Conceptual Encoding and Advanced Management of Leonardo da Vinci's Mona Lisa: Preliminary Results

Alessia Amelio [1] and **Gian Piero Zarri** [2,*]

[1] Dipartimento di Ingegneria Informatica, Modellistica, Elettronica e Sistemistica, University of Calabria, 87036 Rende (CS), Italy; aamelio@dimes.unical.it

[2] STIH Laboratory, Sorbonne University, 75005 Paris, France

* Correspondence: gianpzarri@gmail.com

Abstract: This paper describes a preliminary experiment concerning the use of advanced Artificial Intelligence/Knowledge Representation techniques to improve the present formalization/digitization procedures of Cultural Heritage assets—with reference, in particular, to all types of Cultural Heritage "iconographic" entities. In this context, in agreement with the recent proposal to characterize the digital description of Cultural Heritage items making use of the notion of "Cultural Heritage Digital Twin", we are mainly concerned with the possibility to consider not only the external, "physical", aspects of these iconographic items but also the "message" they convey in a more or less explicit way. For our experiment, some aspects of the Mona Lisa painting by Leonardo da Vinci have been formalized, along with their context, making use of NKRL, the "Narrative Knowledge Representation Language". NKRL is, in reality, both a Knowledge Representation language and a full Computer Science environment, used to represent/manage in an advanced way "narrative" (in the widest meaning of this word) information. The initial results of the experiment are described in the paper, along with some thoughts about their possible interest and developments.

Keywords: ontology; semantic web; pattern recognition; Leonardo da Vinci; Mona Lisa

1. Introduction

What is illustrated in this paper about the possibility of experimenting with new techniques of knowledge representation and querying/inferencing in the Cultural Heritage domain seems to be echoed in some recent wishes recently expressed by the European Commission. For example, in the description of the specific 2020 challenge associated with the topic "DT-TRANSFORMATIONS-12-2018–2020: Curation of digital assets and advanced digitisation" of Horizon 2020, the Commission says that, "*So far, digitisation focused mainly on capturing the visual appearance of individual objects, collections or sites. There is a real need to establish a comprehensive picture of the studied assets, capturing and re-creating not only visual and structural information, but also stories and experiences (stored in language data), together with their cultural and socio-historical context, as well as their evolution over time*" [1], p. 41. This should mean that the new generation of digitization procedures must focus not only on the accurate reproduction of the "physical" characteristics of the Cultural Heritage entities but also on the modelling of the "*message*" these entities convey in a more or less explicit way along with their historic, social and cultural background. Note that, for the sake of simplicity, we englobe under the term "physical" all those "features/properties" normally used to characterize/identify a Cultural Heritage item, i.e., dimensions, support, execution technique, style, name of the artist, information about the owner, places and collections the item is or was in the past, etc. A recent interesting proposal has been put forward in this context, i.e., to make use of the notion of "*Cultural Heritage Digital Twin*"—derived from that of Digital Twin recently introduced in a software engineering setting [2]—to characterize

the digital description that should be ideally associated with every Cultural Heritage item. One of the two components of the twin should then be devoted to the "physical" (according to the above general meaning of this word) description of the item, while the second should be used to represent the associated, "intangible (immaterial) message" transmitted by this item.

Within this general framework, this paper presents some suggestions about the use of advanced Artificial Intelligence/Knowledge Representation techniques to improve the present formalization/digitization procedures of Cultural Heritage assets—with reference, in particular, to all types of Cultural Heritage "iconographic" entities. The Mona Lisa ("La Gioconda") painting—or better, as usual in the context of the computerized enquiries about Cultural Heritage items, some of the multiple photographic records about this work of art—has been selected for a first, preliminary experiment. According to the above Digital Twin option, the digital description of this picture should go beyond the simple account of its "physical" features and immediate documentary and historical background. It should allow for an (at least partial) digital representation of the "immaterial" component of the associated twin, in particular, of those *emotional factors* that have greatly contributed to the universal fame of this picture. Although useful insights in this context have already been provided from a "visual computing" perspective (see [3]) no sufficiently exhaustive (and interoperable) computer-suitable descriptions of these emotional factors seem to be concretely available. Thanks to its ability to represent in a simple but rigorous and efficient way complex situations and events, behaviours, attitudes, etc. (see Section 2.1 below), the "Narrative Knowledge Representation Language", NKRL [4] has been chosen for representing the Mona Lisa twin and its immaterial component in particular.

The paper is organized as follows. Section 2 supplies the general background of this work, introducing first a short state of the art about the use of digitalized procedures in a Cultural Heritage context and recalling then the main features of NKRL. Section 3 is the main Section of the paper, showing how NKRL could be used to represent appropriately, in computer suitable form, a sample of those "hypotheses" [3] that have flourished about the origin and the meaning of "La Gioconda". A short conclusion, Section 4, closes the paper.

2. The General Background

2.1. Short Discussion about the Cultural Heritage Digitization Procedures

These procedures have now a long tradition of work that goes back to the sixties. They were limited, initially, to some elementary descriptions of the Cultural Heritage items that could be likened, in practice, to standard annotations created making use of keywords (metadata) extracted from thesauri as the Art and Architecture Thesaurus (AAT, http://www.getty.edu/research/tools/vocabularies/aat/index.html), ICONCLASS (http://www.iconclass.nl) or the Union List of Artist Names, ULAN (http://www.getty.edu/research/tools/vocabularies/ulan/index.html). They were typically focused, then, on the utilization—mainly for documentary/information retrieval purposes—of the "physical" aspects of the Cultural Heritage Digital Twins introduced above. The use of (relatively simple) data models very popular in the eighties-nineties as the original Dublin Core (http://dublincore.org/) proposal—defined as "pidgin metadata" in [5]—or of more sophisticated models like VRA (Visual Resources Association) Core 4 XML Schema (https://www.loc.gov/standards/vracore/schemas.html) has contributed to improve the efficacy and generality of these annotation procedures. The iconographic items descriptions obtained thanks to this kind of processes correspond then with reference, e.g., to the VRA model—to the use of unstructured sequences of binary entities of the property-value type (like "name-Peter Paul Rubens" or "culture-Flemish") linked to simple XML attributes. The conversion of several of these models into RDF-compatible tools under the influence of the Semantic Web initiative—see the DCMI (Dublin Core Metadata Initiative) Abstract Model [6] or the RDF(S) version of the CIDOC CRM tool (http://www.cidoc-crm.org/rdfs/5.0.4/cidoc-crm)—has further increased their interoperability and standardization potential. CIDOC CRM (Conceptual Reference Model)—see [7] for the last version (6.2.6)—is a well-known, powerful and outstanding tool that aims at providing definitions and a

formal, ontological-oriented structure for describing the main implicit and explicit concepts (99 in the last version) and relationships (191) used in the Cultural Heritage documentation. Since 9/12/2006, it is the official standard ISO 21127:2006.

