Representing Narratives in Digital Libraries: The Narrative Ontology

Carlo Meghini [a], Valentina Bartalesi [a,*] and Daniele Metilli [a,b]
[a] *ISTI-CNR, Pisa, Italy*
[b] *Department of Computer Science, University of Pisa, Italy*
E-mails: carlo.meghini@isti.cnr.it, valentina.bartalesi@isti.cnr.it, daniele.metilli@isti.cnr.it

Editors: Antonis Bikakis, University College London, UK; Béatrice Markhoff, University of Tours, France; Alessandro Mosca, Free University of Bozen-Bolzano, Italy; Stéphane Jean, University of Poitiers – ENSMA, France; Eero Hyvönen, University of Helsinki, Finland
Solicited reviews: Béatrice Markhoff, University of Tours, France; Seven anonymous reviewers

Abstract. Digital Libraries (DLs), especially in the Cultural Heritage domain, are rich in narratives. Every digital object in a DL tells some kind of story, regardless of the medium, the genre, or the type of the object. However, DLs do not offer services about narratives, for example it is not possible to discover a narrative, to create one, or to compare two narratives. Certainly, DLs offer discovery functionalities over their contents, but these services merely address the objects that carry the narratives (e.g. books, images, audiovisual objects), without regard for the narratives themselves. The present work aims at introducing narratives as first-class citizens in DLs, by providing a formal expression of what a narrative is. In particular, this paper presents a conceptualisation of the domain of narratives, and its specification through the Narrative Ontology (NOnt for short), expressed in first-order logic. NOnt has been implemented as an extension of three standard vocabularies, i.e. the CIDOC CRM, FRBRoo, and OWL Time, and using the SWRL rule language to express the axioms. On the basis of NOnt, we have developed the Narrative Building and Visualising (NBVT) tool, and applied it in four case studies to validate the ontology. NOnt is also being validated in the context of the Mingei European project, in which it is applied to the representation of knowledge about Craft Heritage.

Keywords: Narrative, Digital Libraries, Semantic Web, Ontology, Cultural Heritage, Craft Heritage

1. Introduction

Digital Libraries (DLs) abound with narratives, in the sense that every digital object in a DL tells some kind of story, regardless of the medium, the genre, or the type of the object. This is especially true for DLs in the Cultural Heritage domain [1]. However, there is no track of narratives in the services offered by today's DLs. It is not possible, e.g., to discover a narrative, to create one, or to compare two narratives. Of course, any DL offers a discovery service over its content; but this service addresses the objects that carry the narratives, whether books, audio-visual messages and the like; narratives *per se* are not addressed. It may be said, in short, that DLs ignore their contents.

Yet, narratives are central to the documentation of human activity, whether in the cultural, the scientific, or the social area. An art historian willing to tell the reconstructed story surrounding the creation of a painting; a scientist wishing to describe the phases of development and validation of a theory; a sociologist wishing to recount the impact of a social medium in time. All these knowledge operators would take great advantage of a narrat-

*Corresponding author. E-mail: valentina.bartalesi@isti.cnr.it.

0000-0000/0-1900/$00.00 © 0 – IOS Press and the authors. All rights reserved

ive service. And so would a librarian wishing to provide an account of the process of curating a certain type of collection, or an archivist giving a historical record of the preservation of an item. The only option available to them is to use text, or an analogous medium, to tell their story. But once so encoded, the narrative is lost to the DL.

Until machines will exhibit the human ability to interpret media contents, one way to overcome the present status is to make narratives emerge as objects of an autonomous data type, different from any other data type, and amenable to (narrative-aware) machine processing. In other words, to make narratives emerge as formal objects, much in the same way other documentation artefacts, such as bibliographic records, ontologies and terminologies, have emerged as formal objects in time. But to be most effective, formal narratives should not *replace* traditional, informal narratives: rather, they should *enhance* them, by adding a formal dimension to the existing one.

The study of narrative goes back to Aristotle [2] and to the fourth century BC, and has been further elaborated by many philosophers afterwards. The Russian formalists, around the 20s of the last century, have offered an account of narratives that has been used for a systematic study of narrative structure [3]. This account has finally given rise to narratology as an autonomous scientific discipline. According to the Russian formalists, a narrative consists of:

- the *fabula,* i.e., the story itself as it happened, in reality or in fiction;
- the *narrations,* i.e., one or more expressions, each in its own language and *medium,* that narrate the fabula. Each narration corresponds to Bal's definition of *presentation* [4];
- the *plot,* i.e., the story as it is narrated by the narrator. The plot corresponds to the *syuzhet* of the Russian formalists and to Aristotle's *logos.*

Current DLs contain only the *narration* level of the narrative, i.e. the expression of the narrative through a media object. To enhance the representation of narratives in DLs, we propose adding a formal expression of the *fabula* and the *plot*. The resulting representations would enter the information space of a DL as first-class citizens, enabling an entirely new set of services, able to exploit both the informal and the formal dimension of narratives, and the relation between them. Needless to say, knowledge extraction methods from media objects are central to our proposal, as will be argued in due course.

The paper presents a research work that significantly extends a previous study [5]. We have defined a conceptualisation of the domain of narratives, and we have provided its specification through the Narrative Ontology (NOnt for short), expressed in first-order logic. The ontology has been implemented as an extension of three standard vocabularies, i.e. the CIDOC CRM, FRBRoo, and OWL Time, and using the SWRL rule language to express the axioms. On the basis of the ontology, we have developed the Narrative Building and Visualising (NBVT) tool [6], and applied it in four case studies[1]. NOnt is also being validated in the context of the Mingei European project[2], in which the ontology is used for the representation of knowledge about Craft Heritage [7].

The paper is structured as follows: after describing our methodological approach (Section 2), we report a review of existing works about narrative modelling (Section 3). Section 4 presents a detailed conceptualisation of narratives based on narratology, followed by a discussion of narratives in DLs (Section 5). Section 6 presents the NOnt ontology, i.e. a specification of the conceptualisation in first-order logic. Section 7 discusses an implementation of NOnt using Semantic Web languages. Section 8 reports the usage of the ontology in two practical scenarios: (i) the Narrative Building and Visualising tool, and (ii) the Mingei European project. Section 9 concludes and outlines further developments.

2. Methodological Approach

The methodological approach we followed to introduce narratives as a new functionality in DLs is very similar to the one that characterises a common workflow to develop an algorithm in Computer Science [8], that is:

1. Formalisation of the problem
2. Computational analysis
3. Development of a new algorithm
4. Experimentation with a case study

[1] https://dlnarratives.eu/narratives.html
[2] http://www.mingei-project.eu

5. Evaluation

The phases of algorithm development were adapted to our aim. In particular, the adopted methodological approach consists of the following phases:

1. Creation of a conceptualisation of the domain, in which the issue is described and analysed in its main parts.
2. Development of an ontology as the specification of the conceptualisation in terms of a logical theory whose axioms admit as models those licensed by the conceptualisation.
3. Development of an inference engine for reasoning on knowledge bases conforming to the ontology.
4. Implementation of the ontology and of the inference engine, using Semantic Web technologies.
5. Evaluation of the ontology.

In this paper we cover all these phases. The first and second phases are described in Sections 4 and 5, respectively. In most cases, the third phase is not necessary, as Semantic Web technologies provide inference engines for ontologies whose axioms are expressible in one of the Semantic Web languages. These engines perform two fundamental tasks: (1) they answer queries and (2) they check the consistency of a knowledge base; in carrying out both tasks, they reason over the ontology, i.e. they deduce implicit knowledge from the knowledge that is explicitly given.

Our ontology includes predicates for representing qualitative temporal knowledge about the intervals of occurrence of events. Using these predicates, it can be stated, for instance, that an event occurred before, or during, another event. For this purpose, our ontology relies on 13 basic temporal relations (BTRs for short), proposed in the seminal work of James Allen [9] to capture all possible ways in which two intervals can relate to each other. Allen also provided transitivity rules that allow deriving implicit temporal relations from known ones. These transitivity rules can indeed be expressed in one of the Semantic Web languages, the SWRL rule language. However, the temporal relations between two intervals become exponentially many if disjunctions of BTRs are used to state temporal knowledge; and in the case of narratives, these disjunctions are needed, as will be shown in Section 6.5. Consequently, transitivity rules become exponentially many, making reasoning over narratives practically impossible. Tractable subsets of BTRs are known to exist for certain domains, but unfortunately, none of these known subsets can be applied to our case, as it will also be shown in Section 6.5. Therefore, we had to search for a tractable subset of BTRs, so as to provide a complete axiomatisation of our ontology in SWRL, enabling efficient reasoning on narratives. This research is precisely the third phase of the above methodology, and the good news is that we were able to find a tractable subset of BTRs for our ontology, as reported in Section 6.5 of this paper.

We discuss the fourth phase in Section 7, providing an implementation of NOnt using Semantic Web technologies.

Concerning evaluation, which is the last phase of the methodology, we note that our ontology is endowed with a reasoning algorithm that is sound, complete, and efficient, as just explained. We consider this as a qualitative validation of our work, and a crucial one: without such a reasoning algorithm, our ontology would not be usable, given that we aim at *computable* narratives. Of course, a theoretical validation is not sufficient, it must be complemented by a pragmatic one.

In this respect, it should be noted that the very notion of narrative has been recognised as a complex one, and that narratologists have not yet reached a commonly-shared understanding of the fundamental aspects of narrative [10]. Under these circumstances, no ontology can be validated as *the* ontology of narrative. Pragmatically, then, all that can be done is to test a tentative ontology that is promising, to see whether it responds reasonably well to the requirements at hand. This is what we have done with our ontology, by applying it in two evaluation contexts: one in the laboratory and the other on the field in the context of the Mingei project.

3. Related Works

To define a conceptualisation, we started from the study of narratology, in order to identify the fundamental concepts of narratives. Narratology is a discipline in the Humanities "dedicated to the study of the logic, principles, and practices of nar-

rative representation" [11]. In this research field, the concept of *event* is a core element of narrative. Event is generally intended as an occurrence taking place at a certain time at a specific location. Despite its antecedents in classical theories of aesthetics [2], the theoretical principles of narratology derive from linguistic-centred approaches to literature defined by Russian formalists in the early 20th century. Russian formalism identified two structural levels of narrative: (i) the *fabula*, i.e. the sequence of events of the narrative in chronological order; (ii) the *syuzhet* (or *plot*), that is the way in which these events are presented in a narrative [3]. In more recent years, Bal [4] defined a third level, called *presentation*, that constitutes the concrete representation of the content that is conveyed to the audience (e.g. the text in a novel). Narratology also identifies *characters* as a fundamental constituent in a story. Aristotle [2] affirms that characters appear in every type of tale. McKee [12] claims that it is not possible to talk about the plot without the characters and vice versa. According to Chatman [13], the elements of a story can be distinguished in: (i) characters, (ii) elements in the scenario. Characters are usually humans or humanoid beings, while the elements in the scenario are places and objects. In our work, we employed the structural levels of narrative, as defined by Russian formalism, as the base elements of our conceptualisation.

After analysing the narratology literature, we reviewed the Artificial Intelligence literature, and in particular the Event Calculus theory [14–16], in order to understand whether the components of narrative had been formally defined in this field of research. The Event Calculus (EC) is a logic language for representing actions that have duration and can overlap with each other. In the EC, we found the basic elements for representing the fundamental concepts of narrative. The first is the concept of *Fluent*, that identifies a function or a predicate that varies over time, and is used to describe the effects of actions [17]. Two other key concepts are *Events* and *Actions*. In the EC, the terms Actions and Events are interchangeable, and represent changes performed over time. On the other hand, Davidson's theory [18] defines action as a particular subclass of event that is endowed with intentionality. The last core concept is the *Generalised event*, i.e. a space-time chunk which generalises concepts like actions, locations, times, and physical objects such as things, animals, agents, humans. The fundamental concepts of narrative extracted from the EC were represented as core elements of our conceptualisation.

