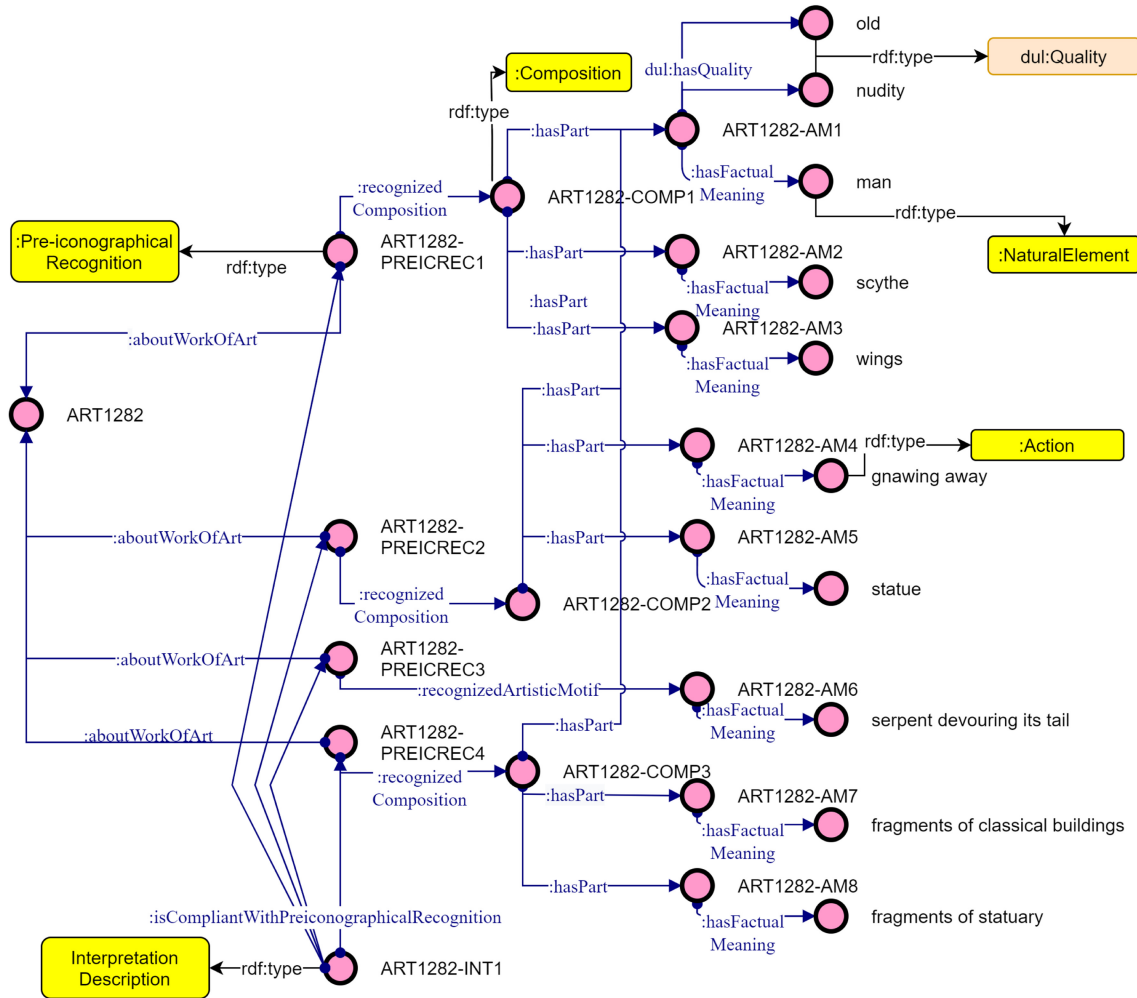


Fig. 6. Description of the example using ICON (level 1).

such as attributes, personifications, symbols, places, and characters (Figure 8), 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” (Figure 9). 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 due to the fact that 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 previously allow a