If the adoption of a Semantic Web (RDF) perspective has indubitably produced important beneficial effects with respect to the digital encoding of the "physical" properties of the iconographic items, it does not seem to have introduced real progresses in the search for concrete solutions for representing the second, "immaterial" size of the Cultural Heritage Digital Twins. A canonical and well-known example in this context is represented by the RDF-based description of Claude Monet's Garden at Sainte-Addressee painting included in the Image Annotation on the Semantic Web W3C Incubator Group Report [8]. This description includes an impressive amount of documentary/bibliographic/descriptive/ "physical" details. However, the description of the proper "content/deep meaning" of the painting and of the "message" it could transmit to the observer is reduced to three extremely sketchy statements—e.g., <vra:subject>Adolphe Monet (artist's father)</vra:subject>—relating the presence on the picture of three persons (three Claude Monet's relatives). No information is given about the real semantic content of the painting, for example, the mode of sitting of the personages in front of the see, their mutual relationships, their attitude about the peaceful and bright landscape, the impressive number of flowers, etc. The Incubator Group Report goes back to 2007, and we could suppose that the state of the art about the formal description of iconographic items described in this report is now largely outdated. This is not always the case. We can see for example, in the context of the very complex and convoluted formalization of the Mona Lisa painting carried out in 2013 in an Europeana context [9] making use of the RDF-oriented European Data Model (EDM), that the description of the semantic content of this painting seems to be several times reduced to flat statements in the style of "the subject of the painting is a woman" [9], p. 19.

A number of proposals concerning the implementation of specialised systems for the description/management of iconographic items that are based on extensions of CIDOC CRM have been advanced these last years. In [10] for example, the authors suggest to expand CIDOC CRM making use of the Situations&Descriptions (S&D) module of the DOLCE Semantic Web tool. In this context, CIDOC CRM should be used to represent the properties and the basic relationships of these items, while the S&D module should supply an in-depth formal description of their "informative content"; the iconographic domain of application concerns some temple scenes and reliefs pertaining to the Meroitic civilization in Sudan. From the paper, it appears clearly through some concrete examples how CIDOC CRM, in association with other ontologies, could be employed to formalise the relationships between temples, scenes and reliefs, highlighting in particular some spatial relationships. On the contrary, the contribution of S&D to the formalization of the informative content is not concretely illustrated in the paper and seems, therefore, to belong more to the domain of intent than to that of tangible results—even if the authors mention, in the "Conclusion", the existence of some undefined "successful tests". Another recent proposal of extension of CIDOC CRM concerns the implementation of an ontological tool called VIR (Visual Representation) [11]—see also https://ncarboni.github.io/vir/. VIR expands the key entities and properties of CIDOC CRM by introducing seven new classes (e.g., "Iconographical Atom") and 20 new relationships (e.g., "portray/is portrayed") corresponding to the needs of the visual and art historical community. According to the authors, the resulting model is able to provide a clear distinction between "denotation" and "signification" of a Cultural Heritage element, permitting, in particular, the definition of diverse denotative criteria for the same representation. Taking as an example the representation of a widely known iconographical character, Saint George, the VIR-supported ontological/graphical representations allow us to follow the evolution in time of the set of attributes illustrating this character. These are simply reduced to "horse" and "spear" in the frescos of the Panagia Phorbiotissa church in Cyprus, to which are added "Castle, Princess, Lake and Dragon" when the "Saint George killing the dragon" painting by Vittore Carpaccio is taken into consideration. This work is of course interesting, but it contributes scarcely to the representation of the "deep meaning" of the iconographic items examined. For example, the modelling of the Carpaccio's

picture is reduced to a simple "binary" representation where an individual, Saint George, is denoted by a set of static properties, horse, spear, castle etc., and the highly complex, dramatic and dynamic illustration of the fight with the dragon is totally ignored.

We can conclude this short discussion about the present state of the art with respect to the modelling of the "inner meaning" of Cultural Items by mentioning a recent paper [12], which concerns the project Cultural-ON (Cultural Ontology) resulting from a collaboration between the Italian Ministry for Cultural Heritage and the Italian National Council for Research. The project has been recently restructured (see http://dati.beniculturali.it/cultural_on/to) to adapt it to other "ontological" initiatives of the Ministry, but its basic approach did not change. In this paper we find, e.g., the formal description of a Giambologna's sculpture denoting the kidnapping of a Sabine woman by a Roman warrior, on display at the Capodimonte museum in Naples in the framework of the so-called "Collezione Farnese". The description of this sculpture is restricted to purely documentary/information retrieval-oriented statements like cis: isMemberOf collection: Collezione_Farnese or cis: isInSite: Sede_del_Museo_di_Capodimonte, without any attempt to model the frightened woman, the kidnapper, the man who tries to prevent the kidnapping, the reciprocal attitudes of the different protagonists of the scene, etc.

Note, eventually, that finding a solution to the above difficulties is not facilitated by the choice of systematically choosing RDF as implementation support, because of the well-known "lack of expressiveness" problem that bothers all the Semantic Web (SW) tools from the beginning. This problem is linked to the choice of making use only of the quite limited "binary" knowledge representation model. In this, in fact, a concept is simply defined through a set of properties/attributes; when the concept is instantiated into a concrete individual, each associated property can only link this individual to another individual or a value, individual1-property-individual2/value. Such an approach renders then particularly difficult the setting up of complete and effective formal descriptions of real-world, complex information structures like spatio-temporal data, contexts, reified situations, human intentions and behaviours, etc. The evident solution to this problem is that of making use, instead of binary models, of n-ary ones where a given predicate can be associated with multiple arguments—for example, the n-ary "purchase" relation/property concerns events where at least a seller, a buyer, a good, a price, and a timestamp are involved. Unfortunately, it is not so evident that an n-ary solution should be used, given the associated implementation difficulties. To give only an example, a common misunderstanding consists in asserting that the use of specific n-ary knowledge representation structures is not necessary given that any n-ary relationship can be simply reduced to a set of binary relationships. This is in principle true, and this sort of decomposition can be also useful for, e.g., dealing with very practical problems like storing efficiently n-ary relationships into standard databases. However, binary and n-ary relationships are conceptually irreconcilable, and an n-ary relation cannot be reduced to the simple addition of binary elements without losing its "deep meaning"—it is impossible to reason about, e.g., the possible reasons and consequences, the context, etc., of a purchase without considering the purchase event in its whole conceptual entirety, i.e., by taking all its arguments simultaneously into account.