In the Semantic Web field, various models have been developed for representing the core concept of event. These models include the Event Ontology [19], the Linking Open Descriptions of Events (LODE) [20], the Event-Model-F ontology [21], and the Simple Event Model (SEM) [22], among others. More general models for semantic data organisation are the CIDOC CRM [23], the Europeana Data Model [24], and the DOLCE upper level ontology [25]. Among the models reported above, we chose the CRM as reference vocabulary for our ontology for narratives, and we took inspiration in particular from the LODE and SEM ontologies in order to represent the factual components of events [26].

In the Digital Libraries field, narratives have been proposed as functionalities to improve the information discovery and exploration of the contents of DLs. In the following, we report several projects that employed narratives as tools to explore digital objects, and that we took into account in the development of our ontology and software. CultureSampo [27] is a portal and a publication channel for Finnish cultural heritage based on Semantic Web technologies. It uses an event-based model that allows linking events with digital objects, although it does not define the specific semantic relations that connect events to objects. BiographySampo [28] is a project that aims to develop a system to extract narratives from biographical dictionaries, represent them in a formal way using the CIDOC CRM and other ontologies, and publish them on the Web as Linked Data. The system has been used to build a portal containing more than 13,000 biographies of historical Finnish people. Another example is Bletchley Park Text [29], an application that helps users to explore the collections of museums. Visitors express their interests on some specific topics using SMS messages containing keywords. The semantic description of the resources is used to organise a collection into a personalised website based on the keywords chosen by the user. The PATHS system [30] allows creating a personalised tour guide through existing DL collections. The system allows the definition of events linked to each other by semantic similarity relations. The

Storyspace system [31] allows describing stories based on events that span museum objects. The focus of the system is the creation of curatorial narratives from an exhibition. Each digital object has a linked creation event in the story of a heritage object. The Labyrinth 3D system [32] integrates the semantic annotation of cultural objects with the interaction style of 3D games. The system immerses the user into a virtual reality, where the user can explore the collection using paths representing the semantic relations over cultural objects.

In comparison to the above systems, our idea is to develop a software that allows the user to create semantic networks endowed with the events that compose the narratives, along with their formal components and the related digital objects. The events are linked to each other through semantic relations.

4. Conceptualisation

This Section presents our view of a computable representation of narrative, as informed by the background reported in Section 3. We introduce the relevant notions both at an informal level and more formally in set-theoretic terms. An initial version of this conceptualisation has been reported in [33]. The present version extends the initial one in several important ways.

Narrative We view a narrative as a story told by a narrator, which may be an individual person or a group of persons taking up the role of the narrator. Each narrative reflects the point of view of its narrator. The stories in the scope of our work are generally real stories of the present or the past. Fictional stories may also be expressed in our ontology. However, since it is more important for us to support science rather than fiction, these stories have to be consistent with the axioms on physical reality that our ontology is able to capture. This excludes stories in which, for instance, effects precede causes, events nest circularly, or objects bilocate.

A narrative consists of three main elements:

1. the *fabula,* i.e., the story itself as it happened, in reality or in fiction;
2. the *narrations,* i.e., one or more expressions, each in its own language and *medium,* that narrate the fabula. Each narration correspond to Bal's definition of *presentation* [4];
3. the *reference,* i.e., a relation that connects (fragments of) the narrations to (fragments of) the fabula, allowing the derivation of the *plot* (or *syuzhet*) of the narrative.

Fabula A fabula consists of events, each of which encompasses a significant fragment of the story. We define an event as a group of coherent phenomena situated in space and time. An event is contextualised in terms of the entities that participate in it. In addition to space and time, the other entities that participate in the event can be identified as having persistent characteristics of structural nature. They may be physical entities (e.g. people, physical objects) or conceptual entities (e.g. concepts, ideas). Actions, and more generally activities, are special cases of events. In the fabula, the events are ordered chronologically, as defined by Russian formalism. Moreover, the events in a fabula participate in three main relations:

- a *mereological* relation, connecting events to other events that include them as parts, e.g. the birth of a person is part of the broader event of the life of that person.
  The event composition relation is a strict partial order, i.e. it is an irreflexive and transitive relation over the fabula's events; consequently, it is asymmetric, and more generally acyclic, so that no event is a sub- or super-event of itself or some other event.
- a *temporal occurrence* relation, associating each event with a time interval during which the event occurs. As such, the temporal occurrence relation is a total function. In turn, time intervals are connected to each other through the 13 relations of Allen's interval algebra [9]. These relations are jointly exhaustive and mutually exclusive, therefore each pair of events is connected by one and only one Allen relation. Each time interval has exactly one starting time point and one ending time point. Time points are connected to each other through before, after, or equals relations;
- a *causal dependency* relation, relating pairs of events such that the occurrence of the former causes the occurrence of latter, e.g. the eruption of the Vesuvius caused the de-

struction of Pompeii. Clearly, a formal account of the causal dependency relation requires a complete knowledge of the laws governing reality, and is therefore out of question. Instead, we will confine ourselves to assert that causal dependency is a *strict partial order*. In this case, acyclicity guarantees that no event is at the same time cause and effect of itself or some other event.

In addition to the features of the individual relations in a fabula stated so far, the following conditions are met by every fabula:

1. The period of occurrence of an event is included in the period of occurrence of any of its super-events.
2. The beginning of occurrence of an event precedes the beginning of occurrence of any event that causally depends on it.

The expression of the inclusion and precedence relations among time intervals will be dealt with in Section 6.4, upon considering the representation of temporal knowledge in narratives.

Narrations In a narrative, a fabula may have any number of narrations, each of which has the obvious characteristic of being *about* the fabula. Intuitively, this aboutness is a notion of *representation* between fabula and narration, in the sense that any narration of the fabula must somehow represent the fabula, in whole or in part. Logically, this amounts to say that any proposition in the narration, whether explicitly or implicitly stated, must be true in the fabula.

Each narration has one or more narrators, the authors of the narration, and of a *narration content*. In general, the narration content is a message, and it may take any form in which a fabula can be communicated, ranging from text, to audio-visual message, to theatrical enactment, etc. For obvious reasons, we are interested in narrations that have at least one digital representation, whether such representation is only the carrier of a non-digital narration (e.g. an audio-visual recording of a theatrical piece) or a born-digital narration (e.g. a born-digital text or a video game). In our conceptualisation of a narration, the content will therefore be any media object, i.e. a text, an image, an audio-visual object, or any complex multimedia object that a particular narrator, or group of narrators, choose to tell their version of the fabula.

Reference In our conceptualisation, the reference is a relation that connects regions of narrations, which we call *narrative fragments* (or simply *fragments*), to events of the fabula. Each fragment is maximal, in that it comprises all portions of the narration that narrate the same event.

A fragment is identified in ways that depend on the structure of the narration. For instance, a textual fragment will be a set of disjoint intervals, each giving the boundaries of texts narrating the same event. A fragment that narrates an event necessarily narrates any of its super-events, and no other event.

Using the reference relation, it is possible to reconstruct the plot of the narrative, i.e. the sequence of fragments in the order established in the narration by the narrator.

Because a fabula is identified by its composing events, two narrations of the same fabula may differ for any combination of the following:

1. the set of events of the fabula narrated by the narrations; each narration may pick a different subset of events, as a way of giving more emphasis to certain aspects of the story;
2. the order in which the selected events are narrated;
3. the expressions used for the narration.

Two narrations offering accounts of the same story that are incompatible, in the sense expressed above, are not narrations of the same fabula. This fact does not prevent comparing the narrations, for instance to appreciate their differences.

5. Representing narratives in DLs

In our view, a Digital Library (DL) should provide digital representations of narratives as first-class citizens of the DL. For simplicity, we will call such digital representations *narratives* whenever no ambiguity can arise.

For completeness, the narratives in a DL should encompass all aspects discussed in the previous Section, i.e. narrations, fabulae and reference relations. While it is expected that a DL already possesses narrations in digital form, our work is motivated by the target of lifting such narrations into narratives, endowing them with a formal representation of the corresponding fabulae, acting as a semantical counterpart of those narra-

tions. Clearly, this two-level representation of the narrative allows supporting the union of the use cases supported by the purely syntactical (i.e. based solely on narrations) and the purely semantical (i.e. based solely on fabulae) representations. From now on, when there is no ambiguity, we will speak of the fabula of a narrative meaning the representation of the fabula, as we do for narratives.

A narrative may be constructed in at least two different ways:

- starting from a narration and associating it to a fabula, or
- starting from a fabula and associating it to a narration of the fabula.

In the former case, the involved process is *formalisation:* the narration is decomposed into meaningful events and each event is formally represented via statements drawn from the narration; the reference relation is used to establish the proper connection between fragments of the narration and the corresponding formalising events. In the latter case, the involved process is *documentation:* the events of the fabula are given, and the narrative is constructed by linking each of them to a narration fragment that illustrates the event, using for that purpose (the inverse of) the reference relation. In either case, automatic or semi-automatic methods can be devised to support the process and make it scale.

It must be noted that either the narration or the fabula of a narrative may provide an incomplete, or even inaccurate, account of the story that the narrative is about. In each of them, events may be reported by omitting or mistaking their temporal or spatial occurrence; likewise, the participation of persons in events or the causal dependencies between events may be omitted or mistaken. For this reason, the fabula of a narrative must be treated as a knowledge base (KB for short), i.e. a set of statements giving the best available approximation of the fabula according to the narrator of the narrative.

The relationship between the real fabula and its representation may be precisely characterised from a logical point of view as follows. A real fabula f may be seen as a set of possible worlds, namely of the worlds that are compatible with the events in the fabula and the relationships that link these events to each other and to their factual components. Let S_f be the maximal set of formal fabula statements that are true in every world in f. A language for expressing these statements will be introduced in the next Section, but for now it suffices to assume that such language exists. Let k be a non-empty KB with the formal representation of f. Then,

- k is an *accurate* representation of f iff every statement in k is true in the fabula, formally iff $k \models S_f$, where $\models$ is the logical implication relation.
- k is a *complete* representation of f iff k says everything about f, formally iff $S_f \models k$.

Accurate and complete accounts of the fabula are therefore knowledge bases k that are equivalent to S_f,. Needless to say, such accurate and complete accounts are idealisations that real representations can only try to approximate.

As a consequence of the inaccuracy or incompleteness of fabulae, and therefore of narratives in general, it may be the case that two narratives provide different versions of the same story, making different statements about the same events, possibly leading to contradiction. For instance, a narrative about the life of Dante Alighieri may include a travel to France as an event, while another narrative may deny the occurrence of that event, e.g. by placing Dante at a different location at the same time. Needless to say, the presence of different versions of the same story is not to be seen as accidental or undesirable in a DL. To the contrary, it manifests different points of view that it is important (in some cases vital) to document. On the other hand, the arising of logical contradictions in a KB is highly undesirable, because it makes the KB unusable: since everything logically follows from an inconsistent KB, the answers to queries performed against an inconsistent KB will not be reliable.

In order to enable a DL to hold incompatible narratives while at the same time avoiding the rise of inconsistencies, we view each narrative as a separate KB, and a DL as a set of narratives, possibly sharing a common set of factual components that occur in the fabulae of these narratives.

In the present study, we focus on the structure and the operation of single narratives, because they present challenging aspects in their own right, as will be shown in the rest of the paper.

6. The NOnt ontology

This Section presents an ontology of narratives, called NOnt, which specifies the conceptualisation given in the previous Section.

As already pointed out, narrations will be represented by digital media objects; each such object gives a narration of some part of, or possibly all, the narrative. Our ontology will not provide machinery to deal with narrations, since they are strongly medium-dependent and as such outside the scope of our work. Narrations will be treated as "black boxes", each represented by a different identifier and characterised as an instance of a special class. Such class will be an extension point of NOnt, in the sense that it is the part of the ontology where the classes and properties for representing narrations can be plugged, for example by drawing them from other standard ontologies.