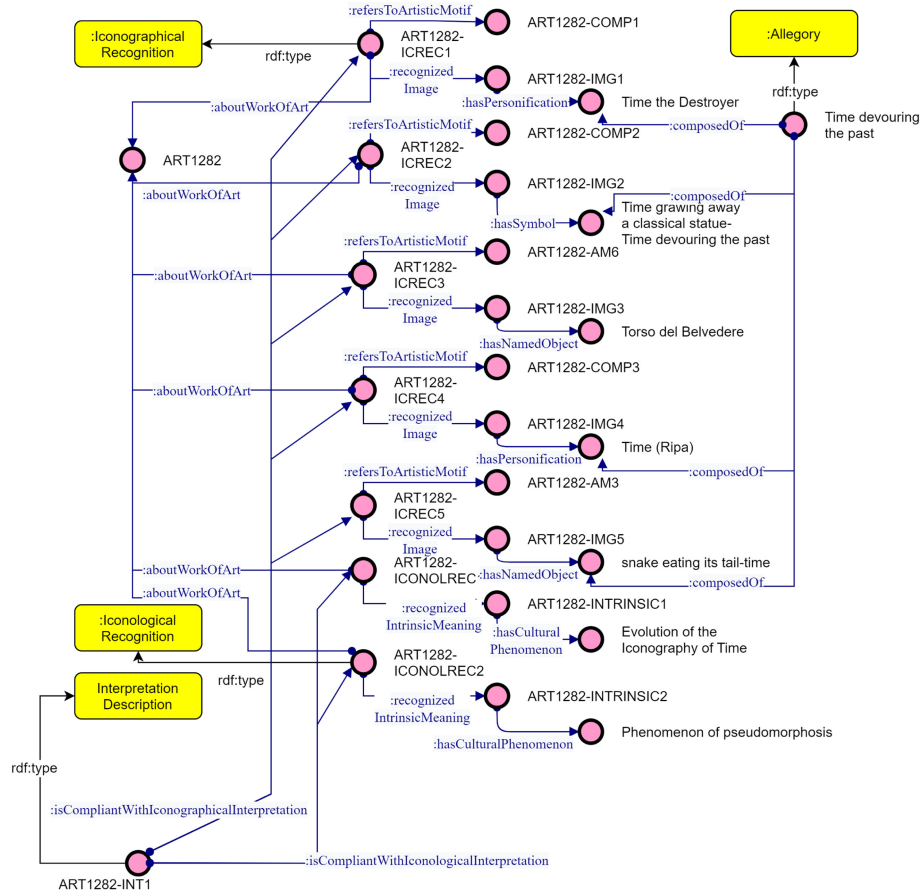


Fig. 7. Description of the example using ICON (levels 2 and 3).

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:Iconological Recognition relating a third-level meaning (concept or cultural phenomenon) to the whole artwork or to its specific parts.

Concerning the interpretation’s attribution of responsibility, both VIR and Wikidata allow registering 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, 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 subject recognized. In this way, more agreeing claims can be expressed in the same recognition description, to better represent the realistic case in which an art historian agrees with the claim of others, quoting them and adding further interpretations. In addition, this structure fosters the interoperability of online sources and data integration.