2.2. Using NKRL in the Context of the Cultural Heritage Domain

NKRL is both a "structured n-ary" (see below) knowledge representation language and a fully implemented [4], computer science environment built up thanks to several European projects. It has been successfully used in many different narrative-based domains like terrorism news, analysis of industrial incidents, sentiment analysis, conceptual IoT, etc.

The most important NKRL features are highlighted below.

From an ontological point of view, the most striking characteristic of the language concerns the addition of an "ontology of elementary events" to the usual ontology of concepts. In NKRL this last—called "HClass", hierarchy of classes—presents some interesting aspects, with respect in particular to the modelling of difficult notions like colour and substance [4], pp. 123–137. However,

its architecture is a traditional one, and the HClass concepts are represented according to the usual "binary" model.

The nodes of the ontology of events are represented, on the contrary, by n-ary knowledge patterns—called "templates" in an NKRL context—that denote formally general classes of elementary events/states/situations/actions/episodes/experiences. Examples of these general classes can "be present in a place", "move a physical object", "have a specific attitude towards someone/something", "send/receive messages", "be characterized by a given property", etc. The instances of these templates—called "predicative occurrences" in an NKRL context—describe then formally the "meaning/inner content" of specific elementary events/states/situations, etc., pertaining to one of these classes.

A "conceptual predicate" denotes the main structuring element of a specific template or predicative occurrence; it defines the general semantic class to which the template/occurrence pertains. Seven conceptual predicates are used in NKRL: BEHAVE, EXIST, EXPERIENCE, MOVE, OWN, PRODUCE, RECEIVE; see [4], pp. 57–59 for the reasons of this reduced choice.

The different entities (human being, physical objects, but also properties, locations, events, situations etc.) involved in the original elementary events correspond, within templates/occurrences, to the n-ary "arguments of the predicate". These arguments are created by using concepts or instances of concepts (individuals) proper to the HClass hierarchy; the two NKRL ontologies are then formally/functionally different, but strictly integrated from an operational point of view.

"Structured arguments (expansions)" under the form of lists of HClass terms can be created making use of the four operators of the so-called AECS sub-language; see [4], pp. 68–70. AECS includes the disjunctive operator ALTERN(ative) = A, the distributive operator ENUM(eration) = E, the collective operator COORD(ination) = C and the attributive operator SPECIF(ication)—their interweaving within a structured argument is controlled by a "priority rule" that forbids, e.g., to use a list of the ALTERN type within the scope of a list COORD.

To specify exactly the "function/role" of each argument in the context of a specific template/occurrence, these arguments are introduced by the following functional roles: SUBJ(ect), OBJ(ect), SOURCE, BEN(e)F(iciary), MODAL(ity), TOPIC, CONTEXT [13]. They individuate then, e.g., the main protagonist (SUBJ(ect)) of an elementary event, the beneficiary (BEN(e)F(iciary)) or the origin (SOURCE) of a given action, etc.

The whole n-ary expression, template or occurrence, is then unequivocally identified making use of a "symbolic label" that, through special reification mechanisms, allows this template/occurrence to be associated to other templates/occurrences in order to allow, in the end, the representation of complex events/scripts/scenarios, etc. These reification mechanisms are implemented according to two main techniques, the "binding occurrences" (see Table 1) and the "completive construction" (see the NKRL clause denoted as gio3.c13 in Table 3).

"Determiners" can be added to templates or predicative occurrences to introduce further details about the basic core of their formal representation, see Equation (1) below. For example, determiners of the "location" type can be associated through the colon operator, ":", with the arguments of the predicate introduced by the SUBJ, OBJ, SOURCE and BENF functional roles. Another important category of determiners associated, in this case, to a full template or predicative occurrence concerns some constants called "modulators", for example, the temporal modulators (begin, end, obs(erve)). These last, associated with the determiners date-1 and date-2, are of a particular importance in the framework of the temporal representation system of NKRL; a recent paper on this subject is [14].

Eventually, "querying/reasoning" in NKRL ranges from the direct questioning of NKRL knowledge bases making use of "search patterns" (NKRL equivalents of natural language queries) to high-level inference procedures [15]. In this last context, e.g., the "transformation rules" try to "adapt" the original query/queries (search patterns) that failed to the real contents of the KB. This means using rules to automatically "transform" the original query into one or more different queries that are not strictly "equivalent" but only "semantically close" to the original one making use of a sort of "analogical"

reasoning. "Hypothesis rules" allow us to build up realistic answers according to a number of reasoning schemata, e.g., "causal" ones.

Equation (1) denotes the "core" of the formal representation of an NKRL template/predicative occurrence:

$$\text{(Li (Pj (R1 a1) (R2 a2) ... (Rn an))).} \tag{1}$$

In Equation (1), Li is the symbolic label identifying (reifying) the whole template/occurrence, Pj is the conceptual predicate, Rk is a generic functional role and ak is a (simple or structured) predicate argument introduced by the role Rk.

3. The Experiment

3.1. A First (Quite Simple) Result

Table 1 supplies the NKRL representation of a fragment of the whole formal encoding of the Mona Lisa picture, which is presented here, for illustrative purposes, under the form of an independent "narrative". This last could be conceived as a didactic description, in a pure documentary style, of basic information about this picture. The formal representation into NKRL terms presented in this table and in the following ones is basically a "manual" type of encoding facilitated, however, by the use of a powerful editor included in the standard NKRL software; see in this context [4].

In Table 1, gio1.c1 and gio1.c3 denote "binding occurrences" (see [4], pp. 91–98), different then from the "predicative" ones like gio1.c2, gio1.c5, gio1.c6 etc., which represent elementary events and are structured according to the general syntax of Equation (1). The binding occurrences, as already mentioned, represent one of the two modalities used in NKRL for grouping together several predicative occurrences. Their general formal expression is (Lbk (Bni L1 L2 ... Ln)), where Lbk is the symbolic label (e.g., gio1.c1 in Table 1) identifying the binding structure. The Bnj operators are: ALTERN(ative), COORD(ination), ENUM(eration), CAUSE, REFER(ence), the weak causality operator, GOAL, MOTIV(ation), the weak intentionality operator, COND(ition).