The ontology is expressed in first-order logic [34] for maximum expressivity. Due to the fact that a DL includes a *global* KB, i.e. a set of statements that document the narratives encompassed in the DL, NOnt will be split in two parts: (i) NOntNar, including the classes, properties and axioms for expressing individual narratives, and (ii) NOntDL, including the classes, properties and axioms for expressing the knowledge in the global KB of a DL. Before delving into the definition of the ontology, we will discuss some epistemic aspects at the basis of NOnt.

6.1. The $\mathcal{L}_n$ language

Our task requires the identification of a specific first-order language $\mathcal{L}_n$ that is able to capture the intended meaning of our ontology for narratives.

$\mathcal{L}_n$ is derived from the $\mathcal{L}$ presented in [35]. It includes the sentences that are required in order to axiomatise the narratives. As customary in logic, the alphabet of $\mathcal{L}_n$ includes two kinds of symbols: logical and non-logical. Logical symbols are the symbols whose usage and interpretation are fixed, while non-logical symbols are the symbols that are domain-dependent.

The logical symbols of $\mathcal{L}_n$ are:

- countably many variables $x, y, z \ldots$;
- the equality symbol $=$ naming the well known equality relation;
- the connectives $\neg$ and $\vee$ and the existential quantifier $\exists$.

The non-logical symbols of $\mathcal{L}_n$ are:

- countably many constant symbols, or simply constants: $a, b, \ldots$;
- unary and binary predicate symbols.

$\mathcal{L}_n$ also includes predicate symbols required to represent and reason about time in narratives. We defer the discussion of those symbols, and the axioms that define them, until Section 6.4.

The terms of $\mathcal{L}_n$ are constants and variables. The atoms of $\mathcal{L}_n$ are expressions of the form $P(t_1, \ldots, t_k)$, where each t_i is a term. A ground atom is an atom $P(t_1, \ldots, t_k)$ where each t_i is a constant. A formula of $\mathcal{L}_n$ is one of the following:

- an atom;
- a co-reference formula of the form $(t_1 = t_2)$, where t_1 and t_2 are terms;
- the negation of a formula $\neg\alpha$;
- the disjunction of two formulas $(\alpha \vee \beta)$;
- an existential quantification of the form $\exists x.\alpha$

A sentence of $\mathcal{L}_n$ is a formula whose variables, if any, are each bound to one quantifier, i.e. a formula with no free variables. As customary, we will consider as part of $\mathcal{L}_n$ sentences including the universal quantifier $\forall$ and the connectives $\wedge$ ("and") and $\rightarrow$ ("implies"), obtained as abbreviations of the equivalent sentences using the previously introduced symbols. Furthermore, to simplify the notation we omit universal quantifiers in formulae. All predicate symbols denote pairwise disjoint sets, i.e.:

$$A(x) \rightarrow \neg B(x) \tag{1}$$

$$P(x, y) \rightarrow \neg R(x, y) \tag{2}$$

where A and B stand for any two different unary predicate symbols, and P and R stand for any two different binary predicate symbols.

The following equality axioms hold in $\mathcal{L}_n$:

$$x = x \tag{3}$$

$$(x = y) \rightarrow (y = x) \tag{4}$$

$$[(x = y) \wedge (y = z)] \rightarrow (x = z) \tag{5}$$

$$(x = y) \rightarrow [A(x) \equiv A(y)] \tag{6}$$

$$[(x_1 = y_1) \wedge (x_2 = y_2)] \rightarrow [P(x_1, y_1) \equiv P(x_2, y_2)] \tag{7}$$

Table 1
The predicate symbols of NOntNar

Unary Predicate Symbols	
Ev(e)	e is an event
Interval(t)	t is a time interval
Place(p)	p is a place
Partic(c)	c is an entity that may participate in an event (e.g. a person, an object, a concept)
Fab(f)	f is a fabula
Nar(a)	a is a narration
MObj(o)	o is a media object
MOFrag(r)	r is a media object fragment
Binary Predicate Symbols	
EP(e_1, e_2)	event e_1 is part of event e_2
EC(e_1, e_2)	event e_1 is causally dependent on event e_2
ETI(e, t)	event e occurs at time interval t
EPlace(e, p)	event e occurs in place p
EPartic(e, c)	event e has participant c
FE(f, e)	fabula f has event e
FN(f, a)	fabula f has narration a
Cont(n, o)	narration n has content o
OF(o, r)	media object o has fragment r
Ref(r, e)	fragment r is about event e
TINC(t_1, t_2)	interval t_1 includes interval t_2
TIP(t_1, t_2)	interval t_1 starts before interval t_2

where A and P are as above.

We adopt the standard first-order semantics to assign meaning to the formulas of $\mathcal{L}_n$.

6.2. *The axioms of the NOntNar ontology*

Table 1 lists the unary and binary predicates of the NOntNar ontology. Figures 1 and 2 show a graphical representation of these predicates. A concrete example of narrative is reported at the end of this Section.

In the following, we list all the axioms holding on the unary and binary predicates of NOntNar. The following axioms provide the domain and range of each binary predicate:

$$\mathsf{EP}(x, y) \rightarrow \mathsf{Ev}(x) \wedge \mathsf{Ev}(y) \quad (8)$$

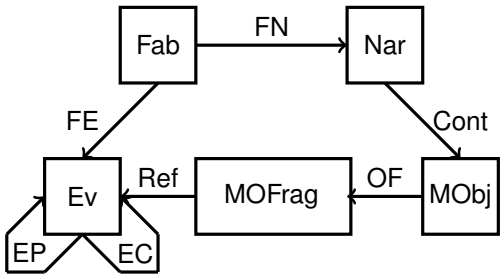

Figure 1. The NOntNar ontology. Partial view to be completed by the predicates related to events, shown in Figure 2

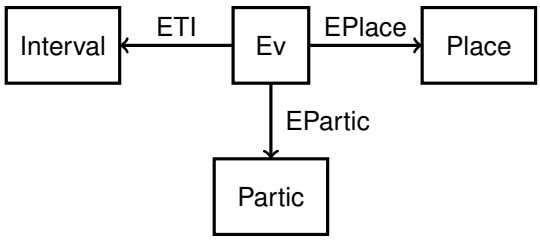

Figure 2. A view of the NOntNar predicates related to events

$$\mathsf{EC}(x, y) \rightarrow \mathsf{Ev}(x) \wedge \mathsf{Ev}(y) \quad (9)$$

$$\mathsf{ETI}(x, y) \rightarrow \mathsf{Ev}(x) \wedge \mathsf{Interval}(y) \quad (10)$$

$$\mathsf{Cont}(x, y) \rightarrow \mathsf{Nar}(x) \wedge \mathsf{MObj}(y) \quad (11)$$

$$\mathsf{OF}(x, y) \rightarrow \mathsf{MObj}(x) \wedge \mathsf{MOFrag}(y) \quad (12)$$

$$\mathsf{Ref}(x, y) \rightarrow \mathsf{MOFrag}(x) \wedge \mathsf{Ev}(y) \quad (13)$$

The following cardinality restrictions apply:

– An event has exactly one time interval:

$$\mathsf{Ev}(x) \rightarrow (\exists y)\mathsf{ETI}(x, y) \quad (14)$$

$$\mathsf{ETI}(x, y_1) \wedge \mathsf{ETI}(x, y_2) \rightarrow y_1 = y_2 \quad (15)$$

– An event has exactly one place:

$$\mathsf{Ev}(x) \rightarrow (\exists y)\mathsf{EPlace}(x, y) \quad (16)$$

$$\mathsf{EPlace}(x, y_1) \wedge \mathsf{EPlace}(x, y_2) \rightarrow y_1 = y_2 \quad (17)$$

– An event has one or more participants:

$$\mathsf{Ev}(x) \rightarrow (\exists y)\mathsf{EPartic}(x, y) \quad (18)$$

– A fabula has one or more events:

$$\mathsf{Fab}(x) \rightarrow (\exists y)\mathsf{FE}(x, y) \quad (19)$$

– A fabula has one or more narrations:

$$\mathsf{Fab}(x) \rightarrow (\exists y)\mathsf{FN}(x, y) \quad (20)$$

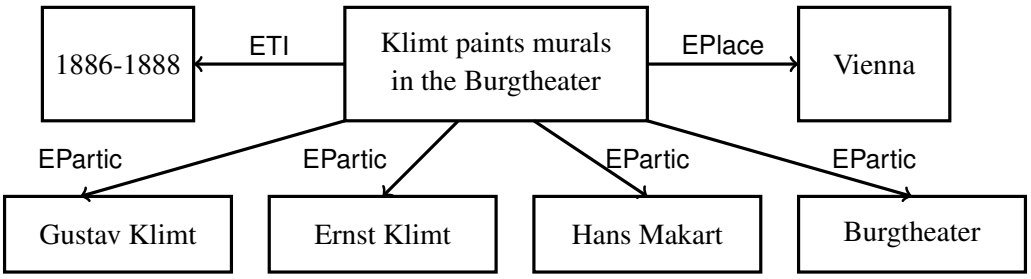

Figure 3. A graph-based representation of the event "Klimt paints murals in the Burgtheater"

– A narration has exactly one content:

$$\mathsf{Nar}(x) \rightarrow (\exists y)\mathsf{Cont}(x,y) \tag{21}$$

$$\mathsf{Cont}(x,y_1) \wedge \mathsf{Cont}(x,y_2) \rightarrow y_1 = y_2 \tag{22}$$

– A fragment belongs to exactly one media object:

$$\mathsf{MOFrag}(x) \rightarrow (\exists y)\mathsf{OF}(y,x) \tag{23}$$

$$\mathsf{OF}(y_1,x) \wedge \mathsf{OF}(y_2,x) \rightarrow y_1 = y_2 \tag{24}$$

We do not admit as consistent the narratives in which event parthood and causal dependency are cyclic, i.e. in which an event is a sub- or super-event of itself or some other event, or in which an event is at the same time a cause and an effect of itself or some other event. Since the relations corresponding to these symbols are transitive, by imposing irreflexivity we obtain acyclicity:

$$\mathsf{EC}(x,y) \rightarrow \neg(x = y) \tag{25}$$

$$\mathsf{EC}(x,y) \wedge \mathsf{EC}(y,z) \rightarrow \mathsf{EC}(x,z) \tag{26}$$

$$\mathsf{EP}(x,y) \rightarrow \neg(x = y) \tag{27}$$

$$\mathsf{EP}(x,y) \wedge \mathsf{EP}(y,z) \rightarrow \mathsf{EP}(x,z) \tag{28}$$

The next two axioms rule the interaction of event parthood and causal dependency with time. They state that the period of occurrence of an event is included in the period of occurrence of any of its super-events:

$$\mathsf{EP}(x,y) \wedge \mathsf{ETI}(x,i_x) \wedge \mathsf{ETI}(y,i_y) \rightarrow \mathsf{TINC}(i_y,i_x) \tag{29}$$

and that the period of occurrence of an event starts before the period of occurrence of any event that causally depends on it:

$$\mathsf{EC}(x,y) \wedge \mathsf{ETI}(x,i_x) \wedge \mathsf{ETI}(y,i_y) \rightarrow \mathsf{TIP}(i_y,i_x) \tag{30}$$

Finally, a fragment that narrates an event x narrates any super-event of x :

$$\mathsf{EP}(x,y) \wedge \mathsf{Ref}(z,x) \rightarrow \mathsf{Ref}(z,y) \tag{31}$$

In the following, we report an example of a simple narrative composed of two events, and we show a graph-based representation of this narrative using the axioms defined above. The narrative is based on the following text from Wikipedia: "Between 1886 an 1888 Gustav Klimt, along with his brother Ernst and his friend Hans Makart, painted the murals in the Burgtheater of Vienna. In 1888, Klimt received the Golden Order of Merit from Emperor Franz Josef I of Austria for his contributions to murals painted in the Burgtheater in Vienna"[3]. This text constitutes the *narration* of the narrative, and is expressed in our ontology using the predicate Nar. The *media object* that contains the narration is the section of the Wikipedia page from which the text was extracted, and it is represented using the predicate MObj.