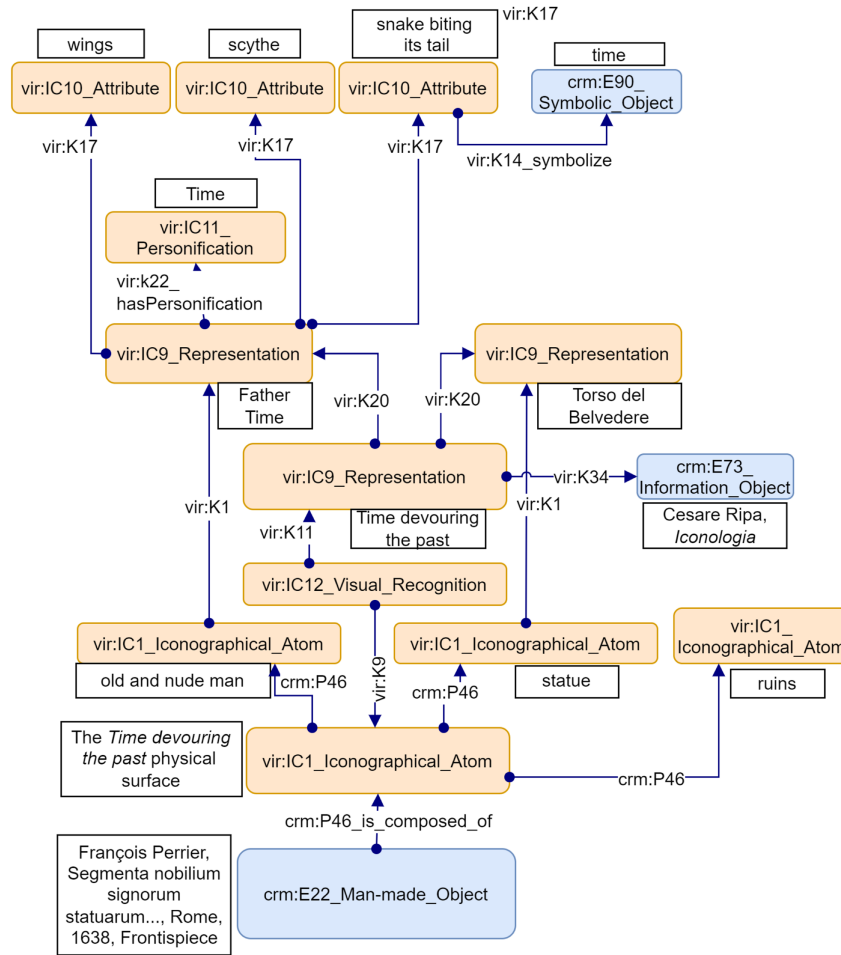


Fig. 8. Description of the example using VIR.

6.4 Automatic Evaluation

A part 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 by the WRC RDF Validation Service.⁴³ No syntax problems were highlighted by this tool. Then, we evaluated the logical consistency of our ontology through the OOPS [41] tool⁴⁴ 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 (e.g., no ranges and domain in properties or inconsistent labeling). The only highlighted pitfall that was directly linked with 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`

⁴³<https://www.w3.org/RDF/Validator/>.

⁴⁴We 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>).

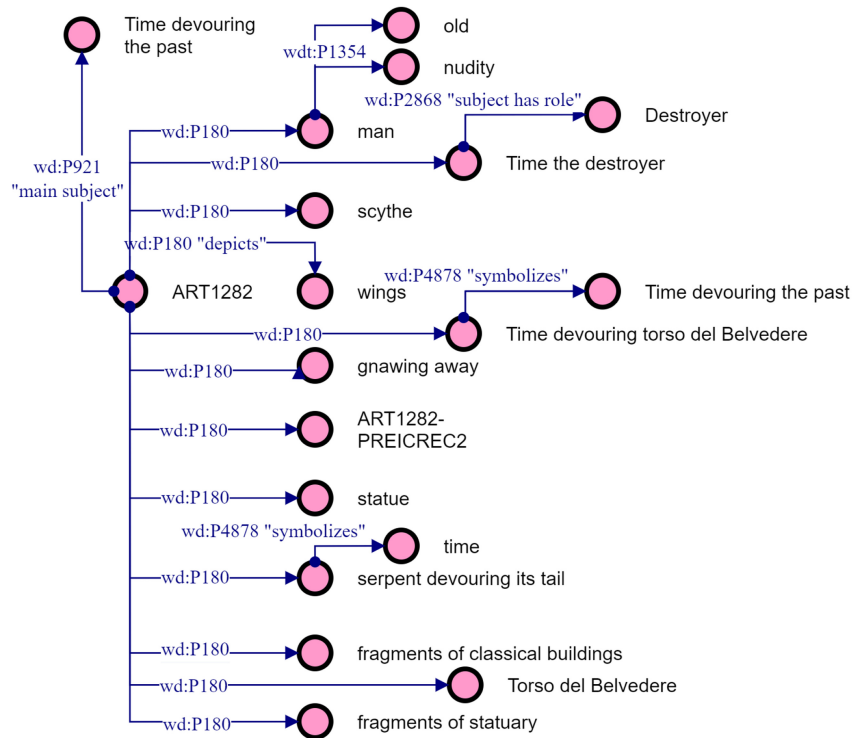


Fig. 9. Description of the example using Wikidata.