As usual in an ontological context, the ontology of concepts (HClass in our case) must be updated whenever a new application must be activated—in the tables of this paper, concepts are in lower case and individuals in upper case. This process did not pose particular problems for the set-up of Table 1 given that most of the concepts utilized there are standard HClass concepts (identified_with, a concept pertaining, as labelled_as, to the binary_relational_property sub-hierarchy of HClass, wife_, a concept of extended_family-role, etc.) and some "new" concepts have been created simply as specialization of existing concepts (florentine_merchant IsA merchant_). The only really "new" concepts are those specifically linked to the Cultural Heritage domain, like painted_on and poplar_plank. Table 1 illustrates the importance of using structured arguments/expansions (see Section 3.1), instead of simple HClass concepts/individuals, as fillers of functional roles in the predicative occurrences. SPECIF(ication) is used, in particular, to better "specify" the HClass term representing the first element (the "head") of a SPECIF list, as portrait_ (gio1.c2) or FRANCESCO_DEL_GIOCONDO (gio1.c6), by adding to it further explaining terms. These lists, as all the AECS lists, can be used in a recursive way see, e.g., gio1.c4 and gio1.c5—note, in this last occurrence, the use of (SPECIF cardinality_ several_), the standard NKRL expression to denote plural entities.

None of the (about 150) HTemp templates needed some particular update/customization procedures to be used for the setup of Table 1. Table 2 reproduces the template Behave:Kinship that gives rise to gio1.c6 in Table 1.

Table 1. Formalization of basic information about the Mona Lisa picture under the form of a simple "narrative".

gio1.c1)		(COORD gio1.c2 gio1.c3 gio1.c4)	
The formalization of the gio1.c1 *narrative is represented by three predicative and binding (see below) occurrences.*			
gio1.c2)	EXIST	SUBJ	PAINTING_1: (LOUVRE_MUSEUM)
		TOPIC	(SPECIF portrait_ WOMAN_52)
		date-1:	1/1/1797, 31/12/1797
		date-2:	29/1/2020
Exist:EntityBePresent (2.11)			
A painting concerning (TOPIC) *the portrait of a woman (conventionally:* WOMAN_52) *is located at the Louvre Museum from an (unknown) date in 1797.*			
gio1.c3)		(COORD gio1.c5 gio1.c6)	
The component gio1.c3 *of the narrative is formed of two coordinated predicative occurrences.*			
gio1.c5)	EXPERIENCE	SUBJ	WOMAN_52
		OBJ	(SPECIF identified_with LISA_GHERARDINI)
		SOURCE	(SPECIF art_historian (SPECIF cardinality_ several_))
		{ mult }	
		date-2:	29/1/2020
Experience:GenericSituation (3.1)			
The woman of the Mona Lisa portrait has been often (modal modulator mult) *identified with Lisa Gherardini.*			
gio1.c6)	BEHAVE	SUBJ	LISA_GHERARDINI: (FLORENCE_)
		MODAL	wife_
		TOPIC	(SPECIF FRANCESCO_DEL_GIOCONDO florentine_merchant)
		{ poss }	
		date-1:	1/1/1495, 31/12/1495
		date-2:	15/7/1542
Behave:Kinship (1.113)			
Lisa Gherardini has been the wife of the Florentine merchant Francesco del Giocondo from an unknown date in 1495 to July 15, 1542 (date of death).			
gio1.c4)	OWN	SUBJ	PAINTING_1: (LOUVRE_MUSEUM)
		OBJ	property_
		TOPIC	(SPECIF painted_on (SPECIF POPLAR_PLANK_1 (SPECIF length_ (SPECIF centimetre_ 53)) (SPECIF height_ (SPECIF centimetre_ 77))))
		{ obs }	
		date-1:	29/1/2020
		date-2:	
Own:CompoundProperty (5.42)			
The painting has been realized on a poplar plank; its dimensions are: length 53 cm, height 77 cm.			

In the templates, the formalism *(Ri) means that the use of the particular role Ri is forbidden, while the components in square brackets may or may not appear in the corresponding predicative occurrences. In a real template, the arguments of the predicate (the ak terms in Equation (1)) are actually represented by variables (vari) with associated constraints. These last are expressed as HClass terms or combinations of these terms and are used to specify the legal sets of concepts/individuals that can be substituted for these variables within the predicative occurrences—see, e.g., the correspondence between the constraints of the template of Table 2 and the terms used as arguments of the predicate in gio1.c7.

Table 2. The Behave:Kinship template.

name: Behave:Kinship		
father: Behave:Role		
position: 1.113		
natural language description: "An Individual Has a Given Kinship Relationship with a Human Being Filling the TOPIC Role"		
BEHAVE	SUBJ	*var1*: [(*var2*)]
	*(OBJ)	
	[SOURCE	*var3*: [*var4*]]
	(BENF)	
	MODAL	*var5*
	TOPIC	*var6*
	[CONTEXT	*var7*]
	{ [modulators] }	
var1 =	human_being	
var1 =	individual_	
var3 =	human_being_or_social_body	
var4 =	human_being_or_social_body	
var5 =	extended_family_role	
var6 =	human_being_or_social_body	
var7 =	situation_, symbolic_label	
var2, *var4* =	location_	

3.2. A More Advanced Result: Isabella d'Aragona's Hypothesis

The few predicative occurrences of Table 1 can already be used to ask some simple questions in a Mona Lisa context through the use of "search patterns" pi. These are, in practice, partially instantiated templates that correspond to formal queries posed by the user to a knowledge base of NKRL narratives but that can also be generated automatically by the system during the execution of high-level inference procedures; see [4], pp. 183–202. Examples of these questions could be "Who was Lisa Gherardini's spouse?", "What are the dimensions of the painting?", "Where is the painting presently located?", etc. They can present some interest but, being so simple, do not justify by themselves the use of NKRL.

We have already mentioned the proliferation, well described in [3], of all the possible "hypotheses" concerning the origin, the development, the hidden meaning, etc., of the La Gioconda painting. For example, a problem that has troubled scholars for a long time concerns the identification of the woman represented in the "hidden painting", i.e., the portrait, visible to us only in x-rays, indubitably painted by Leonardo and that lies beneath Mona Lisa on the same poplar panel. According, among others, to Lillian Feldmann Schwartz [16] (see also Amelio [3], p. 77), the woman on the hidden painting has nothing to do with Mona Lisa and represents probably Isabella d'Aragona, wife of the Duke of Milan Gian Galeazzo Sforza—at the end of the 15th century, Leonardo had worked for the Sforza family. Feldmann Schwartz's hypothesis is based, essentially, on the lack of correspondence between the facial characteristics of the hidden painting woman and Mona Lisa. A correspondence exists, on the contrary, between this woman and that represented on a Leonardo's "cartoon", a (very retouched) preparatory study (a charcoal and graphite drawing on laid paper) for a portrait of Isabella d'Aragona. Specifically, the eyes of the woman on the hidden painting were "lined up with the eyes on

the cartoon, the hairlines fell into place and overlapped exactly" [16]. Moreover, it was observed that eyes, mouth, nose tips, hairlines and chins did not match between Mona Lisa and the cartoon. In fact, the subject in the original cartoon has a "rounder jaw, a broader nose, lower dress, graver expression, and eyes which look directly at the viewer" [16] than Mona Lisa's subject. Table 3 reproduces the NKRL "narrative", gio2, that formalizes (only partly, for reasons of space) the Isabella d'Aragona hypothesis.