The *fabula* of the narrative is composed of two events, each represented using the predicate Ev. The first event, "Klimt paints murals in the Burgtheater", is represented by the following textual fragment (MOFrag): "Between 1886 an 1888 Gustav Klimt, along with his brother Ernst and his friend Hans Makart, painted the murals in the Burgtheater of Vienna". This event has exactly one time interval (1886-1888), and one place (Vienna). Furthermore, the event has four participants: three people (Gustav Klimt, Ernst Klimt and Hans Makart) and one building (the Burgtheater). A graph-based representation of this event is shown in Figure 3.

[3]Text from the English Wikipedia, https://en.wikipedia.org/wiki/Gustav_Klimt

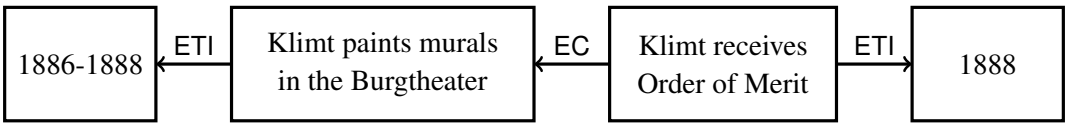

Figure 4. A graph-based representation of two events and the causality relation between them

The second event, "Klimt receives Golden Order of Merit", is represented by the following textual fragment (MOFrag): "In 1888, Klimt received the Golden Order of Merit from Emperor Franz Josef I of Austria for his contributions to murals painted in the Burgtheater in Vienna". This event has exactly one time interval (the year 1888), and one place (Vienna). The event has three participants: Emperor Franz Josef I, Gustav Klimt and the Golden Order of Merit. The two events are connected through the binary predicate EC, which expresses a causality relation between them. This relation is shown in Figure 4.

6.3. The axioms of the NOntDL ontology

Table 2 lists the unary and binary predicates of the NOntDL ontology. A graphical representation of these predicates is provided in Figure 5.

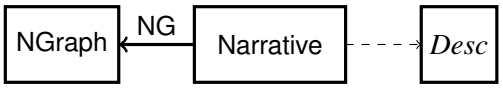

Figure 5. A view of the NOntDL ontology

In the following, we list all the axioms holding on the unary and binary predicates of NOntDL .

The two unary predicate symbols are pairwise disjoint:

$$A(x) \rightarrow \neg B(x) \tag{32}$$

Narratives and graphs are one-to-one:

$$\mathsf{NG}(n, g_1) \wedge \mathsf{NG}(n, g_2) \rightarrow g_1 = g_2 \tag{33}$$

$$\mathsf{NGraph}(g) \rightarrow (\exists n)\mathsf{NG}(n, g) \tag{34}$$

$$\mathsf{NG}(n_1, g) \wedge \mathsf{NG}(n_2, g) \rightarrow n_1 = n_2 \tag{35}$$

$$\mathsf{Nrt}(n) \rightarrow (\exists g)\mathsf{NG}(n, g) \tag{36}$$

A digital library is any KB that includes the above axioms and a set of assertions that connect each narrative to the corresponding NGraph through the NG property. As such, these assertions link the digital library to the graphs containing the formal representations of the narratives that are part of it.

Table 2
Unary and Binary Predicates of NOntDL

Unary Predicates	
Nrt(n)	n is a narrative
NGraph(g)	g is a narrative graph
Binary Predicates	
NG(n, g)	narrative n has graph g

6.4. Temporal reasoning in narratives

As stated in Section 4, we represent time in narratives using intervals. Sometimes, the time points giving the beginning and the end of such intervals are known, and the total ordering relation between time points can be used to express and reason over temporal knowledge in a narrative. However, this is not always the case: in many situations, only the relative relation between intervals is known, such that an event occurs before, or during another event. In these cases, a relative form of representation is the only viable option. We therefore need a conceptualisation of time that supports both time points and intervals, and absolute and relative relations between them.

Our conceptualisation includes both time instants and time intervals, along with the following relations:

- two functions connecting a time interval to its beginning and ending time instants, respectively;
- the total ordering between instants;
- the 13 jointly exhaustive and pairwise disjoint relations in Allen's algebra [9] capturing all possible ways in which two intervals can stand to each other in relative terms. In what follows, we shall call these 13 relations *basic temporal relations* (BTRs, for short). They are as follows: equals (abbreviated as e); before (b); after (bi); meets (m); metBy (mi); overlaps (o); overlappedBy (oi); during (d); include (di); starts (s); startedBy (si); finishes (f); finishedBy (fi).

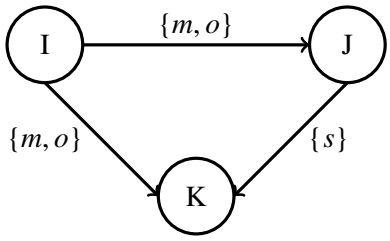

Figure 6. An example of QTK network

In the previous Section, two more relations between time intervals have been introduced, named in $\mathcal{L}_n$ by the TINC and TIP predicate symbols. These relations can be expressed as the union of BTRs as follows (for simplicity we abuse notation and use the predicate symbols also for the respective relation):

$$\mathsf{TINC} = \cup \{e, d, s, f\} \tag{37}$$

$$\mathsf{TIP} = \cup \{b, m, o, di, fi\} \tag{38}$$

Reasoning over Allen's temporal relations has been extensively studied in the literature. For reasons of space, we hereby report only the results of these studies that are relevant to the present context. The interested reader may consult e.g.[36] for a general treatment of this subject, and [37] for a discussion of temporal reasoning in the context of Semantic Web languages and technologies.

Following Allen's seminal work, the relations between the time intervals in a narrative are maintained in a network, which will be called Qualitative Temporal Knowledge network (QTK for short). The nodes of a QTK represent the time intervals in the narrative, while the arcs represent relations between the intervals corresponding to the conjoined nodes. The arcs are labelled with non-empty sets of BTRs, and each such set represents the union of its member relations. Specifically, an arc between nodes I and J is labelled by a set L of BTRs if and only if the temporal knowledge stored in the network implies that I and J are related by one of the relations in L. For example, the QTK given in Figure 6 stores knowledge about three intervals I, J and K, such that I meets or overlaps both J and K, while J starts K.

At the beginning, a QTK is empty. When a set of relations R between two intervals I and J must be asserted, two nodes corresponding to I and J are created, and the arc between them is added, labelled by R. Now suppose relation set S between nodes J and K needs to be asserted (such as $\{s\}$ in Figure 6). Correspondingly, node K is added to the network and S is used as a label of the arc connecting J and K. But node K must also be connected to all the other nodes of the network by adding the corresponding arcs, each with the appropriate label. In the example, K must be connected to node I with the appropriate label. In absence of any knowledge, the label is the complete set of Allen's relations, meaning that I and K can be related in *any possible way*. However, the known relations between nodes I and J and between J and K may *restrict* the possible relations between I and K. In order to compute these restrictions, *composition* rules are used. A composition rule is a statement about three intervals I, J and K. The statement has a *premise* and a *conclusion* as follows:

1. the premise gives a set of temporal relations between intervals I and J, and a set of temporal relations between intervals J and K;
2. the conclusion gives a set of temporal relations between intervals I and K.

The meaning of a composition rule is the following: if the relations in the premise hold between nodes I and J and between nodes J and K, then only one of the relations in the conclusion may hold between nodes I and K. Every time a QTK network is updated with additional knowledge, the composition rules are applied in order to restrict the relations labelling the arcs of the network.

Given that there are $2^{13} - 1$ non-empty temporal relation sets, and that there is a different composition rule for every pair of such sets, there are millions of composition rules. However, composition rules enjoy a nice mathematical property: the conclusions of the rules having non-singleton premises can be efficiently computed from those of the rules having singleton premises. Since the latter kind of rules are of the order of dozens (they are given in [9]), we have a method to efficiently compute the label of any arc of a QTK. However, since the number of possible labels grows exponentially with the number of labelled arcs, and labels need to be re-computed at each update, it may take an exponential amount of time to compute the QTK resulting from an update. This combinatorial explosion is one problem with QTKs.

A second problem is given by the rise of inconsistencies in QTK. These inconsistencies can

be detected by applying *path consistency* [36], a technique based on the application of the iterative formula for computing R(I,J)$^{n+1}$, that is the label on the arc between any two nodes I and J at step $n+1$, given the labels between any two nodes at step n. The formula is given by:

$$R(I,J)^{n+1} = R(I,J)^n \cap (\cup_{\mathrm{K}}(R(I,K)^n \circ R(K,J)^n)) \quad (39)$$

where $\circ$ denotes the composition of sets of BTRs. The formula is applied iteratively to a QTK until a fixed point is reached. If some label equals the empty set, then the QTK is inconsistent. Otherwise, the QTK resulting from path consistency contains a set of labels that is no larger than the set of labels of the initial QTK, and that embodies the temporal knowledge currently held in the network. Path consistency can achieve its task in a polynomial amount of time, therefore the second problem does not prevent the efficient management of a QTK.

In order to address the former problem, it is necessary to seek tractable sets of temporal relations, i.e. sets $\mathcal{T}$ including disjunctions of BTRs such that $\mathcal{T}$ is closed under intersection and composition, so that the application of path consistency always yields a relation in $\mathcal{T}$. This property prevents the combinatorial explosion of the time needed to compute a QTK following an update, while guaranteeing detection of inconsistencies. In order to perform temporal reasoning over narratives, we have derived a tractable set of temporal relations including the 13 Allen's BTRs and the disjunctions TINC and TIP that we need in order to axiomatise narratives, as explained in Section 6.5. This set, which we call $\mathcal{T}_n$, only includes 81 disjunctions; in the remaining part of this Section we briefly describe its composition and the way it has been derived.

6.5. Minimal tractable set of BTRs

We started from the minimal tractable set of BTRs computed in [37]. The set consists of 28 relations, including the 13 primitive ones plus 15 disjunctions. This set includes the TINC disjunction $\{b, m, o, di, fi\}$ representing the precedence relation, but unfortunately it does not include the TIP disjunction $\{e, d, s, f\}$ representing the inclusion relation. Therefore, this set is not sufficient for our purposes.

In order to solve this issue, we re-computed the minimal tractable set that includes TIP and TINC, using the path consistency algorithm described and implemented in [37].

The path consistency algorithm starts from an initial set of relations, and from the known transitivity table expressing their compositions. Each time a composition results in a new disjunction not present in the set, the algorithm adds a new row to the transitivity table and computes the composition between this disjunction and each other relation. When no new disjunctions are generated, the execution of the algorithm is stopped and the resulting set of relations is returned to the user. In our case, the initial set given as input to the algorithm contains the 13 primitive BTRs, plus TIP and TINC. At the end of the process, the resulting set contains 81 relations[4].

In order to reason on these 81 relations, it is necessary to explicitly express as rules all the possible compositions and intersections between each pair of relations contained in the set. In theory, this process should yield 6,561 composition rules plus 6,561 intersection rules, for a total of 13,122 rules. In practice, however, many rules can be safely removed because they involve the disjunction of all basic relations. This disjunction always holds between two intervals, thus it does not add any new information to the graph. By removing the rules involving this disjunction, the final number of rules is reduced to 7,671.

6.6. Defining and axiomatising temporal primitives

We can now complete the expression of NOnt by introducing and axiomatising the symbols for temporal representation and reasoning.

Table 3 gives the unary and binary temporal predicate symbols. As the Table shows, $\mathcal{L}_n$ also provides time points, for usability in realistic contexts. Consequently, the symbols modelling the ordering of time points and those for linking intervals to their beginning and ending time points are added as well. The binary predicate T stands

[4] These relations are defined at the beginning of the SWRL rule file at the following address: https://dlnarratives.eu/ontology/swrl-rules.owl

Table 3
The temporal predicate symbols of NOntNar

Unary Predicate Symbols	
Ev(e)	e is an event
Interval(t)	t is a time interval
Fab(f)	f is a fabula
Nar(a)	a is a narration
MObj(o)	o is a media object
MOFrag(r)	r is a media object fragment
Binary Predicate Symbols	
IB(t, p)	interval t begins at point p
IE(t, p)	interval t ends at point p
$p_1 < p_2$	point p_1 precedes point t_2
$p_1 = p_2$	point p_1 is equal to point t_2
$p_1 > p_2$	point p_1 follows point t_2
T(t_1, t_2)	interval t_1 is in relation T with interval t_2
$\perp$	the empty relation symbol

for each of the 81 binary predicate symbols that are one-to-one with the relations in $\mathcal{T}_n$, allowing users to exploit the full power of the temporal language in manipulating narratives. Finally, the special predicate symbol $\perp$ stands for the empty relation.