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`. 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 considered instead 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 originally created classes only). Finally, we analyzed our ontology with the FOOPS [26] 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, a score of 8 out of 9 on findability, a score of 3 out of 3 on interoperability, and a score of 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).⁴⁵ 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 to our ontology did not cause any inconsistency in `ICON` by running the Hermit Reasoner [28].

7 ONTOLOGY RELEASE AND DOCUMENTATION

The current version of `ICON` ontology contains 28 classes and 58 properties (of which one is a data property). Additionally, it reuses classes and properties taken from six ontologies. Its release comes with a documentation

⁴⁵<https://lov.linkeddata.es/dataset/lov/>.

that gives a definition for all its classes and properties, along with examples on the use of the ontology. Several permanent URIs were provided by the w3id service:

- <https://w3id.org/icon/docs/>: that links to the documentation of the ontology generated with the widoco tool [25].
- <https://w3id.org/icon/ontology/>: that leads to the OWL file of the ontology (RDF/XML serialization)
- <https://w3id.org/icon/development/>: that links to the GitHub repository where the SAMOD iterations and eXtreme Design unit tests can be found.

The preferred prefix of the ontology (icon) was registered on <http://prefix.cc>.

8 CONCLUSION, LIMITATIONS, AND FUTURE WORK

In this article, we presented ICON, an ontology dedicated to the conceptualization of 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 a documentation that guides potential users in formalizing art interpretations using our model. ICON was evaluated on its extraction potential, syntax, metadata, and FAIRness. Moreover, its granularity was highlighted through a comparison to 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 Semantic Web by providing a granularity level that was not reached by other ontologies on this topic. Finally, its effectiveness of describing the iconographical and iconological complex domain is confirmed by the results of the proposed CQs, formalized in SPARQL queries, ran on a test dataset containing artistic interpretations.

Although ICON provides users the option to describe a plethora of concepts related to art interpretation, the ontology presented here is still a work in progress. Future work will be directed to a more thorough conceptualization of cultural phenomena, personifications, and allegories in the same way that cultural symbols were defined by the Simulation Ontology [50]. Furthermore, modeling artistic interpretations required the use of several agglomerated n-ary relationship classes, drawing a long path from the artwork to its meaning. On the one hand, this modeling offers a very high granularity level, as interpretations can be dissected into recognitions representing different levels (pre-iconographic, iconographic, iconological), allowing the potential extraction of very specific information as shown by the testing of the CQs. On the other hand, potential users might only be interested into separating what is depicted in an artwork into the aforementioned levels, without having to describe the whole process of interpretation. Future work will be devoted to the creation of property chains that can be used to declare that an artwork represents elements of a pre-iconographic, iconographic, or iconological level. Both solutions (series of n-ary classes) and the property chains could be adopted, in a knowledge graph, at the same time, as in ArCo [10], which uses property chains as shortcuts, and n-ary relationship classes to describe the same information with more granularity. As for the alignments, although, as Table 5 shows, there are more than 50 alignments to external ontologies, we decided to keep a conservative approach for this version of ICON. Future versions will foster the interoperability of ICON even more by (i) finding other ontologies that cover similar aspects to increase alignments (e.g., [14]), and (ii) by specializing the alignments with the currently aligned ontologies by finding more specific classes and properties. The `icon:Invenzione` class, for instance, is aligned to a very generic `crm:E1_CRM_Entity`. Future versions of CIDOC might introduce new classes that are closely related to ours, leading to a more specific alignment.

On another note, the current debate of interpretation provenance and the difference between “asserting and expressing” [17] in a Semantic Web context using recently introduced technologies of RDF-star⁴⁶ and conjectures [17] could use a dataset of contrasting interpretations described with our ontology as a case study.

⁴⁶<https://w3c.github.io/rdf-star/cg-spec/2021-12-17.html>.