We can remark about Table 3:

The gio2.c7 and gio2.c8 predicative occurrences are intimately associated in the framework of a "completive construction" (symbol #). This construction [4], pp. 87–91 consists in using as filler of a functional role in a predicative occurrence pci the symbolic label Lj of another (predicative or binding) occurrence cj. Only the OBJ, MODAL, TOPIC and CONTEXT roles of pci can accept as filler the symbolic label Lj of a cj in a completive construction context, and only one symbolic label of the #Lj type can appear within a predicative occurrence. In NKRL, the completive construction mechanism is the privileged vehicle for denoting any sort of "transmission of information", see gio2.c7.

Occurrence gio2.c6 is a typical example of formalization of the notion of "negation" in an NKRL context. This consists in adding to a complete and well-formed predicative occurrence pci a specific "modal modulator", negv ("negated event"), to point out that the corresponding elementary event did not take place—in our example, WOMAN_53 has not been recognized as WOMAN_52. Modulators represent an important category of NKRL operators that apply to a template or predicative occurrence to particularize its meaning [4], pp. 71–86. Besides negv, other modulators utilized in this paper are, e.g., obs(erve)—a given elementary event is observed to occur at a given date (conventionally, here, the date of the VIPERC workshop, 29 January 2020)—mult(iple), poss(ible), for, etc.

Eventually, we can also note that the logical arrangement of a generic narrative (like, e.g., that of Table 3) can always be represented as some sort of complex tree structure (a "knowledge graph" then), see Figure 1.
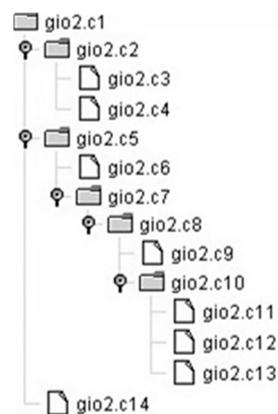

Figure 1. Tree-shaped representation of the Table 3 formalism.

Table 3. Formalization of the Isabella d'Aragona hypothesis.

gio2.c1)	(COORD gio2.c2 gio2.c5 gio2.c14)		
The top level of the gio2 *narrative is represented by three occurrences, two binding occurrences and a predicative one.*			
gio2.c2)	(COORD gio2.c3 gio2.c4)		
The first component of the narrative consists of two predicative occurrences.			
gio2.c3)	PRODUCE	SUBJ	LEONARDO_DA_VINCI
		OBJ	PAINTING_2
		TOPIC	(SPECIF portrait_ WOMAN_53)
		date-1:	1/1/1497, 31/12/1503
		date-2:	
Produce:Entity (6.2)			
The portrait of a woman (conventionally: WOMAN_52) has been produced by Leonardo in the temporal interval 1497–1503.			
gio2.c4)	OWN	SUBJ	PAINTING_2
		OBJ	property_
		TOPIC	(COORD1 (SPECIF painted_on POPLAR_PLANK_1 (SPECIF under_ PAINTING_1)) (SPECIF labelled_as HIDDEN_PAINTING))
		{ obs }	
		date-1:	29/1/2020
		date-2:	
Own:CompoundProperty (5.42)			
PAINTING_2 has been painted under the plank of PAINTING_1, and is known as the HIDDEN_PAINTING.			
gio2.c5)	(CAUSE gio2.c6 gio2.c7 #gio2.c8)		
The elementary event modelized by the predicative occurrence gio2.c6 *has been caused by what is collectively described in the completive construction involving occurrences* gio2.c7 *and* gio2.c8.			
gio2.c6)	BEHAVE	SUBJ	(SPECIF WOMAN_53 (SPECIF identified_with WOMAN_52))
		{ obs, negv }	
		date-1:	29/1/2020
		date-2:	
The elementary event represented by gio2.c7 is a "negated event" (modulator negv), i.e., WOMAN_53 is not WOMAN_52.			
gio2.c7)	MOVE	SUBJ	LILLIAN_FELDMANN_SCHWARTZ
		OBJ	#gio2.c8
		MODAL	(SPECIF SCIENTIFIC_PAPER_2 (SPECIF published_on THE_VISUAL_COMPUTER_JOURNAL))
		date-1:	1/1/1988, 31/1/1988
		date-2:	

Table 3. *Cont.*

Lillian Feldmann Schwartz has spread what described in gio2.c8 by means of a paper in "The Visual Computer" Journal.

gio2.c8)	(COORD gio2.c9 gio2.c10)		

The message transmitted by L.F. Schwartz through her paper is formed of two coordinated components.

gio2.c9)	OWN	SUBJ	(COORD1 (SPECIF eye_ (SPECIF cardinality_ 2) WOMAN_53))
		OBJ	property_
		TOPIC	(SPECIF different_from (SPECIF eye_ (SPECIF cardinality_ 2) WOMAN_52))
		MODAL	x_ray_analysis
		{ obs }	
		date-1:	29/1/2020
		date-2:	

Own:CompoundProperty (5.42)

For the sake of simplification, only the dissimilarity between the eyes of WOMAN_52 *and* WOMAN_53 *has been taken into account here. In reality, also the mouth, nose, chin and hairline of the two women are incompatible.*

gio2.c10)	(COORD gio2.c11 gio2.c12 gio2.c13)		

The second component of L.F. Schwartz's message is formed of three predicative occurrences.

gio2.c11)	PRODUCE	SUBJ	LEONARDO_DA_VINCI
		OBJ	CARTOON_DRAWING_1
		TOPIC	(SPECIF portrait_ ISABELLA_D_ARAGONA)
		date-1:	1/1/1490, 31/12/1495
		date-2:	

Produce:Entity (6.2)

In the period 1490–1495, Leonardo da Vinci realized, as a cartoon drawing, the portrait of Isabella d'Aragona.

gio2.c12)	OWN	SUBJ	(SPECIF eye_ (SPECIF cardinality_ 2) ISABELLA_D_ARAGONA): (CARTOON_DRAWING_1)
		OBJ	property_
		TOPIC	(SPECIF coincident_with (SPECIF eye_ (SPECIF cardinality_ 2) WOMAN_53))
		MODAL	x_ray_analysis
		{ obs }	
		date-1:	29/1/2020
		date-2:	

Own:CompoundProperty (5.42)

