

# Semantic Data Retrieval and Integration for Supporting Artworks Interpretation Through Integrative Narrative Networks

Sofia Baroncini<sup>1</sup>, Luc Steels<sup>2,3</sup> and Remi van Trijp<sup>3</sup>

<sup>1</sup>*DHARC-Digital Humanities Advanced Research Centre, University of Bologna, Italy*

<sup>2</sup>*Venice International University, Italy*

<sup>3</sup>*Sony Computer Science Laboratory Paris, France*

## Abstract

Significant recent advances in AI are progressively giving Digital Humanities a range of powerful tools to analyse and contextualise artworks using techniques from computer vision, pattern recognition, ontology engineering, natural language processing, and the semantic web. These tools help to analyse artworks and link them to insightful descriptions. However, to obtain the full potential of these tools we need to tackle two issues: (i) how to integrate the fragmented and sometimes contradictory information these various tools provide, and (ii) how to make it much easier for art historians, curators, and artists to use and extend these tools.

This paper addressed these questions with a focus on semantic web information retrieval and integration. It introduces a data structure called an Integrative Narrative Network (INN) that supports the integration of information from different knowledge sources, which formally represents the process of understanding as a question-answering approach. It further introduces the ongoing development of a tool by which an art historian can build up narrative networks by retrieving and selecting information queried from online available Knowledge Graphs. In particular, we show how semantic web resources can help to raise questions and find answers to them, through the real case study of a Late Renaissance artwork interpretation.

## Keywords

Art interpretation, AI tools, Integrative Narrative networks, Semantic data retrieval, semantic data integration.

## 1. Narrative-based Art Interpretation

### 1.1. Understanding as question-answering

There has been remarkable progress the past decade in tools to analyse and contextualise artworks based on techniques from computer vision, pattern recognition, ontology engineering, natural language processing, and the semantic web.<sup>1</sup> What is missing however are good ways

---

*SWODCH'23: International Workshop on Semantic Web and Ontology Design for Cultural Heritage, November 7, 2023, Athens, Greece*

✉ [sofia.baroncini4@unibo.it](mailto:sofia.baroncini4@unibo.it) (S. Baroncini); [steels@arti.vub.ac.be](mailto:steels@arti.vub.ac.be) (L. Steels); [Remi.Vantrijp@sony.com](mailto:Remi.Vantrijp@sony.com) (R. v. Trijp)

🆔 0000-0002-5636-8328 (S. Baroncini); 0000-0001-9134-3663 (L. Steels); 0000-0002-0475-8367 (R. v. Trijp)



© 2023 Copyright for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).

CEUR Workshop Proceedings (CEUR-WS.org)

<sup>1</sup>S. Baroncini is responsible for sections 1.2, 3.1, 3.2, and 3.4. L. Steels is responsible for sections 1.1, and 2. R. van Trijp is responsible for section 3.3. All the authors are responsible for section 4.

to combine the outcomes of these various tools into one coherent interpretation and to make the results available both to art historians and to viewers.

The human interpretation process strives for narrative-based understanding. When interpreting an artwork the viewer tries to construct a narrative that integrates the visual depictions, memory of past experiences, knowledge of artistic styles, the personal history and prior work of the artist, and general world knowledge, in order to answer a series of questions (e.g. Who is depicted? Why?) and to put an artwork in context (e.g. Who is the painter? What was the function of the work?).

Narrative-based understanding has to be conceived of as a spiraling process.[1, 2]. Starting from an initial examination of some input elements with a lot of ambiguity, uncertainty and indeterminacy, hypotheses of the whole are constructed, which then provide top-down expectations to be tested by a more detailed examination of the same or additional elements, leading to a clearer view of the whole, which then leads back to the examination of additional elements, etc., until the narrative 'makes sense' and resonates with the personal episodic memory of the viewer, reaching a state known as *narrative closure* in literature studies [3].