In conclusion, by combining our model with the external ontologies to which it is aligned, it is now possible to thoroughly describe artworks both in their standard metadata (i.e., creator, date of creation, dimensions, place of creation) and on their content side based on interpretations. The result is that, finally, these two evenly important types of information are now treated equally and can both exploit the potentiality offered by the Semantic Web. This contribution opens up the possibility to link artworks in their content level, allowing content-based research questions in the field of art history to cross into the linked open data realm.

REFERENCES

- [1] Laurie Adams. 2010. *The Methodologies of Art: An Introduction* (2nd ed.). Westview Press.
- [2] Sören Auer, Christian Bizer, Georgi Kobilarov, Jens Lehmann, Richard Cyganiak, and Zachary Ives. 2007. Dbpedia: A nucleus for a web of open data. In *The Semantic Web*. Springer, 722–735.
- [3] S. Baroncini, M. Daquino, and F. Tomasi. 2021. Modelling art interpretation and meaning: A data model for describing iconology and iconography. *arXiv:2106.12967* (2021).
- [4] Sofia Baroncini, Bruno Sartini, Marieke van Erp, Francesca Tomasi, and Aldo Gangemi. 2023. Is dc:subject enough? A landscape on iconography and iconology statements of knowledge graphs in the semantic web. *Journal of Documentation* 79, 7 (2023), 115–136. <https://doi.org/10.1108/JD-09-2022-0207>
- [5] C. Bekiari, G. Bruseker, M. Doerr, C.-E. Ore, S. Stead, and A. Velios. 2021. *Definition of the CIDOC Conceptual Reference Model, Version 7.1.1*. CIDOC. <https://doi.org/10.26225/FDZH-X261>
- [6] Stefano Borgo, Roberta Ferrario, Aldo Gangemi, Nicola Guarino, Claudio Masolo, Daniele Porello, Emilio M. Sanfilippo, and Laure Vieu. 2022. DOLCE: A descriptive ontology for linguistic and cognitive engineering. *Applied Ontology* 17, 1 (2022), 45–69.
- [7] Editors of Encyclopaedia Britannica. 2020. Art. Retrieved May 3, 2023 from <https://www.britannica.com/art/visual-arts>.
- [8] Nicola Carboni and Livio de Luca. 2019. An ontological approach to the description of visual and iconographical representations. *Heritage* 2, 2 (2019), 1191–1210. <https://doi.org/10.3390/heritage2020078>
- [9] Valentina Carriero, Marilena Daquino, Aldo Gangemi, Andrea Nuzzolese, Silvio Peroni, Valentina Presutti, and Francesca Tomasi. 2020. The landscape of ontology reuse approaches. In *Applications and Practices in Ontology Design, Extraction, and Reason*, Vol. 49. Studies on the Semantic Web. IOS Press, 21–38.
- [10] Valentina Anita Carriero, Aldo Gangemi, Maria Letizia Mancinelli, Ludovica Marinucci, Andrea Giovanni Nuzzolese, Valentina Presutti, and Chiara Veninata. 2019. ArCo: The Italian cultural heritage knowledge graph. In *The Semantic Web—ISWC 2019*, Chiara Ghidini, Olaf Hartig, Maria Maleshkova, Vojtěch Svátek, Isabel Cruz, Aidan Hogan, Jie Song, Maxime Lefrançois, and Fabien Gandon (Eds.). Springer International Publishing, Cham, Switzerland, 36–52. https://doi.org/10.1007/978-3-030-30796-7_3
- [11] Valentina Anita Carriero, Aldo Gangemi, Maria Letizia Mancinelli, Andrea Giovanni Nuzzolese, Valentina Presutti, and Chiara Veninata. 2021. Pattern-based design applied to cultural heritage knowledge graphs. *Semantic Web Journal* 12, 2 (2021), 313–357.
- [12] Fabio Ciotti. 2016. Toward a formal ontology for narrative. *MATLIT: Materialidades da Literatura* 4, 1 (2016), 29–44.
- [13] H. van de Waal, Rudi H. Fuchs, L. D. Couprie, and E. Tholen. 1973. *Iconclass: An Iconographic Classification System*/H. van de Waal. L. D. Couprie, R. H. Fuchs, and E. Tholen (Eds.). North-Holland Pub. Co., Amsterdam.
- [14] Rossana Damiano and Antonio Lieto. 2013. Ontological representations of narratives: A case study on stories and actions. In *Workshop on Computational Models of Narrative 2013*, Mark A. Finlayson, Bernhard Fisseni, Benedikt Lowe, and Jan Christoph Meister (Eds.). Schloss Dagstuhl—Leibniz-Zentrum fuer Informatik GmbH, Wadern/Saarbruecken, Germany, 76–93.
- [15] Marilena Daquino, Francesca Mambelli, Silvio Peroni, Francesca Tomasi, and Fabio Vitali. 2017. Enhancing semantic expressivity in the cultural heritage domain: Exposing the Zeri Photo Archive as linked open data. *Journal on Computing and Cultural Heritage* 10, 4 (July 2017), Article 21, 21 pages. <https://doi.org/10.1145/3051487>
- [16] Marilena Daquino, Valentina Pasqual, and Francesca Tomasi. 2020. Knowledge representation of digital Hermeneutics of archival and literary sources. *JLIS.it, Italian Journal of Library and Information Science* 11, 3 (2020), 59–77.
- [17] Marilena Daquino, Valentina Pasqual, Francesca Tomasi, and Fabio Vitali. 2021. Expressing without asserting in the arts. In *Proceedings of the Italian Research Conference on Digital Libraries*.
- [18] Marilena Daquino and Francesca Tomasi. 2015. Historical context ontology (HiCO): A conceptual model for describing context information of cultural heritage objects. In *Research Conference on Metadata and Semantics Research*. Springer, 424–436.
- [19] Chris Dijkshoorn, Lizzy Jongma, Lora Aroyo, Jacco Van Ossenbruggen, Guus Schreiber, Wesley Ter Weele, and Jan Wielemaker. 2018. The Rijksmuseum collection as linked data. *Semantic Web* 9, 2 (2018), 221–230.
- [20] Martin Doerr, Chryssoula Bekiari, Athina Kritsotaki, Gerald Hiebel, and Maria Theodoridou. 2014. Modelling scientific activities: Proposal for a global schema for integrating metadata about scientific observation. In *Access and Understanding—Networking in the Digital Era: The 6th Annual Conference of CIDOC, the International Committee for Documentation of ICOM, Dresden, Germany*.
- [21] Umberto Eco. 2016. *Trattato di Semiotica Generale*. La Nave di Teseo Editore spa.
- [22] Ingo Frank. 2019. Multi-perspectival representation of historical reality. In *Proceedings of the Joint Oncology Workshops (JOWO’19)*.