Table 3. *Cont.*

There is a correspondence between the eyes of Isabella d'Aragona and those of WOMAN_53.			
gio2.c13)	OWN	SUBJ	(SPECIF hairline_ ISABELLA_D_ARAGONA): (CARTOON_DRAWING_1)
		OBJ	property_
		TOPIC	(SPECIF coincident_with (SPECIF hairline_ WOMAN_53))
		MODAL	x_ray_analysis
		{ obs }	
		date-1:	29/1/2020
		date-2:	
Own:CompoundProperty (5.42)			
There is a correspondence between the hairline of WOMAN_53 *and the hairline of* WOMAN_54.			
gio2.c14)	BEHAVE	SUBJ	(SPECIF individual_person art_historian (SPECIF cardinality_ several_))
		OBJ	LILLIAN_FELDMANN_SCHWARTZ
		MODAL	endorsement_
		TOPIC	(SPECIF ISABELLA_D_ARAGONA_HYPOTHESIS LILLIAN_FELDMANN_SCHWARTZ)
		{ for, obs }	
		date-1:	29/1/2020
		date-2:	
Behave:FavourableAttitude (1.311)			
Several art historians agree with Lillian Feldmann Schwartz about her Isabella d'Aragona hypothesis.			

3.3. Approximating Some Erudite Theories about Mona Lisa's "Meaning"

Since its creation, all the details of the Mona Lisa painting have been the object of the most ingenious and erudite interpretations. In particular, the clear imbalance between the two parts of the background landscape—chaotic mountains and flooding river on the right, peaceful flatland on the left—has always excited the creativity of the scholars, who have been tempted to associate this dissymmetry with some particularities of the proper Mona Lisa's portrait. For example, John F. Asus [17] writes that " . . . the background on the left is a metaphor of the external Mona Lisa. The background on the right flows out from the figure and is a metaphor of the inner life . . . Order and conformity characterize the left side and the form of Mona Lisa. Disorder and chaos are associated with the right side, and it is this side that is tied to the interior of Mona Lisa": in particular, the "left boundary between the background and the figure is very sharp. On the right, they merge. The shoulder mantle becomes the bridge and the turbulent vortex of the river deluge". It is clear that we leave here the domain of a pure descriptive endeavour to enter that of a high-level interpretative sphere. This last, to be fully rendered into NKRL terms, would require a long-term effort in collaboration with specialists of the domain, and this is totally off-topic in the context of our (very preliminary) work.

To show, however, that NKRL can be of some help also in the preliminary phases of the construction of these interpretative theories, we have modelized in the best possible way some basic aspects of both the landscape and the Mona Lisa's portrait; see Table 4. These could be then used to create some (preliminary and naive) associations between the two classes of elements on the basis of the (assumed) reactions of Louvre's visitors in face of them.

Assuming now that the Table 4 data have been introduced in an NKRL knowledge base about the Mona Lisa painting, we could ask, following John F. Asmus' line of reasoning, whether it could be possible to retrieve in this base some information about the existence of relationships between the right part of the landscape (LANDSCAPE_1) and the Mona Lisa's portrait. In an NKRL context, the direct questioning of a knowledge base is implemented by means of a "search pattern" pi—as already stated, a partially instantiated template—that tries to match/unify information in the base thanks to a Filtering Unification Module (FUM) [4], pp. 183–201. In our case, pi could be expressed as in Table 5, i.e., like a partial instantiation of the Exist: SituationBePresent template. In this pattern, feeling_ is an HClass term having generic_animate_condition as generic term and, among others, neutral_feeling and emotionally_charged_feeling as specific terms. emotionally_charged_feeling has in turn, as specific terms, positive_feeling and negative_feeling; discomfort_feeling, see Table 4 above, is a specific term of negative_feeling. Looking at Table 4, it appears clearly that the search pattern of Table 5 is unable to find direct unifications with these data. Before giving up, we can try to see if we could find, within the "transformation rules" associated with a hypothetical Mona Lisa NKRL system, a (sufficiently general) rule allowing us to reformulate the question of Table 5 in terms more consistent with those used in the formulation of Table 4.

Table 4. The right landscape and Mona Lisa's behaviour.

gio3.c1)	(COORD gio3.c2 gio3.c3 gio3.c4 gio3.c5)		
The top level of the gio3 *narrative is represented by four occurrences, a predicative occurrence and three binding ones.*			
gio3.c2)	OWN	SUBJ	(SPECIF LANDSCAPE_1 (SPECIF right_hand_part landscape_))
		OBJ	property_
		TOPIC	(SPECIF different_from greatly_ (SPECIF LANDSCAPE_2 (SPECIF left_hand_part landscape_)))
		{ obs }	
		date-1:	29/1/2020
		date-2:	
Own:CompoundProperty (5.42)			
The two parts of the landscape are very different.			
gio3.c3)	(CAUSE gio3.c6 gio3.c7)		
The reason of what is described in occurrence gio3.c6 *is provided in occurrence* gio3.c7.			
gio3.c6)	PRODUCE	SUBJ	(SPECIF physical_aspect LANDSCAPE_2)
		OBJ	neutral_feeling
		BENF	(SPECIF individual_person (SPECIF visitor_ LOUVRE_MUSEUM))
		{ obs }	
		date-1:	29/1/2020
		date-2:	
Produce:CreateCondition/Result (6.4)			
The left part of the landscape generates no particular feelings in the Louvre Museum's visitors.			
gio3.c7)	OWN	SUBJ	LANDSCAPE_2
		OBJ	property
		TOPIC	(SPECIF regular_ greatly_)
		{ obs }	
		date-1:	29/1/2020
		date-2:	
Own:CompoundProperty (5.42)			
The left part of the landscape appears as perfectly regular.			
gio3.c4)	(CAUSE gio3.c8 gio3.c9)		
The reason of what is described in occurrence gio3.c8 *is provided in occurrence* gio3.c9.			
gio3.c8)	PRODUCE	SUBJ	(SPECIF physical_aspect LANDSCAPE_1)
		OBJ	anomaly_feeling
		DEST	(SPECIF individual_person (SPECIF visitor_ LOUVRE_MUSEUM))
		{ obs }	

Table 4. *Cont.*

		date-1:	29/1/2020
		date-2:	

Produce:CreateCondition/Result (6.4)

The vision of the right part of the landscape engenders a feeling of strangeness in the Louvre Museum's visitors.

gio3.c9)	OWN	SUBJ	LANDSCAPE_1
		OBJ	property_
		TOPIC	(COORD1 (SPECIF characterized_by (SPECIF mountain_ (SPECIF cardinality_ several_) wild_uninhabited_)) (SPECIF characterized_by (SPECIF bridge_ (SPECIF flooded_by (SPECIF river_ turbulent _)))))
		{ obs }	
		date-1:	29/1/2020
		date-2:	

Own:CompoundProperty (5.42)

The right part is characterized by the presence of wild and inhabited mountains and by a bridge flooded by a tumultuous river.