When considering partial digital support for interpretation we need not work not only on the individual knowledge sources that can answer or raise questions but also on the integration process and this requires as a first, fundamental step a data structure that supports integration.

For this purpose, we present a data structure called **Integrative Narrative Network (INN)**. An INN acts as a kind of blackboard on which different knowledge sources write partial descriptions and through which they can consult the information provided by other knowledge sources to advance understanding. We operationalise the understanding process in terms of narrative questions and answers, defining questions as open slots which can be filled by knowledge sources by (i) evoking new questions, (ii) introducing answers to questions, (iii) constraining the answers to questions, or (iv) shrinking the set of questions by realizing that the answers to two different questions are in fact the same.

## 1.2. Case Study

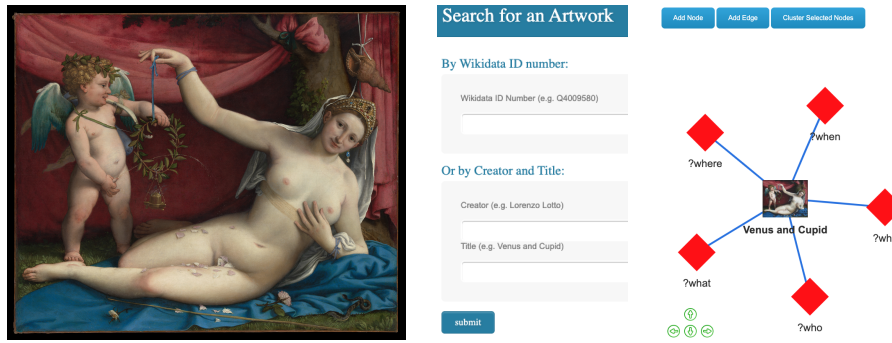
We have already applied the methodology and techniques described in this paper to other domains (e.g. history [4]) and to the interpretation of the contemporary artwork 'Secrets' by the painter Luc Tuymans. This experiment was the subject of an exhibition at the BOZAR cultural centre in Brussels in 2011<sup>2</sup>. For that project we focused on integrating results coming from pattern recognition and image processing as reported in [5] and from NLP techniques applied to text from the catalogue and other textual resources.

In this paper, we explore historical artworks from the late Renaissance period by the Venetian painter Lorenzo Lotto and we focus on integrating results from semantic web resources. More concretely we report on a use case for the painting "Venus and Cupid" by Lorenzo Lotto<sup>3</sup> (see Fig. 1, left). This painting was chosen as a feasible case study for art understanding since the lack of textual documentation and its complex iconography constitute an interpretative challenge for art historians [6].

---

<sup>2</sup><https://readymag.com/u3083945729/secrets-guide/>

<sup>3</sup>Catalogue entry available at <https://www.metmuseum.org/en/art/collection/search/436918>



**Figure 1:** Left: Lorenzo Lotto, *Venus and Cupid*, 1520s, Metropolitan Museum of Art, New York. Public domain. Right: Interface for art historians to steer the INN expansion process. The historian can click on certain questions to expand them further and zoom in and out.

At first glance, the observer’s attention is driven by some visual indicators, such as the brightness contrast, to the main characters, i.e. a putto (little angel) and a woman. If the observer is familiar with Italian Renaissance themes, s/he will recognise immediately these figures respectively as Cupid and Venus. Nevertheless, the curious actions performed and the presence of uncommon objects may lead the observer to raise new questions about their meaning. Supported by a possible knowledge of Lotto’s artistic practice of depicting symbols in his artworks, the observer, particularly if s/he is an art historian, will search for possible meanings of the objects.

Narrative closure is reached when answers to all these questions are coherently integrated with each other and grounded in external and past experiences. For example, the art historian Christiansen recognises that the characters’ actions and the majority of the objects belong to the sphere of marital love, concluding that the artwork expresses a wedding wish<sup>4</sup>.

The focus of the interpreter has gone beyond the main focal point of the painting (i.e. Venus and Cupid), to explore secondary details, which are now seen as signs, triggering a deeper interpretation of the scene.