The following axioms provide domain, range and cardinality of the symbols linking intervals and time points:

$$\mathsf{Interval}(x) \rightarrow (\exists b)\mathsf{IB}(x,b) \quad (40)$$

$$\mathsf{IB}(x,b) \rightarrow \mathsf{Interval}(x) \wedge \mathsf{TPoint}(b) \quad (41)$$

$$\mathsf{IB}(x,b_1) \wedge \mathsf{IB}(x,b_2) \rightarrow b_1 = b_2 \quad (42)$$

$$\mathsf{Interval}(x) \rightarrow (\exists e)\mathsf{IE}(x,e) \quad (43)$$

$$\mathsf{IE}(x,e) \rightarrow \mathsf{Interval}(x) \wedge \mathsf{TPoint}(e) \quad (44)$$

$$\mathsf{IE}(x,e_1) \wedge \mathsf{IE}(x,e_2) \rightarrow e_1 = e_2 \quad (45)$$

The axioms on the symbols for ordering time points are not given, since the corresponding relations are constants and are available in any implementation.

The axioms on the symbols standing for the relations in $\mathcal{T}_n$ are given by the following sets of formulas:

- the set $\mathcal{C}_t$ containing the sentences for the composition of the temporal predicate symbols, each having the form

$$\mathsf{R}_1(x,y) \wedge \mathsf{R}_2(y,z) \rightarrow \mathsf{R}_3(x,z) \quad (46)$$

where each R_i is a temporal predicate symbol;
- the set $\mathcal{I}_t$ containing the sentences for the intersection of the temporal predicate symbols, each having the form

$$\mathsf{R}_1(x,y) \wedge \mathsf{R}_2(x,y) \rightarrow \mathsf{R}_3(x,y) \quad (47)$$

where each R_i is a temporal predicate symbol, and R_3 can be $\perp$;
- the set $\mathcal{P}_t$ containing the sentences relating the symbols of the 13 BTRs on time intervals and those on time points. Each such sentence is an if-and-only-if statement, expressing equivalence of one of the Allen's BTRs with a conjunction of atoms on the symbols $<$, $=$, and $>$ on time points. All these share the sub-formula

$$\mathsf{IB}(x,b_x) \wedge \mathsf{IE}(x,e_x) \wedge \mathsf{IB}(y,b_y) \wedge \mathsf{IE}(y,e_y) \quad (48)$$

which binds the six involved variables in the appropriate way. By abbreviating this formula as $\alpha(x,y)$, the sentence for the before BTR is given by:

$$\alpha(x,y) \rightarrow (\mathsf{before}(x,y) \equiv e_x < b_y) \quad (49)$$

which is equivalent to the two implications

$$(\alpha(x,y) \wedge \mathsf{before}(x,y)) \rightarrow e_x < b_y \quad (50)$$

$$(\alpha(x,y) \wedge e_x < b_y) \rightarrow \mathsf{before}(x,y) \quad (51)$$

As we have seen in Section 6.4, the binary predicate TIP representing a causality relation between two events implies that the corresponding temporal intervals are linked by one of the relations from the set $\{b, m, o, di, fi\}$. For example, suppose that for the two events reported in Figure 4 we do not know the exact beginning and end of the temporal intervals, but we know that the first event causes the second one. From this knowledge, we can infer that the BTR existing between the two time intervals must be one of the relations from the aforementioned set.

7. Implementing NOnt using Semantic Web technologies

Ontologies have long been recognised to be a crucial component of the Semantic Web [38]. The recommendation of languages for expressing ontologies is a core activity of the World Wide Web Consortium, which has produced a whole family of powerful such languages, collectively known as Ontology Web Language (OWL for short) [39], directly derived from Description Logics. The OWL family has now reached its second generation, OWL 2. It is therefore natural to consider the most expressive decidable language of the OWL family, OWL 2 DL, as a candidate for implementing the narrative ontology NOnt.

In this respect, unary predicate symbols would be implemented as OWL 2 DL classes, while binary predicate symbols would be implemented as OWL 2 DL object or data properties, depending on whether the range of a property is a class or a datatype. A wide array of datatypes are also available in OWL 2 DL, including the XML Schema datatype `dateTime`, which would be a most natural candidate for the implementation of time points. Based on this correspondence, the axioms of NOnt would have to be translated into OWL 2 DL axioms, by relying on the rich variety of operators that OWL 2 DL offers to this end. Before considering such translation, however, there are two immediate reasons why OWL 2 DL is not sufficient for implementing NOnt:

1. Properties corresponding to the EC and EP predicate symbols would have to be declared as irreflexive and transitive, to correctly reflect axioms 25 to 28 of NOnt. However, transitive properties are composite in an OWL 2 DL ontology, and as such they cannot be declared to be irreflexive, so as not to violate the global restrictions on the axioms of an OWL 2 DL ontology [40].
2. Path consistency requires axioms for the composition of temporal properties (given in set $\mathcal{C}_t$). These axioms can be expressed in OWL 2 DL as complex role inclusions. Now, the properties that occur in the right-hand side of complex role inclusions are composite, and this would prevent the expression of important axioms on these properties, for instance the axioms stating disjointness from other properties.

Furthermore, declaring the composition of the 81 temporal properties in $\mathcal{T}_n$ would require thousands of complex role inclusion axioms, and it would most certainly be impossible to avoid circular definitions, as required by the global restrictions on the axioms of an OWL 2 DL ontology[5].

An alternative to OWL 2 DL, which has also been considered in [37], is the Semantic Web Rule Language (SWRL)[6], a language of the Semantic Web family for specifying Horn clauses [41].

In order to implement NOnt using SWRL, the axioms of the ontology must be expressed as definite program clauses (DPCs). In fact, most of these axioms are *already* DPCs (such as for instance the axioms in $\mathcal{C}_t$, in $\mathcal{I}_t$, or in $\mathcal{P}_t$). Some of the remaining axioms can be easily transformed into DPCs. This is the case of axioms that are implications with a conjunction in their consequent, such as axioms 8 to 13. Each such axiom is equivalent to DPCs that have as body the antecedent of the implication, and as head a different conjunct in the consequent of the implication. For instance, axiom

$$\mathsf{EP}(x,y) \rightarrow \mathsf{Ev}(x) \wedge \mathsf{Ev}(y) \tag{52}$$

is equivalent to the DPCs

$$\mathsf{EP}(x,y) \rightarrow \mathsf{Ev}(x) \tag{53}$$

$$\mathsf{EP}(x,y) \rightarrow \mathsf{Ev}(y) \tag{54}$$

In addition, axioms that have an equivalence in the head can also be easily transformed into DPCs. These axioms are of the form

$$B_1 \wedge \ldots \wedge B_n \rightarrow (A \equiv A') \tag{55}$$

like axioms 6 and 7. Each such axiom is equivalent to the pair of DPCs

$$B_1 \wedge \ldots \wedge B_n \wedge A \rightarrow A' \tag{56}$$

$$B_1 \wedge \ldots \wedge B_n \wedge A' \rightarrow A \tag{57}$$

as we have already argued concerning the axioms in $\mathcal{P}_t$. Finally, the reflexivity axiom for equality can be replaced by the DCP

$$\neg(x = x) \rightarrow \bot \tag{58}$$

[5] https://www.w3.org/TR/owl2-syntax/#Global_Restrictions_on_Axioms_in_OWL_2_DL

[6] https://www.w3.org/Submission/SWRL

which produces a contradiction whenever an irreflexive axiom is violated.

However, axioms containing negation (such as axiom 1) or the existential quantifier (such as axiom 14) are not trivially reduced to DPCs. The remaining part of this Section shows that these axioms can be dealt with using SWRL, which is chosen as the implementation language of the temporal representation aspects of NOnt.

Time points will be implemented as values of the `dateTime` datatype of XML Schema[7], thereby equating the unary predicate symbol TPoint with that datatype.

7.1. Eliminating negation

Since it does not appear in the body of any rule, negation can be handled without resorting to the techniques devised in Datalog, such as stratification [42]. A much simpler approach is indeed possible [43], which consists in introducing a new set of predicate symbols, called *complements,* that are one-to-one with the predicate symbols in $\mathcal{L}_n$ and that stand for the negation of the corresponding predicate symbols. Technically, for every predicate symbol P in $\mathcal{L}_n$, we introduce a new predicate symbol called the complement of P. As customary, the complement of the equality symbol $=$ will be denoted as $\neq$, while the complement of any other predicate symbol P will be denoted as $\overline{P}$. We then modify the set of NOnt axioms as follows:

1. replace any instance of the axiom schema 1

$$A(x) \rightarrow \neg B(x) \tag{59}$$

by the corresponding instance of the schema:

$$A(x) \rightarrow \overline{B}(x) \tag{60}$$

and add

$$A(x) \wedge \overline{A}(x) \rightarrow \bot \tag{61}$$

2. replace any instance of the axiom schema 2

$$P(x,y) \rightarrow \neg R(x,y) \tag{62}$$

by the corresponding instance of the schema:

$$P(x,y) \rightarrow \overline{R}(x,y) \tag{63}$$

and add

$$P(x,y) \wedge \overline{P}(x,y) \rightarrow \bot \tag{64}$$

By so doing, a new set of axioms is obtained, which is intuitively equivalent to the initial set, since the two sets state the same sentences in different ways.

7.2. Dealing with existential quantification

As it is well-known, the typical technique for eliminating existentially quantified variables from first-order formulae is Skolemisation. Skolemisation is performed by replacing every existentially quantified variable y in the scope of n universally quantified variables $x_1, \ldots, x_n$ with a term $f(x_1, \ldots, x_n)$ where f is a new function symbol.

However, Skolemisation cannot be applied to reduce a set of axioms to SWRL rules, because function symbols are not allowed in SWRL rules. As a consequence, the existentially quantified axioms of NOnt , which are:

$$\mathsf{Ev}(x) \rightarrow (\exists y)\mathsf{ETI}(x,y) \tag{65}$$

$$\mathsf{Nar}(x) \rightarrow (\exists y)\mathsf{Cont}(x,y) \tag{66}$$

$$\mathsf{MOFrag}(x) \rightarrow (\exists y)\mathsf{OF}(x,y) \tag{67}$$

$$\mathsf{Interval}(x) \rightarrow (\exists y)\mathsf{IB}(x,y) \tag{68}$$

$$\mathsf{Interval}(x) \rightarrow (\exists y)\mathsf{IE}(x,y) \tag{69}$$

cannot be transformed into SWRL rules and must therefore be expunged from the SWRL implementation of NOnt. The negative effect of this elimination can be mitigated by considering that the individuals denoted by the existential variables in the above axioms are all unique, as guaranteed by the corresponding cardinality axioms:

$$\mathsf{ETI}(x,y_1) \wedge \mathsf{ETI}(x,y_2) \rightarrow y_1 = y_2 \tag{70}$$

$$\mathsf{Cont}(x,y_1) \wedge \mathsf{Cont}(x,y_2) \rightarrow y_1 = y_2 \tag{71}$$

$$\mathsf{OF}(y_1,x) \wedge \mathsf{OF}(y_2,x) \rightarrow y_1 = y_2 \tag{72}$$

$$\mathsf{IB}(x,y_1) \wedge \mathsf{IB}(x,y_2) \rightarrow y_1 = y_2 \tag{73}$$

$$\mathsf{IE}(x,y_1) \wedge \mathsf{IB}(x,y_2) \rightarrow y_1 = y_2 \tag{74}$$

Moreover, the individuals implied by the first three axioms, i.e. the time interval of an event, the content of a narration, and the media object con-

[7] https://www.w3.org/TR/xmlschema-2

Table 4
Mapping of NOnt classes with reference ontologies

Class	Linked class
Nrt	subclass of E73 Information Object
Fab	subclass of E4 Period
Nar	subclass of F14 Individual Work
Ev	equivalent to E5 Event
MObj	subclass of F22 Self-Contained Expression
MOFrag	subclass of E90 Symbolic Object
Interval	equivalent to E52 Time-Span
	and to Proper Interval of OWL Time

Table 5
Mapping of NOnt properties with reference ontologies

Property	Linked property
FN	subproperty of P148 has component
FE	subproperty of P9 consists of
Cont	subproperty of R9 is realised in
OF	subproperty of R15 has fragment
Ref	subproperty of P129 is about
EP	subproperty of P9 consists of
EC	superproperty of P15 was influenced by
ETI	equivalent to P4 has time-span
EPartic	equivalent to P12 occurred in the presence of
EPlace	equivalent to P7 took place at

taining a fragment, are all known at the time when the corresponding ETI, Cont, and OF atoms are asserted, therefore we will design the interface of the system in a way that forces the user to specify those individuals.