gio3.c5)		(CAUSE gio3.c10 gio3.c11)	

The reason of what is described in occurrence gio3.c10 is provided in occurrence gio3.c11.

gio3.c10)	PRODUCE	SUBJ	(SPECIF part_of (SPECIF physical_aspect WOMAN_52))
		OBJ	anomaly_feeling
		DEST	(SPECIF individual_person (SPECIF visitor_ LOUVRE_MUSEUM))
		{ obs }	
		date-1:	29/1/2020
		date-2:	

Produce:CreateCondition/Result (6.4)

Some elements of the physical appearance of WOMAN_52 engender a feeling of strangeness in the Louvre Museum's visitors.

gio3.c11)		(COORD gio3.c12 #gio3.c13 gio3.c14 gio3.c15)	

Three aspects of this physical appearance are noticed here; gio3.c12 + #gio3.c13 = completive construction.

gio1.c12)	EXPERIENCE	SUBJ	WOMAN_52
		OBJ	(SPECIF need_for instinctive_)
		MODAL	#gio.c13
		TOPIC	(SPECIF reinforcement_ (SPECIF self_protection WOMAN_52))
		{ obs }	
		date-1:	29/1/2020
		date-2:	

Experience:NegativeSituation (3.21)

Table 4. *Cont.*

The woman manifests an instinctive demand for the reinforcement of her own protection making use (MODAL) *of what is specified in the* gio.c13 *binding occurrence.*

gio.c13)	(COORD gio3.c16 gio3.c17)		

The modalities of auto-protection of WOMAN_52 *are detailed in two predicative occurrences.*

gio3.c16)	OWN	SUBJ	(SPECIF right_arm WOMAN_52))
		OBJ	property_
		TOPIC	(SPECIF crossing_ foreground_ (SPECIF chest_ WOMAN_52))
		{ obs }	
		date-1:	29/1/2020
		date-2:	

Own:CompoundProperty (5.42)

The right arm of the woman crosses in the foreground her chest.

gio3.c17)	OWN	SUBJ	(SPECIF hand_ (SPECIF right_arm WOMAN_52))
		OBJ	property_
		TOPIC	(SPECIF grabbing_ strong_ (SPECIF left_arm WOMAN_52))
		{ obs }	
		date-1:	29/1/2020
		date-2:	

Own:CompoundProperty (5.42)

The left arm of the woman is clasped by the hand of the right arm.

gio3.c14)	OWN	SUBJ	(SPECIF hand_ (SPECIF cardinality_ 2) WOMAN_52)
		OBJ	property_
		TOPIC	(SPECIF appear_as (SPECIF extraneous_ (SPECIF body_ WOMAN_52)))
		{ obs }	
		date-1:	29/1/2020
		date-2:	

Own:CompoundProperty (5.42)

The hands of WOMAN_52 *seem to be extraneous to her body.*

gio1.c15)	BEHAVE	SUBJ	WOMAN_52
		MODAL	(SPECIF half_smiling slight_ ambiguous_)
		{ obs }	
		date-1:	29/1/2020
		date-2:	

Behave:HumanProperty (1.1)

The woman produces an imperceptible and mysterious half-smile.

Table 5. A search pattern about the possible relationships between the landscape and the portrait.

EXIST
SUBJ: relationship_
TOPIC: (COORD1 (SPECIF feeling_ (SPECIF related_to (physical_appearance LANDSCAPE_1))) (SPECIF feeling_ (SPECIF related_to (physical_appearance WOMAN_52))))
{ obs }
date1: 29/1/2020
date2:

As already stated, transformation rules are used to adapt, from a semantic point of view, a search pattern pi that "failed"—i.e., that was unable to find a unification within an NKRL knowledge base—to the actual contents of this base using a sort of analogical reasoning. This means that transformations try to automatically convert pi into one or more different patterns p1, p2 . . . pn that are not strictly equivalent but only "semantically close" to the original one. A transformation rule that provides an appropriate reformulation of the pattern of Table 5 is represented in Table 6.

A transformation rule can be conceived as a type of generalized "production rule" formed, then, of a left-hand side, the "antecedent"—i.e., a general formulation of the failed question and of all the analogous questions that can find a solution making use of this rule—and of one or more right-hand sides, the "consequent(s)". These provide one or more new search pattern(s) to be substituted for the original one. The rule is activated by unifying (using FUM) the antecedent of the rule and the pi to be transformed. This allows us to (i) verify that the rule represents really a valid framework for the transformation of the original query, and (ii) give some initial values to the variables derived of the rule. In our case, e.g., var1 and var2 take, respectively, the values LANDSCAPE_1 and WOMAN_52. These values are transmitted to the consequent part of the rule and are used—in conformity with the constraints associated with the variables and carrying out, in case, backtracking operations—to build up the new search patterns. For example, the pattern derived from conseq1 will be characterized by the value emotionally_charged_feeling associated with OBJ, allowing then us to avoid matching occurrence gio3.c6 in Table 4 that includes neutral_feeling in the same position. The equality constraints shown in conseq2 impose that the type of reaction in front of the landscape and of particular aspects of the portrait, anomaly_feeling in this case (see Table 4), must be the same; moreover, the "producer" of this reaction must also be the same, a visitor of the museum, or pertaining to the "visitor" category. The two predicative occurrences, gio3.c8, originated from conseq1 and gio3.c10, originated from conseq2, will be provided to the user as "transformed" answers to the original query. A detailed description of the philosophy and implementation of the NKRL inference engines can be found in [4], pp. 201–239.

Table 6. A transformation rule example.

T54: "identical reaction of specific people in front of different visions" transformation.		
antecedent:		
EXIST	SUBJ	relationship_
		(COORD1 (SPECIF feeling_ (SPECIF related_to (SPECIF
	TOPIC	physical_appearance *var1*))) (SPECIF feeling_ (SPECIF related_to
		(SPECIF physical_appearance *var2*)))
var1 =	location_	
var2 =	individual_person	
first consequent schema (*conseq1*):		
PRODUCE	SUBJ	(SPECIF physical_appearance *var1*)
	OBJ	*var3*
	BENF	*var4*
var3 = emotionally_charged_feeling		
var4 = individual_person		
var4 ≠ *var2*		
second consequent schema (*conseq2*):		
PRODUCE	SUBJ	(SPECIF physical_appearance *var2*)
	OBJ	*var5*
	BENF	*var6*
var5 = emotionally_charged_feeling		
var6 = individual_person		
var5 = *var3*		
var6 ≠ *var2*		
var6 = *var4*		
If the same person (see the equality constraint between var4 and var6 in conseq2) feels the same emotion (see the equality constraint between var3 and var5 in the same conseq2) in front of two different entities—here a location, LANDSCAPE_1 and an individual, Mona Lisa—it is possible that some type of relationship exists between these two entities.		