## 2. Integrative datastructure (INN)

### 2.1. Narrative questions and answers

As explained, we frame art interpretation as a process whereby the agent raises questions, finds answers to questions, and interlinks questions and answers. And we focus in particular on how semantic web resources can help to raise questions and find answers to them.

**What are the questions?** Questions are computationally operationalised as variables. Following AI tradition, the name of the question is written as a symbol with a question mark in front, as in ?question-name (e.g. ?Where).

**What are the answers?** The answers to a question are entities in the domain of discourse. Entities are objects, events or (reified) concepts. They either refer to real-world observational

<sup>4</sup>A more thorough artwork’s interpretation description is available on catalogue entry of the MET Museum.

data (for example a physical painting, a region in an image), to virtual entities (which may or may not exist in reality), or to entities in a knowledge graph in which case we use the URI (Universal Resource Identifier) as a unique identifier.

INN makes the relation between a question and an answer computationally operational in terms of binding between questions (which are technically variables) and identifiers of entities or constants.

**Where do questions come from?** We use a frame-based approach which was pioneered in research on knowledge representation and object-oriented programming, starting in the mid-seventies with the proposal by [7] and technical realisations such as KRL [8], KRS [9] or CLOS [10]. In the present INN implementation, we use CLOS. In Minsky's original conception, a frame is a bundle of important questions to be asked about a particular type of object. A frame has a set of slots with values, which are in effect the questions that can be asked about a particular entity represented by a frame and respectively the answers to these questions. An entity which is described with a certain frame is called a frame-instance of that frame.

**Visualisation** The INN represents all questions that have been posed during the understanding process and all entities that appear as answers to these questions. The questions are visualized with green or red diamonds. The diamond is green if the question could be answered and red if it is still open. The answers to questions are represented with squares and can be frame-instances or constants (e.g. a number or a boolean value).<sup>5</sup>

### 3. Case study for the Lotto painting

This section illustrates the INN and its use in interpretation for a concrete case study on the painting *Venus and Cupid* by Lorenzo Lotto introduced earlier in section 1.2. The case study has been implemented in the sense that all the access to diverse knowledge sources is done through queries and then integrated into the INN.

#### 3.1. Queries to knowledge sources

We formalised the set of narrative questions raised during art interpretation with a top-down approach from the didactic art literature<sup>6</sup> [13]. The suggested questions were grouped as sub-questions of five wh-questions, namely *Who? What? Where? When? Why?*. In terms of the INN structure, the art interpretation frame raised the 5Wquestions, which raise further more detailed questions which can be expressed as SPARQL queries to be performed over an initial selection of Knowledge Bases (KBs) (illustrated in table 1). For example to answer the narrative question *?Who* (i.e. who is the painter?) the following SPARQL query can be launched to the Wikidata KB<sup>7</sup>: `SELECT ?author ?authorLabel WHERE {wd:Q4009580 wdt:P170 ?author. ?author rdfs:label ?authorLabel}`. The answer from a query (if successful) is then integrated as

---

<sup>5</sup>Our paper is supported by an open-source software implementation of integrative narrative networks for Art Interpretation (Apache 2.0 license) as part of the Babel cognitive software suite [11],[12] which is available at <https://github.com/SonyCSLParis/art-network>.

<sup>6</sup>University of Manchester, Art historical research. Url: [https://library.wcupa.edu/\\_art\\_history](https://library.wcupa.edu/_art_history)

<sup>7</sup>The SPARQL queries performed are available at [https://inn-painting-art-interpretation.streamlit.app/Narrative\\_Questions\\_and\\_Queries](https://inn-painting-art-interpretation.streamlit.app/Narrative_Questions_and_Queries)

**Table 1**

Knowledge bases included in the case study.