- [23] Aldo Gangemi, Nicola Guarino, Claudio Masolo, Alessandro Oltramari, and Luc Schneider. 2002. Sweetening ontologies with DOLCE. In *Proceedings of the International Conference on Knowledge Engineering and Knowledge Management*. 166–181.
- [24] Aldo Gangemi and Peter Mika. 2003. Understanding the Semantic Web through descriptions and situations. In *On The Move to Meaningful Internet Systems 2003: CoopIS, DOA, and ODBASE*, Robert Meersman, Zahir Tari, and Douglas C. Schmidt (Eds.). Springer, Berlin, Germany, 689–706.
- [25] Daniel Garijo. 2017. WIDOCO: A wizard for documenting ontologies. In *Proceedings of the International Semantic Web Conference*. 94–102. https://doi.org/10.1007/978-3-319-68204-4_9
- [26] Daniel Garijo, Oscar Corcho, and María Poveda-Villalón. 2021. FOOPS!: An ontology pitfall scanner for the FAIR principles. In *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'21)*. https://foops.linkeddata.es/assets/iswc_2021_demo.pdf.
- [27] Richard Gartner. 2020. Towards an ontology-based iconography. *Digital Scholarship in the Humanities* 35, 1 (April 2020), 43–53. <https://doi.org/10.1093/lc/fqz009>
- [28] Birte Glimm, Ian Horrocks, Boris Motik, Giorgos Stoilos, and Zhe Wang. 2014. HermiT: An OWL 2 reasoner. *Journal of Automated Reasoning* 53, 3 (Oct. 2014), 245–269. <https://doi.org/10.1007/s10817-014-9305-1>
- [29] E. H. Gombrich. 1948. Icones symbolicae: The visual image in Neo-platonic thought. *Journal of the Warburg and Courtauld Institutes* 11, 1 (Jan. 1948), 163–192.
- [30] Patricia Harpring. 2010. Development of the Getty vocabularies: AAT, TGN, ULAN, and CONA. *Art Documentation: Journal of the Art Libraries Society of North America* 29, 1 (2010), 67–72. <http://www.jstor.org/stable/27949541>.
- [31] Kerr Houston. 2013. *An Introduction to Art Criticism: Histories, Strategies, Voices*. Pearson, Boston, MA.
- [32] Max Imdahl. 2012. Iconica. L'intuizione delle immagini. *Aisthesis: Pratiche, Linguaggi E Saperi dell'estetico* 5, 2 (2012), 11–32.
- [33] Antoine Isaac and Bernhard Haslhofer. 2013. Europeana linked open data—data.Europeana.Eu. *Semantic Web* 4, 3 (July 2013), 291–297.
- [34] Marion G. Müller. 2014. Iconography and iconology as a visual method and approach. In *The SAGE Handbook of Visual Research Methods*. SAGE Publications Ltd., London, UK, 283–297.
- [35] Inès Osman, Sadok Ben Yahia, and Gayo Diallo. 2021. Ontology integration: Approaches and challenging issues. *Information Fusion* 71 (2021), 38–63.
- [36] Erwin Panofsky. 1955. *Meaning in the Visual Art: Papers in and on Art History*. Anchor Books, New York, NY.
- [37] Erwin Panofsky. 1972. *Studies in Iconology: Humanistic Themes in the Art of the Renaissance*. Westview Press, Boulder, CO. <https://archive.org/details/studiesiniconolo0000pano>.
- [38] Erwin Panofsky and Fritz Saxl. 1933. Classical mythology in mediaeval art. *Metropolitan Museum Studies* 4, 2 (March 1933), 228.