The situation is different for the last two axioms: the starting and ending points of a temporal interval may not be known at the time when the interval is asserted, and this is in fact the reason why NOnt allows the representation and reasoning about qualitative temporal knowledge. In the last two cases, then, failing the user to provide a data value for each of these points, the system will force temporal constants for them, using these constants as placeholders for the corresponding values about which knowledge can be expressed or inferred by the system.

7.3. The Ontology Mapping

The first requirement we took into account to develop our ontology was its semantic interoperability. Semantic interoperability is a two-way concept: on the one hand, we aim at widening the usage of our ontology for narratives, by making it re-usable; on the other, we aim at re-using as much as possible of existing ontologies in developing our own. A natural candidate of this latter category is the CIDOC CRM ontology [44], an ISO standard largely employed in the Digital Libraries and Cultural Heritage domains. The CRM includes temporal entities for capturing time-dependent concepts such as events; moreover, its harmonisation with the FRBR ontology, known as FRBRoo [45], provides fundamental notions for the modelling of text, such as expressions and expression fragments. To represent the temporal dimension, we also integrated NOnt with OWL Time [46], a domain ontology recommended by the W3C for the representation of time.

Tables 4 and 5 report the mapping between NOnt and the three reference ontologies (CIDOC CRM, FRBRoo, and OWL Time), for classes and properties respectively. In the tables, the classes starting with E and the properties starting with P are from the CIDOC CRM, while the classes starting with F and the properties starting with R are from FRBRoo.

To create the mapping, we analysed the definitions of the classes and properties of the three reference ontologies. In particular, we took into account the following versions of the ontologies: CIDOC CRM 6.2.9[8], FRBRoo 3.0[9], OWL Time W3C Recommendation of 19 October 2017[10].

The current implementation of the ontology is available on our website, along with the set of SWRL rules that we have implemented for temporal reasoning[11].

In the following Section, we report two examples of applications that show how the ontology is being used in practice.

8. The Ontology in Practice

In this Section, we report how the ontology is being used in practice in the context of: (i) the Narrative Building and Visualising Tool (NBVT)

[8] http://www.cidoc-crm.org/Version/version-6.2.9
[9] http://www.cidoc-crm.org/frbroo/ModelVersion/frbroo-v.-3.0
[10] https://www.w3.org/TR/2017/REC-owl-time-20171019
[11] https://dlnarratives.eu/ontology

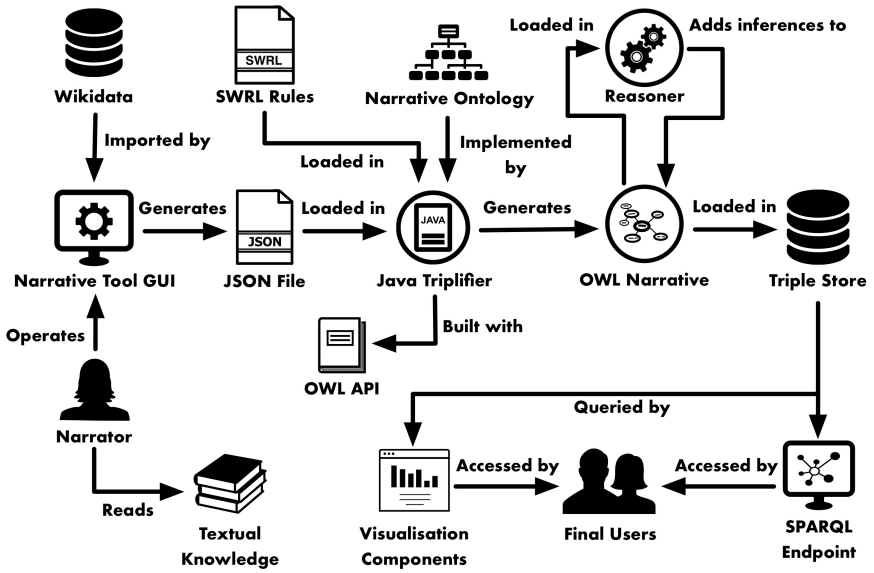

Figure 7. A view of the architecture of NBVT showing also the interactions between its components and its users

[6], and (ii) the Mingei European project, focused on the representation and preservation of Craft Heritage [7].

8.1. *Validating the Narrative Ontology using the Narrative Building and Visualising Tool*

As we explained in Section 1, in our vision, Digital Libraries should provide narratives as answers to the users' queries, which can help users to obtain a more complete knowledge on the subject of their searches. To reach this aim and validate NOnt, we developed the Narrative Building and Visualising Tool (NBVT)[12], a semi-automatic software based on Semantic Web technologies that allows creating narratives and visualising them in several ways [6]. Using NBVT, we have created narratives in four case studies[13]: in particular, two biographical narratives are about the life of Dante Alighieri, the major Italian poet of the late Middle Ages, and the life of the Austrian painter Gustav Klimt. Two other narratives have been created by a researcher in computational biology to narrate the history of the discoveries related to the giant squid, and the evolution of climate change due to human activity.

We have also performed an experiment to explore the integration of our tool with the Europeana digital library, by linking the narrative about Klimt to the digital objects of Europeana [1].

The software we developed is freely available for research aims, and access to the online version of the tool is available on request.

Figure 7 shows the architecture of NBVT, whose main components are the following:

1. *Narrative building interface*. It is used for creating, modifying or visualising a narrative, possibly representing knowledge that has been derived by reading some texts. The user operates through the Graphical User Interface (GUI) of the narrative-building tool, by manually inserting the narrative components (e.g. place, person, object) and, at the same time, importing resources from Wikidata[14] and images from Wikimedia Commons[15]. The created narrative is stored as an intermediate JSON[16] representation;
2. *OWL triplifier*. The triplifier transforms the JSON data into an OWL (Web Ontology Language) KB, encoded as an RDF graph. The organisation of the knowledge in the graph follows the structure defined in NOnt.
3. *Semantic reasoner*. It is used by the triplifier to infer new knowledge. The triplifier takes as input also a file with SWRL rules [47] that

[12] https://dlnarratives.eu/tool.html
[13] https://dlnarratives.eu/narratives.html
[14] https://www.wikidata.org
[15] https://commons.wikimedia.org
[16] https://www.json.org/json-en.html

are used to support reasoning on the narrative, as explained in Section 7.

4. *Triple store*. The triplifier stores the resulting graph into a Blazegraph triple store[17];
5. *Visualisation interface*. Finally, the user can access the knowledge stored in the triple store through a Web interface. The knowledge is extracted using SPARQL queries [48] and visualised in several ways, i.e. as a media-rich timeline, as a set of network graphs, and in table format.

In order to create narratives, the tool assists the user in creating the events that compose the narrative, attempting to minimise the cognitive and technical burden in the selection and identification of the involved entities. In the following list, we report the main functionalities of the tool, and each of these is explained through an example extracted from the narrative about the life of Gustav Klimt. These functionalities are:

- Defining the factual components that characterise the events, linking each event to persons, place, time, physical and conceptual objects through the appropriate semantic relations. For example, the creation of the portrait of Sonja Knips by Gustav Klimt is linked to the persons who participated in the event (Gustav Klimt and Sonja Knips) through the property P12 occurred in the presence of. Furthermore, the event is linked to the place in which it occurred (Vienna) through the property P7 took place at, and to the time span in which it occurred (1889) through the property P4 has time-span.
- Identifying the roles that persons played in an event. For example, in the event described above, Klimt played the role of *painter* and Sonja Knips played the role of *model*. In order to represent the role, in the ontology we define the class ActorWithRole that is linked to the class E39 Actor through the property hasRole and to the class E5 Event through the property hasParticipant.
- Defining the type of each event, choosing from a list of predefined options. In the ontology, we represent the types of event using subclasses of E5 Event. For example, the creation of the portrait of Sonja Knips has type E65 Creation.
- Storing the textual fragment, if any, providing a narration of the event in natural language. For example, the creation of the portrait of Sonja Knips is described in the narrative by the following textual fragment: "With a style reminding of the Belgian artist Fernand Khnopff, Klimt paints a lady from the Viennese élite, who was active with her husband in the circle of the Wiener Werkstätte. The face's plasticity contrasts with the soft inconsistency of the fluffy dress. In this diagonal composition, the evanescence of the chair, the book's red blur, the head surrounded by flowers, all anticipate the portraits of the golden period".[18]

Figure 8 shows the main interface of NBVT. The tool takes as input resources inserted manually by the user or imported automatically from Wikidata. It also initially imports a few default events from Wikidata, such as births, deaths, marriages, and company foundations. Then, the user adds the remaining events of the narrative one by one, by inserting the following information: (i) the title of the event; (ii) the start and end dates of the event; (iii) the event type; (iv) a set of entities imported from Wikidata or defined by the user (e.g. persons, locations, objects); (v) for each entity, one or more primary or secondary sources and, in the case of people, the role they played in the event; (vi) a textual description of the event; (vii) optional textual notes; (viii) one or more digital objects.

In addition to the creation of narrative, the tool supports its management as a digital object, enabling the visualisation and storage of the narrative in a well-understood format for local download or for sharing it on the Web. Concerning visualisation, we developed predefined SPARQL queries to implement the following functionalities:

- Visualising the fabula of the narrative on a timeline, including for each event: (i) a title, (ii) a textual description, (iii) the time span of the event, (iv) one or more related entities, and optionally (v) a set of related digital

[17] https://www.blazegraph.com

[18] Text from the English Wikipedia, https://en.wikipedia.org/wiki/List_of_paintings_by_Gustav_Klimt

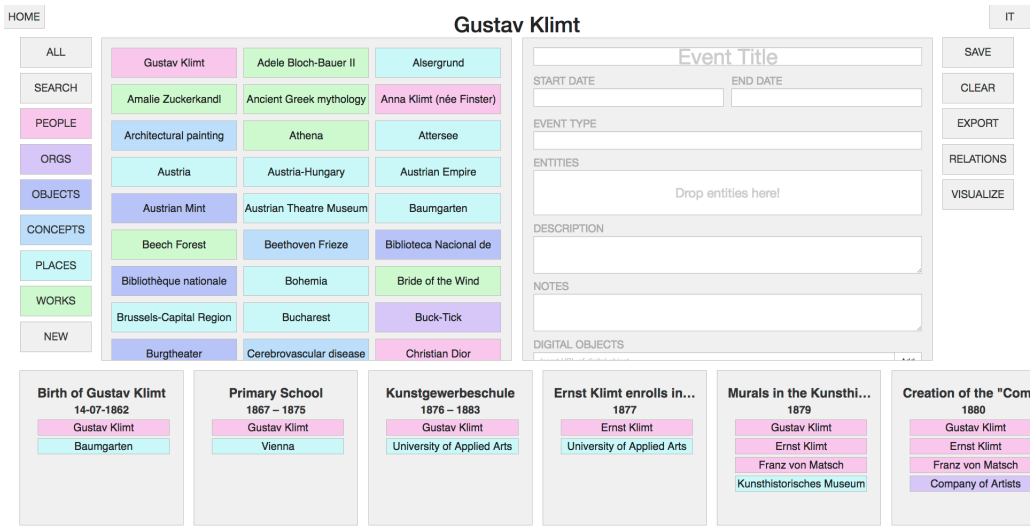

Figure 8. The main interface of NBVT

objects, and (vi) an image that illustrates the event.