Knowledge Base	Description
Wikidata[14]	Open, collaborative KB hosting structured data from other Wikimedia projects
Arco[15]	Italian Cultural Heritage
Zeri&Lode[16]	The photographs about XVI Century art of Federico Zeri's photographic collection
The Iconology Dataset[17]	Iconological interpretations about ca 400 artworks, mainly by the art historian Erwin Panofsky
HyperReal[18]	Symbols and symbolic meanings in different contexts

an answer bound to the question ?Who in the INN, and the colour of the variable is changed to green. The criteria for selecting KBs is the possible presence of data about artworks, their content, and/or possible meanings. Furthermore, results obtained from computer vision analysis were integrated into the Narrative Network through a test on Google API Cloud Vision<sup>8</sup>.

More concretely, we searched for the painting ID on the KBs through SPARQL queries looking for the artist's name and title. We then performed the queries expressing the 5W questions on each KB. During the process, results were represented according to the INN data structure, showing a graph which progressively expanded from the initial artwork node. We visually distinguished the sources of the answers using different shades of yellow colour. Nodes retrieved from different KBs which were possibly referring to the same concept (e.g. the representation of the character "Venus" in Wikidata and in Zeri&Lode) were detected through a label similarity fuzzy ratio, and through Wordnet synonyms, and through the presence of common URI to which they were already aligned.

### 3.2. Results

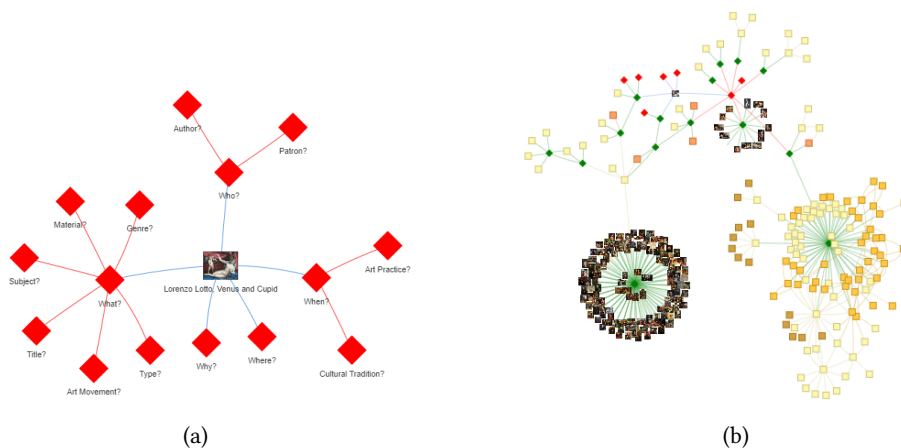
In the initial phase, we retrieved the artwork ID, which was described in 3 of them (Wikidata, the Iconology Dataset, and Zeri&Lode), and we added a starting INN Artwork node to the graph. Following, the 5W narrative questions and their sub-questions were raised<sup>9</sup>. Figure 2(a) illustrates how the network visualisation appears at this stage. The following sections illustrate in detail the network expansion according to each main question.

**?Who** This question concerned the author since no information about other involved people (e.g the patron) was available on the KBs. It was possible to retrieve the author's date of birth and death, places and art movement. Since this information partially answered other open narrative questions, such as ?When and ?Art movement (sub-question of ?What, a relation between them was added. It was also possible to retrieve other paintings by the same author, which is relevant information for art historical research, e.g. for style and subject matter comparison.

**?What** This question included the retrieval of metadata, e.g. the title, type, genre, material, and art movement of the painting. All the sub-questions found an answer on Wikidata and Zeri. We then answered the narrative question '?Subject, sub-question in the ?What variable, which

<sup>8</sup><https://cloud.google.com/vision>

<sup>9</sup>The visualisation of the graph at each stage is available at [https://inn-painting-art-interpretation.streamlit.app/Network\\_Evolution](https://inn-painting-art-interpretation.streamlit.app/Network_Evolution)



**Figure 2:** How the network appears (a) with the 5WH questions and sub-questions still unbound, and (b) at the end of the information retrieval.