4. Conclusions

As stated at length in this paper, the work presented here represents only a very preliminary experiment whose initial results should be confirmed in the framework of larger (European?) initiatives. In any case, we are particularly pleased to have led to a good end a challenge whose positive result was far from being evident at the beginning. In the context of this experiment, NKRL has in fact (at least provisionally) confirmed, both from a specific knowledge representation and an inferential point of view, its well-known generality and efficiency for the in-depth representation of any possible "narrative" environment—in this case, the "immaterial" component (including its emotional factors) of the twins associated with Cultural Heritage "iconographic" entities. To be able, in fact, to describe in a really expressive, complete and meaningful form this immaterial component, we cannot limit ourselves to utilize a simple description under metadata form or to make use of some knowledge representation tools unsuitable for representing complex situations/events, like the semantic web tools according to their traditional "binary" form. Some examples of "lack of expressiveness" of the traditional tools have already been supplied in Section 2.1 above.

The main concrete difficulty encountered in this challenge has concerned the need of adapting its "traditional concepts" component, HClass, to a completely new domain, the Cultural Heritage one—even if, in reality, about 80% of the concepts used in Tables 1, 3 and 4 were already present in HClass or have been easily derived from existing concepts. We can note that, should this work find some form of continuation, this problem should not be considered as really important in a domain where a quantity of "terminological" tools already exist, from WordNet to CIDOC CRM. This is without forgetting, moreover, that the new developments in an AI framework are likely to increasingly facilitate the automatic/semi-automatic set-up of important corpora of new concepts from object-based scene representations (see, to give only an example, a very recent paper like [18]).

References

1. European Commission: Horizon 2020, Work Programme 2018-2020 - 13. Europe in a changing world – Inclusive, innovative and reflective societies, 2019. Available online: https://ec.europa.eu/research/participants/data/ref/h2020/wp/2018-2020/main/h2020-wp1820-societies_en.pdf (accessed on 25 July 2019).
2. Grieves, M. *Origins of the Digital Twin Concept—Working Paper*; Florida Institute of Technology: Melbourne, FL, USA, 2016; Available online: https://www.researchgate.net/publication/307509727 (accessed on 30 July 2019).
3. Amelio, A. Exploring Leonardo Da Vinci's Mona Lisa by Visual Computing: A Review. In Proceedings of the 1st International Workshop on Visual Pattern Extraction and Recognition for Cultural Heritage Understanding (VIPERC). CEUR-WS, Pisa, Italy, 30 January 2019; pp. 25–38.
4. Zarri, G.P. *Representation and Management of Narrative Information—Theoretical Principles and Implementation*; Springer: London, UK, 2009.
5. Lagoze, C.; Hunter, J. The ABC Ontology and Model. In Proceedings of the International Conference on Dublin Core and Metadata Applications, Tokyo, Japan, 24–26 October 2001; pp. 160–176.
6. Nilsson, M.; Powell, A.; Johnston, P.; Naeve, A. *Expressing Dublin Core Metadata Using the Resource Description Framework, RDF (DCMI Recommendation 2008-01-04)*; Dublin Core Metadata Initiative: Silver Spring, MD, USA, 2008; Available online: https://www.dublincore.org/specifications/dublin-core/dc-rdf/ (accessed on 5 August 2019).
7. Le Boeuf, P.; Doerr, M.; Ore, C.E.; Stead, S. (Eds.) *Definition of the CIDOC Conceptual Reference Model (version 6.2.6, June 2019)*; ICOM/CIDOC Documentation Standard Group: Heraklion, Greece, 2019.
8. Troncy, R.; van Ossenbruggen, J.; Pan, J.Z.; Stamou, G. (Eds.) *Image Annotation on the Semantic Web, W3C Incubator Group Report 14 August 2007*; W3C: Cambridge, MA, USA, 2007; Available online: https://www.w3.org/2005/Incubator/mmsem/XGR-image-annotation/ (accessed on 10 August 2019).

9.　Isaac, A.; Clayphan, R. (Eds.) *Europeana Data Model Primer*; Europeana Foundation: The Hague, The Netherlands, 2011; Available online: https://pro.europeana.eu/files/Europeana_Professional/Share_your_data/Technical_requirements/EDM_Documentation/EDM_Primer_130714.pdf (accessed on 10 August 2019).

10.　D'Andrea, A.; Ferrandino, G. Shared Iconographical Representations with Ontological Models. In Proceedings of the 35th International Conference on Computer Applications and Quantitative Methods in Archaeology (CAA), Budapest, Hungary, 2–6 April 2008; pp. 194–199.

11.　Carboni, N.; de Luca, L. An Ontological Approach to the Description of Visual and Iconographical Representations. *Heritage* **2019**, *2*, 1191–1210. [CrossRef]

12.　Lodi, G.; Asprino, L.; Nuzzolese, A.G.; Presutti, V.; Gangemi, A.; Reforgiato Recupero, D.; Veninata, C.; Orsini, A. Semantic Web for Cultural Heritage Valorisation. In *Data Analytics in Digital Humanities Valorisation*; Hai-Jew, S., Ed.; Springer International Publishing AG: New York, NY, USA, 2017; pp. 3–37.

13.　Zarri, G.P. Functional and Semantic Roles in a High-Level Knowledge Representation Language. *Artif. Intell. Rev.* **2019**, *51*, 537–575. [CrossRef]

14.　Zarri, G.P. Modelling and Exploiting the Temporal Information Associated with Complex 'Narrative' Documents. *Int. J. Knowl. Eng. Data Min.* **2019**, *6*, 135–167. [CrossRef]

15.　Zarri, G.P. Advanced Computational Reasoning Based on the NKRL Conceptual Model. *Expert Syst. Appl.* **2013**, *40*, 2872–2888. [CrossRef]

16.　Schwartz, L.F.F. The Mona Lisa Identification: Evidence from a Computer Analysis. *Vis. Comput.* **1988**, *4*, 40–48. [CrossRef]

17.　Asmus, J.F. Mona Lisa Symbolism Uncovered by Computer Processing. *Mater. Charact.* **1992**, *29*, 119–128. [CrossRef]

18.　Mao, J.; Gan, C.; Kohli, P.; Tenenbaum, J.B.; Wu, J. The Neuro-Symbolic Concept Learner: Interpreting Scenes, Words, and Sentences from Natural Supervision. In Proceedings of the Seventh International Conference on Learning Representations (ICLR), New Orleans, LA, USA, 6–9 May 2019; Available online: https://openreview.net/pdf?id=rJgMlhRctm (accessed on 1 September 2019).