- Visualising the events that happened in a range of time specified by the user, in form of table, exportable in CSV format.
- Visualising the entities related to a specific event, in graph form.
- Visualising the events related to a specific entity, in graph form.

In Appendix A, we report the code of SPARQL queries that implement the functionalities listed above.

Figure 9 shows the graph visualisation of the event of the painting of Klimt's murals in the Burgtheater in Vienna. Figure 10 shows an event in the timeline of Klimt's biography, including the textual description of the event, the secondary sources, the related entities, the related digital object and an image from Wikimedia Commons.

An experimental validation describing the satisfaction of the data modelling requirements of NBVT is reported in [6].

8.2. *Validating the Narrative Ontology in the Mingei European Project*

As the partner of the project expert in knowledge representation, we applied NOnt within the Mingei European project[19]. This project aims at representing and preserving knowledge about

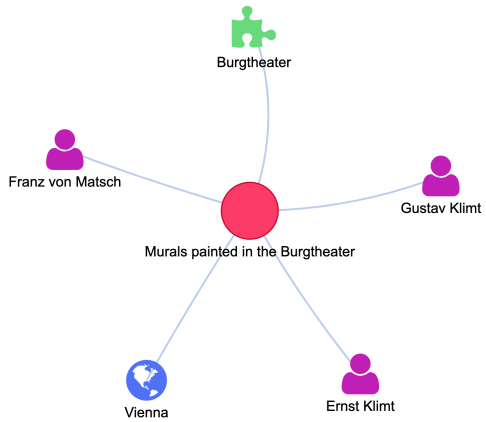

Figure 9. Graph visualisation of the painting of Klimt's murals in the Burgtheater in Vienna

both tangible and intangible aspects of craft as cultural heritage [7].

Heritage Crafts (HCs) involve craft artefacts, materials, and tools and encompass craftsmanship as a form of Intangible Cultural Heritage [49]. Intangible HC dimensions include dexterity, know-how, and skilled use of tools, as well as tradition and identity of the communities in which they are (or were) practiced. HCs are part of the history of the areas in which they flourish, and also have an impact upon their economy. The project selected three pilot themes that exhibit richness in tangible and intangible dimensions and are directly related to European history: (i) glass, represented by the Conservatoire National des Arts et Métiers (CNAM) in Paris, France, (ii) silk, represented by the Haus der Seidenkultur museum of

[19] http://www.mingei-project.eu

Figure 10. Timeline visualisation of the creation event of Klimt's Portrait of Sonja Knips

Krefeld, Germany, and (iii) mastic, represented by the Chios Mastic Museum in Greece.

In Mingei, we developed the Craft Ontology (CrO for short). CrO is an application ontology that uses NOnt as its core ontology. CrO adds to NOnt features that are useful in the context of the Mingei project, such as the notion of process of creation of a craft artefact, and the presentation of a narrative in the context of a museum exhibition. Moreover, CrO contains specific predicates for the chosen pilots. For instance, it specialises the Ev class into three Ev subclasses, each related to one of the three domains addressed by the project. These subclasses are further specialised on the basis of the domain they represent (e.g. Blowing is a subclass of Glass Activity, that is a subclass of Event). The added features and predicates are not reported here because they are specifically focused on the Craft Heritage domain, and they cannot be generalised to all possible kinds of narratives. To ensure interoperability, CrO has been integrated with the CIDOC CRM, FRBRoo and OWL Time, relying on the mapping presented in Section 7.3.

CrO is currently being populated by the experts of the three pilots using the ResearchSpace platform [50], but for reasons of intellectual property rights, we cannot report any of the narratives produced in Mingei in the present paper. Some of the narratives that are being created based on our ontology are about crafts, i.e. they narrate how craft masters blow glass, weave silk, and process mastic. In addition to these, there are also narratives that contextualise crafts, concerning for example the history of the town of Krefeld (a center of silk production in Germany), the history of the silk cloth in the Shrine of Charlemagne and its complex motif including the figure of an elephant, and the teaching of glass production at the CNAM in Paris, the history of the association of mastic growers on the island of Chios in Greece. These narratives will soon be made accessible to scholars and, via special applications, also to the general public visiting the museums involved in the crafts. An overview of the Mingei approach can be found in [51].

The main outcome of the Mingei project that is relevant to the present context is that the narrative ontology presented in this paper, forming the conceptual backbone of CrO, is adequate to support the representation of narratives about Craft Heritage. In particular, the scholars that are collaborating to the project have appreciated the two levels of representation offered by the ontology, i.e. the fabula and the narration. This double representation overcomes the main limitation of the previous projects addressing the same issue, namely the lack of a strong connection between the documents collected about the domain, which form the narration, and the role that these documents

play in the general representation of the craft. In sum, Mingei proves that narratives are a powerful digitisation tool.

9. Conclusions and Future Work

In the context of the Digital Humanities, and in particular of Digital Libraries focusing on the Cultural Heritage domain, the narration of major cultural or historical events is a very central point. In this article we have presented our research aiming at introducing narratives in Digital Libraries using Semantic Web technologies.

In order to do so, we have adopted a methodological approach similar to the one used for developing algorithms in Computer Science. We have followed the following phases: (i) conceptualisation, (ii) mathematical specification, (iii) development of an ontology using Semantic Web languages, and (iv) experimental implementation and validation of the ontology.

Before developing our conceptualisation of narrative, we have reviewed the Narratology and Artificial Intelligence literature in order to identify the formal components of narratives. We have initially expressed our conceptualisation in an informal way, and then we have formalised it using first-order logic. Finally, in order to represent the first-order logic specification through the languages of the Semantic Web, we have implemented an ontology for representing narratives, called Narrative Ontology (NOnt). The ontology is expressed using the OWL 2 DL and SWRL languages, and it has been implemented as an extension of three standard vocabularies: CIDOC CRM, FRBRoo and OWL Time.

On the basis of NOnt, we have developed the Narrative Building and Visualising Tool (NBVT), which has been used to create and visualise narratives in four case studies. Furthermore, we have also applied NOnt within the Mingei European project, in order to represent knowledge about Craft Heritage. We plan to conduct a full evaluation of the ontology as an output of Mingei.

10. Acknowledgments

This work has been partly carried out with funding from the Mingei project, European Union's Horizon 2020 research and innovation programme, grant agreement no. 822336.

References

[1] C. Meghini, V. Bartalesi, D. Metilli and F. Benedetti, Introducing Narratives in Europeana: Preliminary Steps, in: *Advances in Databases and Information Systems*, Springer, Cham, 2017, pp. 333–342, DOI: https://doi.org/10.1007/978-3-319-67162-8_33.

[2] Aristotle, *Poetics*, in *Oxford World's Classics*, A. Kenny, ed., Oxford University Press, Oxford, 2013, ISBN: 978-0191635809.

[3] V. Shklovsky, Art as technique, in: *Twentieth-Century Literary Theory*, K.M. Newton, ed., Palgrave, London, 1997, pp. 3–5, DOI: https://doi.org/10.1007/978-1-349-25934-2_1.

[4] M. Bal, *Narratology: Introduction to the theory of narrative*, University of Toronto Press, Toronto, 1997, ISBN: 978-0802096319.

[5] V. Bartalesi, C. Meghini and D. Metilli, A conceptualisation of narratives and its expression in the CRM, *International Journal of Metadata, Semantics and Ontologies* **12**(1) (2017), 35–46, DOI: https://doi.org/10.1504/IJMSO.2017.087692.

[6] D. Metilli, V. Bartalesi and C. Meghini, A Wikidata-based tool for building and visualising narratives, *International Journal on Digital Libraries* **20**(4) (2019), 417–432, DOI: https://doi.org/10.1007/s00799-019-00266-3.

[7] X. Zabulis, C. Meghini, N. Partarakis, C. Beisswenger, A. Dubois, M. Fasoula, V. Nitti, S. Ntoa, I. Adami, A. Chatziantoniou et al., Representation and Preservation of Heritage Crafts, *Sustainability* **12**(4) (2020), 1461, DOI: https://doi.org/10.3390/su12041461.

[8] A. Levitin, *Introduction to the design & analysis of algorithms*, Pearson Education, Addison-Wesley, Upper Saddle River, NJ, 2007, ISBN: 978-0132316811.

[9] J.F. Allen, Maintaining knowledge about temporal intervals, *Communications of the ACM* **26**(11) (1983), 832–843, DOI: https://doi.org/10.1145/182.358434.

[10] D. Herman, *Basic Elements of Narrative*, Wiley-Blackwell, Hoboken, NJ, 2009, ISBN: 978-1444305937.

[11] J.C. Meister, Narratology, in: *The Living Handbook of Narratology*, P. Hühn, J.C. Meister, J. Pier and W. Schmid, eds, University of Hamburg, Hamburg, 2014. http://www.lhn.uni-hamburg.de/article/narratology.

[12] R. McKee, *Story: Substance, structure, style, and the principles of screenwriting*, Harper Collins, New York, NY, 1997, ISBN: 978-0060391683.

[13] S.B. Chatman, *Story and discourse: Narrative structure in fiction and film*, Cornell University Press, Ithaca, NY, 1980, ISBN: 978-0801491863.

[14] R. Kowalski and M. Sergot, A logic-based calculus of events, in: *Foundations of knowledge base management*, Springer, Berlin, Heidelberg, 1989, pp. 23–55, DOI: https://doi.org/10.1007/978-3-642-83397-7_2.

[15] R. Miller and M. Shanahan, Some alternative formulations of the event calculus, in: *Computational logic: logic programming and beyond*, Springer, Berlin, Heidelberg, 2002, pp. 452–490, DOI: https://doi.org/10.1007/3-540-45632-5_17.

[16] E.T. Mueller, *Commonsense Reasoning: An Event Calculus Based Approach*, Morgan Kaufmann, Burlington, MA, 2014, ISBN: 978-0128014165.

[17] S.J. Russell and P. Norvig, *Artificial intelligence: a modern approach*, Pearson Education, London, 2010, ISBN: 978-0132071482.

[18] D. Davidson, *Essays on actions and events: Philosophical essays*, Vol. 1, Oxford University Press, Oxford, 2001, ISBN: 978-0199246267.

[19] Y. Raimond and S. Abdallah, The Event Ontology, Technical Report, Queen Mary University of London, 2007. http://motools.sourceforge.net/event.

[20] R. Shaw, R. Troncy and L. Hardman, LODE: Linking Open Descriptions of Events, in: *The Semantic Web*, Springer, Berlin, Heidelberg, 2009, pp. 153–167, DOI: https://doi.org/10.1007/978-3-642-10871-6_11.

[21] A. Scherp, T. Franz, C. Saathoff and S. Staab, F: A model of events based on the foundational ontology Dolce+ DnS ultralight, in: *Proceedings of the fifth International Conference on Knowledge Capture*, ACM, 2009, pp. 137–144, DOI: https://doi.org/10.1145/1597735.1597760.

[22] W.R. Van Hage, V. Malaisé, R. Segers, L. Hollink and G. Schreiber, Design and use of the Simple Event Model (SEM), *Web Semantics: Science, Services and Agents on the World Wide Web* **9**(2) (2011), 128–136, DOI: https://doi.org/10.1016/j.websem.2011.03.003.

[23] M. Doerr, C.-E. Ore and S. Stead, The CIDOC Conceptual Reference Model: a new standard for knowledge sharing, in: *Tutorials, posters, panels and industrial contributions at the 26th International Conference on Conceptual Modeling*, Australian Computer Society, Darlinghurst, 2007, pp. 51–56.