retrieved a great amount of information. The graph is enriched with the subjects depicted, and some of them were aligned through the reference to a common Iconclass identifier (92C454 for Venus and Cupid). The Iconology dataset provided a thorough description of the subjects according to three layers of understanding. Also, the results obtained by CV analysis were integrated, despite only the concept of "person" being recognised, and the algorithm erroneously recognising Venus's diadem as a hat. Following, we retrieved the sources of the painting (i.e. an epithalamium, according to Wikidata) and of various objects, e.g. the incense burner described as a proper bridal chamber decoration by the Roman poet Sidonius (source: Iconology Dataset).

**?When** The retrieval of information about the artist partially answered the question, restricting the range of the possible date of the artwork creation to the one of the author's life. Nevertheless, a more precise date was provided by Wikidata (1530) and Zeri (1520-1556).

**?Where** The main questions raised are 1) where the artwork was created, and 2) where the artwork was intended to be displayed. Whereas they could be formally expressed as queries, no further information was available for Lotto's painting. The question is only partially answered by the information about the author's work location previously retrieved in the ?Who question. We consequently added a relation between the narrative questions ?Location and ?Where.

**?Why** The narrative question ?Why had as a main sub-question the retrieval of the patron, namely the person who commissioned the artwork. Although this sub-question can be formally expressed (Wikidata's relation `wdt:P88`), such information is not available for the chosen case study. The narrative question remains unbound.

**Second iteration** We continued the exploration of further knowledge in which a potential user could be interested, discovering further connections or information about a retrieved node. For example, we retrieved more information about the depicted characters and their role, and we looked for objects' potentially embedded symbolism, not yet discovered, by querying HyperReal, discovering, for example, that the conch is a symbol of fertility, a concept supporting the meaning of the artwork as a wedding wish.

### **3.3. Utilisation of the INN by art historians**

Among the possible applications that can be developed around the INN data structure, we operationalised one application intended for art historians. It allows them to direct the expansion of the INN by indicating which questions should be preferentially explored and by choosing which knowledge sources should be utilised preferentially. Figure 1 shows a snapshot of this interface. The user-centered implementation of INN data structure helps resolve conflicting or erroneous information from various sources, as the user can remove erroneous information from his/her own network. Future implementation considers providing feedback to KBs about the correctness of the retrieved data and integrating the newly added experts' knowledge into collaborative graphs (e.g. Wikidata) to improve data quality.

### **3.4. Discussion**

By implementing the questions raised by the literature, the major part of the questions could be expressed as SPARQL queries, performed and answered with currently available information on knowledge graphs. Some answers were bound to multiple questions, as they answered both. Despite multiple information could be aligned through the strategies of reconciliation with a common vocabulary, word similarity, and synonyms detection, challenges in information integration remain open (e.g. when pieces of information slightly differ, as happened with the date of creation). As a solution, we make the INN available as an interactive tool, in which the user can manually select the desired data and perform reconciliations.

## **4. Conclusions and future work**

This paper introduced an application to the art history domain of Integrative Narrative Networks (INN), a data structure for representing progress in interpreting artworks in the form of a graph that connects questions and answers for a specific work. Narrative closure occurs when the main questions of relevance to the human interpreter have been answered. The paper defined the INN and focused on how semantic web resources about art history and interpretation, which are becoming more and more available, can be marshalled to push the interpretation forward. We used a case study of a painting by Lorenzo Lotto to illustrate the proposed methods and techniques. In future work, we plan to integrate additional resources, conduct more case studies, expand the capabilities of the user interface for art historians with the aid of a user study, make advancements toward the generation of KB-independent queries, and develop another interface embedded in an augmented reality device so that viewers seeing an artwork in situ can also interactively explore the semi-autonomous expansion of the Integrative Narrative Network.

## **Acknowledgments**

This research was supported by the EU Pathfinder project MUHAI (EU grant 951846) and by the international exchange program of Bologna University.