- [39] Silvio Peroni. 2016. A simplified agile methodology for ontology development. In *Proceedings of the 13th OWL: Experiences and Directions Workshop and 5th OWL Reasoner Evaluation Workshop (OWLED-ORE'16)*. 55–69. <https://doi.org/10.6084/M9.FIGSHARE.3189769.V2>
- [40] Davide Picca, Alfio Massimiliano Gliozzo, and Aldo Gangemi. 2008. LMM: An OWL-DL MetaModel to represent heterogeneous lexical knowledge. In *Proceedings of the 6th International Conference on Language Resources and Evaluation (LREC'08)*.
- [41] María Poveda-Villalón, Asunción Gómez-Pérez, and Mari Carmen Suárez-Figueroa. 2014. Oops! (Ontology Pitfall Scanner!): An on-line tool for ontology evaluation. *International Journal on Semantic Web and Information Systems* 10, 2 (2014), 7–34.
- [42] Valentina Presutti, Enrico Daga, Aldo Gangemi, and Eva Blomqvist. 2009. eXtreme design with content ontology design patterns. In *Proceedings of the 2009 International Conference on Ontology Patterns (WOP'09)*. Vol. 516. 83–97.
- [43] Donald Preziosi. 2009. *The Art of Art History: A Critical Anthology*. Oxford University Press, Oxford, UK.
- [44] Thomas Rebele, Fabian Suchanek, Johannes Hoffart, Joanna Biega, Erdal Kuzey, and Gerhard Weikum. 2016. YAGO: A multilingual knowledge base from Wikipedia, WordNet, and GeoNames. In *Proceedings of the International Semantic Web Conference*. 177–185.
- [45] Cesare Ripa. 1764. *Iconologia Del Cavaliere Cesare Ripa, Perugino*, Vol. 3. nella stamperia di Piergiovanni Costantini.
- [46] Patrizia Riva and Maja Zumer. 2017. FRBRoo, the IFLA library reference model, and now LRMoo: A circle of development. In *Proceedings of the Annual Congress of the International Federation of Libraries Association (IFLA WLIC'18)*.
- [47] Stecker Robert. 1996. *Artworks: Meaning, Definition, Value*. Penn State University Press.
- [48] Remigio Sabbadini, Georgius Thilo, and Ermannus Hagen. 1888. Servii grammatici qui feruntur in Vergilii carmina commentarii. *Rivista di Filologia e di Istruzione Classica* 16 (1888), 322.
- [49] Bruno Sartini and Aldo Gangemi. 2021. Towards the unchaining of symbolism from knowledge graphs: How symbolic relationships can link cultures. In *Book of Extended Abstracts of the 10th National AIUCD Conference*. AIUCD, Pisa, Italy, 576–580.
- [50] Bruno Sartini, Marieke van Erp, and Aldo Gangemi. 2021. Marriage is a peach and a chalice: Modelling cultural symbolism on the Semantic Web. In *Proceedings of the 11th Knowledge Capture Conference (K-CAP'21)*. ACM, New York, NY, 201–208. <https://doi.org/10.1145/3460210.3493552>
- [51] David Shotton. 2010. CiTO, the citation typing ontology. *Journal of Biomedical Semantics* 1 (2010), 1–18.
- [52] Elena Simperl. 2009. Reusing ontologies on the Semantic Web: A feasibility study. *Data & Knowledge Engineering* 68, 10 (2009), 905–925. <https://doi.org/10.1016/j.datak.2009.02.002>