[24] M. Doerr, S. Gradmann, S. Hennicke, A. Isaac, C. Meghini and H. van de Sompel, The Europeana Data Model (EDM), in: *World Library and Information Congress: 76th IFLA general conference and assembly*, IFLA, 2010, pp. 10–15.

[25] A. Gangemi, N. Guarino, C. Masolo, A. Oltramari and L. Schneider, Sweetening ontologies with DOLCE, *Knowledge engineering and knowledge management: Ontologies and the Semantic Web* (2002), 223–233, DOI: https://doi.org/10.1007/3-540-45810-7_18.

[26] V. Bartalesi and C. Meghini, Formal Components of Narratives, in: *Italian Research Conference on Digital Libraries*, Springer, Cham, 2016, pp. 112–121, DOI: https://doi.org/10.1007/978-3-319-56300-8_11.

[27] E. Hyvönen, J. Takala, O. Alm, T. Ruotsalo and E. Mäkelä, Semantic Kalevala: Accessing Cultural Contents Through Semantically Annotated Stories, in: *Proceedings of the Cultural Heritage on the Semantic Web Workshop at the 6th International Semantic Web Conference (ISWC 2007), Busan, Korea, Nov, 2007*, 2007. https://seco.cs.aalto.fi/publications/2007/hyvonen-et-al-kalevala-2007.pdf.

[28] E. Hyvönen, P. Leskinen, M. Tamper, H. Rantala, E. Ikkala, J. Tuominen and K. Keravuori, BiographySampo: Publishing and Enriching Biographies on the Semantic Web for Digital Humanities Research, in: *European Semantic Web Conference*, Springer, Cham, 2019, pp. 574–589, DOI: https://doi.org/10.1007/978-3-030-21348-0_37.

[29] P. Mulholland and T. Collins, Using digital narratives to support the collaborative learning and exploration of cultural heritage, in: *13th International Workshop on Database and Expert Systems Applications*, IEEE, 2002, pp. 527–531, DOI: https://doi.org/10.1109/DEXA.2002.1045951.

[30] K. Fernie, J. Griffiths, P. Archer, K. Chandrinos, A. de Polo, M. Stevenson, P. Clough, P. Goodale, M. Hall, E. Agirre et al., PATHS: Personalising access to cultural heritage spaces, in: *2012 18th International Conference on Virtual Systems and Multimedia*, IEEE, 2012, pp. 469–474, DOI: https://doi.org/10.1109/VSMM.2012.6365960.

[31] A. Wolff, P. Mulholland and T. Collins, Storyspace: a story-driven approach for creating museum narratives, in: *Proceedings of the 23rd ACM conference on Hypertext and social media*, ACM, 2012, pp. 89–98, DOI: https://doi.org/10.1145/2309996.2310012.

[32] R. Damiano, V. Lombardo, A. Lieto and D. Borra, Exploring cultural heritage repositories with creative intelligence. The Labyrinth 3D system, *Entertainment Computing* **16** (2016), 41–52, DOI: https://doi.org/10.1016/j.entcom.2016.05.002.

[33] V. Bartalesi, C. Meghini and D. Metilli, Steps Towards a Formal Ontology of Narratives Based on Narratology, in: *7th Workshop on Computational Models of Narrative (CMN 2016)*, B. Miller, A. Lieto, R. Ronfard, S.G. Ware and M.A. Finlayson, eds, OpenAccess Series in Informatics (OASIcs), Vol. 53, Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik, Dagstuhl, Germany, 2016, pp. 4:1–4:10, DOI: https://doi.org/10.4230/OASIcs.CMN.2016.4.

[34] R.R. Smullyan, *First-order logic*, Vol. 43, Springer-Verlag, Berlin, Heidelberg, 1968, DOI: https://doi.org/10.1007/978-3-642-86718-7.

[35] H.J. Levesque and G. Lakemeyer, *The Logic of Knowledge Bases*, The MIT Press, Cambridge, MA, 2001, ISBN: 978-0262122320.

[36] J. Renz and B. Nebel, Qualitative spatial reasoning using constraint calculi, in: *Handbook of spatial logics*, Springer, Dordrecht, 2007, pp. 161–215, DOI: https://doi.org/10.1007/978-1-4020-5587-4_4.

[37] S. Batsakis, E.G. Petrakis, I. Tachmazidis and G. Antoniou, Temporal representation and reasoning in OWL 2, *Semantic Web* **8**(6) (2017), 981–1000, DOI: https://doi.org/10.3233/SW-160248.

[38] G. Antoniou and F. van Harmelen, *A Semantic Web primer*, The MIT Press, Cambridge, MA, 2004, ISBN: 978-0262018289.

[39] W3C OWL Working Group (ed.), OWL 2 Web Ontology Language Document Overview (Second Edition), W3C Recommendation, World Wide Web Consortium, 2012. https://www.w3.org/TR/owl2-overview.

[40] B. Motik, P.F. Patel-Schneider and B. Parsia, OWL 2 Web Ontology Language Structural Specification and Functional-Style Syntax (Second Edition), W3C Recommendation, World Wide Web Consortium, 2012. https://www.w3.org/TR/owl-syntax.

[41] J.W. Lloyd, *Foundations of Logic Programming*, Springer-Verlag, Berlin, Heidelberg, 1984, DOI: https://doi.org/10.1007/978-3-642-96826-6.

[42] S. Abitebul, R. Hull and V. Vianu, *Foundations of Databases*, Addison-Wesley, Reading, MA, 1995, ISBN: 978-0201537710.

[43] C. Meghini and M. Doerr, A first-order logic expression of the CIDOC Conceptual Reference Model, *International Journal of Metadata, Semantics and Ontologies* **13**(2) (2018), 131–149, DOI: https://doi.org/10.1504/IJMSO.2018.098393.

[44] M. Doerr, The CIDOC Conceptual Reference Module: An Ontological Approach to Semantic Interoperability of Metadata, *AI Mag.* **24**(3) (2003), 75–92.

[45] M. Doerr, C. Bekiari and P. LeBoeuf, FRBRoo, a conceptual model for performing arts, in: *2008 Annual Conference of CIDOC*, CIDOC – ICOM International Committee for Documentation, 2008, pp. 06–18. https://projects.ics.forth.gr/_publications/drfile.2008-06-42.pdf.

[46] J.R. Hobbs and F. Pan, Time ontology in OWL, W3C Candidate Recommendation, World Wide Web Consortium, 2020. https://www.w3.org/TR/owl-time.

[47] I. Horrocks, P.F. Patel-Schneider, H. Boley, S. Tabet, B. Grosof and M. Dean, SWRL: A Semantic Web rule language combining OWL and RuleML, W3C Member Submission, World Wide Web Consortium, 2004. https://www.w3.org/Submission/SWRL.

[48] W3C SPARQL Working Group (ed.), SPARQL 1.1 Overview, W3C Recommendation, World Wide Web Consortium, 2013. https://www.w3.org/TR/sparql11-overview.

[49] X. Zabulis, C. Meghini, N. Partarakis, D. Kaplanidi, P. Doulgeraki, E. Karuzaki, E. Stefanidi, T. Evdemon, D. Metilli, V. Bartalesi et al., What is needed to digitise knowledge on Heritage Crafts?, *Memoriamedia Review* (2019), DOI: https://doi.org/10.5281/zenodo.3723792.

[50] D. Oldman and D. Tanase, Reshaping the knowledge graph by connecting researchers, data and practices in ResearchSpace, in: *The Semantic Web – ISWC 2018*, Springer, Cham, 2018, pp. 325–340, DOI: https://doi.org/10.1007/978-3-030-00668-6_20.

[51] C. Meghini, The Mingei approach: we collect, connect and open up, Technical Report, Mingei Newsletter, 2020. http://www.mingei-project.eu/the-mingei-approach.

Appendix A. SPARQL Queries

This appendix contains examples of SPARQL queries that we have used to implement the visualisation functionalities of the NBVT tool. The queries extract knowledge from the narratives that have been represented according to our ontology model and are currently stored in our KB.

Query 1. Visualising the fabula of the narrative on a timeline, including for each event: (i) its title, (ii) its textual description, (iii) its time span, (iv) the related entities, (v) the related digital objects, (vi) an image that illustrates the event.

```
PREFIX : <http://dlnarratives.eu/ontology#>
PREFIX rdf: <http://www.w3.org/1999/02/22-rdf-syntax-ns#>
PREFIX rdfs: <http://www.w3.org/2000/01/rdf-schema#>
PREFIX ecrm: <http://erlangen-crm.org/current/>

SELECT ?eventTitle ?eventDesc ?startDate ?endDate
  ?entityName ?digitalObject ?image

WHERE {
  ?eventIRI rdf:type :Event.
  ?eventIRI ecrm:P1_is_identified_by ?eventAppellation.
  ?eventAppellation ecrm:P3_has_note ?eventTitle.
  OPTIONAL {
     ?eventIRI rdfs:comment ?eventDesc.
  }
  ?eventIRI ecrm:P4_has_time-span ?timeSpan.
  ?timeSpan :timeSpanStartedBy ?startDate.
  ?timeSpan :timeSpanFinishedBy ?endDate.
  ?eventIRI ecrm:P12_occurred_in_the_presence_of ?entityIRI.
  ?entityIRI ecrm:P1_is_identified_by ?entityAppellation.
  ?entityAppellation ecrm:P3_has_note ?entityName.
  OPTIONAL {
     ?digitalObject ecrm:P67_refers_to ?eventIRI.
  }
  OPTIONAL {
     ?image ecrm:P138_represents ?eventIRI.
  }
}
```

Query 2. Visualising the events that happened in a range of time specified by the user.

```
PREFIX : <http://dlnarratives.eu/ontology#>
PREFIX rdf: <http://www.w3.org/1999/02/22-rdf-syntax-ns#>
PREFIX ecrm: <http://erlangen-crm.org/current/>

SELECT ?eventTitle ?startDate ?endDate

WHERE {
  ?eventIRI rdf:type :Event.
  ?eventIRI ecrm:P1_is_identified_by ?eventAppellation.
  ?eventAppellation ecrm:P3_has_note ?eventTitle.
  ?eventIRI ecrm:P4_has_time-span ?timeSpan.
  ?timeSpan :timeSpanStartedBy ?startDate.
  ?timeSpan :timeSpanFinishedBy ?endDate.

  FILTER (?startDate > "STARTDATE" && ?endDate < "ENDDATE")
}
ORDER BY ?startDate
```

Query 3. Visualising the entities related to a specific event.

```
PREFIX : <http://dlnarratives.eu/ontology#>
PREFIX ecrm: <http://erlangen-crm.org/current/>

SELECT ?entityIRI ?entityName

WHERE {
  <event_IRI> ecrm:P1_is_identified_by ?eventAppellation.
  ?eventAppellation ecrm:P3_has_note ?eventTitle.
  <event_IRI> ecrm:P12_occurred_in_the_presence_of ?
     entityIRI.
  ?entityIRI ecrm:P1_is_identified_by ?entityAppellation.
  ?entityAppellation ecrm:P3_has_note ?entityName.
}
```

Query 4. Visualising the events related to a specific entity.

```
PREFIX : <http://dlnarratives.eu/ontology#>
PREFIX rdf: <http://www.w3.org/1999/02/22-rdf-syntax-ns#>
PREFIX ecrm: <http://erlangen-crm.org/current/>

SELECT ?eventIRI ?eventTitle

WHERE {
  <entity_IRI> ecrm:P1_is_identified_by ?entityAppellation.
  ?entityAppellation ecrm:P3_has_note ?entityName.
  ?eventIRI ecrm:P12_occurred_in_the_presence_of <entity_IRI
     >.
  ?eventIRI rdf:type :Event.
  ?eventIRI ecrm:P1_is_identified_by ?eventAppellation.
  ?eventAppellation ecrm:P3_has_note ?eventTitle.
}
```