## References

- [1] M. Heidegger, *Being and Time*, SCM Press, 1962.
- [2] H.-G. Gadamer, *Hermeneutics and social science*, *Cultural Hermeneutics* 2(4) (1975) 207–316.
- [3] N. Carroll, *Narrative closure*, *Philosophical Studies*. 135 (2007) 1–15.
- [4] R. van Trijp, I. Blin, *Narratives in historical sciences*, in: L. Steels (Ed.), *Foundations for Incorporating Meaning and Understanding in Human-centric AI*, “MUHAI consortium”, 2022.
- [5] S. Aslan, L. Steels, *Aligning Figurative Paintings With Their Sources for Semantic Interpretation*, *IJIMAI In Press* (2023).
- [6] K. Christiansen, *Lorenzo Lotto and the tradition of epithalamic paintings*, *Apollo* 124 (1986) 166–173.
- [7] M. Minsky, *A framework for representing knowledge*, in: P. Winston (Ed.), *The Psychology of Computer Vision*, McGraw-Hill, New York, 1975, pp. 211–277.
- [8] D. G. Bobrow, T. Winograd, *An overview of KRL, a knowledge representation language*, *Cognitive Science* 1 (1977) 3–46. doi:10.1016/S0364-0213(77)80003-7.
- [9] L. Steels, *Design requirements for knowledge representation systems*, in: J. Laubsch (Ed.), *GWAI-84*, Springer Berlin Heidelberg, Berlin, Heidelberg, 1985, pp. 1–19.
- [10] R. P. Gabriel, J. L. White, D. G. Bobrow, *CLOS: Integrating object-oriented and functional programming*, *Commun. ACM* 34 (1991) 29–38. doi:10.1145/114669.114671.
- [11] M. Loetzsch, P. Wellens, J. De Beule, J. Bleys, R. van Trijp, *The Babel2 Manual*, Technical Report AI-Memo 01-08, AI-Lab VUB, Brussels, 2008.
- [12] L. Steels, M. Loetzsch, *Babel: A tool for running experiments on the evolution of language*, in: S. Nolfi, M. Mirolli (Eds.), *Evolution of Communication and Language in Embodied Agents*, Springer, Berlin, Germany, 2010, pp. 307–313.
- [13] T. Barrett, *Interpreting art: Reflecting, wondering, and responding*, McGraw-Hill New York, 2003.
- [14] D. Vrandečić, M. Krötzsch, *Wikidata: A free collaborative knowledgebase*, *Commun. ACM* 57 (2014) 78–85. doi:10.1145/2629489.
- [15] V. A. Carriero, A. Gangemi, M. L. Mancinelli, L. Marinucci, A. G. Nuzzolese, V. Presutti, C. Veninata, *Arco: The italian cultural heritage knowledge graph*, in: C. Ghidini, O. Hartig, M. Maleshkova, V. Svátek, I. Cruz, A. Hogan, J. Song, M. Lefrançois, F. Gandon (Eds.), *The Semantic Web – ISWC 2019*, Springer International Publishing, Cham, 2019, pp. 36–52.
- [16] M. Daquino, F. Mambelli, S. Peroni, F. Tomasi, F. Vitali, *Enhancing semantic expressivity in the cultural heritage domain: Exposing the zeri photo archive as linked open data*, *J. Comput. Cult. Herit.* 10 (2017). doi:10.1145/3051487.
- [17] S. Baroncini, M. Daquino, F. Tomasi, *Exploring iconographical and iconological content in semantic data through art historians’ interpretations*, in: *Atti del XII convegno annuale AIUCD*, Siena, 2023, pp. 9–14.
- [18] B. Sartini, M. van Erp, A. Gangemi, *Marriage is a peach and a chalice: Modelling cultural symbolism on the semantic web*, in: *K-CAP ’21: Proceedings of the 11th on Knowledge Capture Conference*, ACM, New York, 2021, pp. 201–208. doi:https://doi.org/10.1145/3460210.3493552.