- [53] Stuart Snyderman, Robert Sanderson, and Tom Cramer. 2015. The International Image Interoperability Framework (IIIF): A community & technology approach for web-based images. In *Archiving Conference*, Vol. 2015. Society for Imaging Science and Technology, 16–21.
- [54] Luc Steels. 2022. *Foundations for Meaning and Understanding in Human-Centric AI*. Venice International University. <https://doi.org/10.5281/zenodo.6666820>
- [55] Mari Carmen Suárez-Figueroa, Asunción Gómez-Pérez, and Boris Villazón-Terrazas. 2009. How to write and use the ontology requirements specification document. In *On the Move to Meaningful Internet Systems: OTM 2009*. Springer, Berlin, Germany, 966–982.
- [56] Mischa M. Tuffield, Dave E. Millard, and Nigel R. Shadbolt. 2006. Ontological approaches to modelling narrative. In *Proceedings of the 2nd AKT DTA Symposium*.
- [57] Michael Uschold and Michael Grüninger. 1996. Ontologies: Principles, methods and applications. *Knowledge Engineering Review* 11, 2 (1996), 93–136.
- [58] M. Van Ruymbeke, P. Hallot, and R. Billen. 2017. Enhancing CIDOC-CRM and compatible models with the concept of multiple interpretation. *ISPRS Annals of the Photogrammetry, Remote Sensing and Spatial Information Sciences IV-2/W2* (Aug. 2017), 287–294.
- [59] Roelof van Straten. 2012. *An Introduction to Iconography* (2nd ed.). Routledge, London, England.
- [60] Denny Vrandečić and Markus Krötzsch. 2014. Wikidata: A free collaborative knowledgebase. *Communications of the ACM* 57, 10 (Sept. 2014), 78–85. <https://doi.org/10.1145/2629489>
- [61] Aby M. Warburg. 1999. *The Renewal of Pagan Antiquity—Contributions to the Cultural History of the European Renaissance*. Getty Publications, Los Angeles, CA.
- [62] Rudolf Wittkower. 1987. *Allegory and the Migration of Symbols*. Thames & Hudson, London, England.
- [63] Lei Xu, Albert Merono-Penuela, Zhisheng Huang, and Frank Van Harmelen. 2017. An ontology model for narrative image annotation in the field of cultural heritage. In *Proceedings of the Workshop on Humanities in the Semantic Web*. 15–26.

Received 2 June 2022; revised 19 February 2023; accepted 15 March 2023