Finanziato dall'Unione europea
NextGenerationEU

D3.1 - Observations Modeling: State of the Art

Project report

Emilio M. Sanfilippo & Roberta Ferrario
{name.surname@cnr.it}
CNR ISTC Laboratory for Applied Ontology (LOA), Trento & Catania

MITE
Make It Explicit

Project: Make it explicit: Documenting interpretations of literary fictions with conceptual formal models (MITE), PRIN 2022 PNRR, grant n. P20225MRTS, funded by European Union – Next Generation EU

Website: http://www.loa.istc.cnr.it/mite/

Work-package: WP3

Reference period: December 2023-May 2024 (B1-B3)

Release date: 31 May 2024

Abstract: This document reports on the work of WP3 in the context of the MITE project (**Task T3.1**) relative to the analysis of the state of the art for the formal modeling of scholarly statements.

Contents

1 Introduction

There is an increasing attention in the Digital Humanities (DH) for the design of tools addressing both the *hypothetical* and *partial* nature of data in the humanities, and the (analytic, empirical, etc.) methods or arguments that produce them [1, 2, 3, 4]. These two dimensions of the data could be due to factors like the presence of multiple sources (documents) presenting alternative and sometimes incompatible data about the same phenomenon, the lack or only partial availability of sources, or the presence of multiple viewpoints in the scholarly debate, just to mention some cases. To make some examples, in historical research, scholars might have at their disposal documents reporting *contrasting* data for the birth (death) date of a person, might *not* have this information at all, or might be in situations where they are unsure about the *reliability* of the sources. On the other hand, it is also common for scholars to *disagree* on a shared subject of study, one case for all being that of *aesthetic criticism* in fine arts, literature, and performing arts, among others [5].

To make sense of *observational data*, i.e., hypothetical and partial data that is produced through research investigation [6], and design information systems for their documentation and analysis, models suited for their representation are first of all needed. Ideally, as stressed by different parties [7, 8, 9], once a model in this direction has been developed, one can exploit different sorts of mechanisms to, e.g., compare alternative data about the same phenomena, making in this way sense of competing and conflicting perspectives, debates, and disagreements. In addition, in an era where AI systems are always more pervasive, an approach along these lines could be functional to make artificial agents "aware" of plural and contrasting viewpoints.

In the case of literary studies and criticism, which are at the focus of our investigation in MITE, observational data primarily cover data concerning *interpretations* of the "contents" of literary texts. However, they could also concern other sorts of data, such as information about the production contexts of texts. Interpretation is highly debated under different perspectives, including what it is, what its goals are, how it distinguishes from other activities for the analysis of texts etc. [10, 11]. For our purposes, we understand it as "the formulation of hypotheses about aspects of meaning in literary texts" [8, p.236]. Hence, for this report, we shall focus on observational data mainly in the sense of texts' interpretations – one may talk of *literary interpretational data* – because of the relevance that the latter have in literary contexts. However, the research work we report and the considerations we do are in principle applicable to observational data in a more general sense.

For the modeling and analysis of literary interpretations from a formal conceptual perspectives some of the challenges one needs to face are the following ones.

First, interpretations are not formally structured; they exist as scholarly texts in natural language which need to be processed for data extraction, documentation, and analysis. If one wishes to compare the interpretations provided by scholars across texts, an approach is to structure interpretations through formal models like ontologies [12]. The purpose of the latter in this case is to provide a set of well-defined categories and relations to formally represent interpretations as literary interpretational data, and therefore to make the data available for further analyses. Alternatively, one could think of approaches based on machine learning or natural language processing, including modern large language models (LLMs), that do not necessarily require the representation of textual interpretations through ontologies. However, these systems can notoriously suffer from opacity or even hallucinations in their functioning [13]. Differently, by fixing a conceptual reference system, ontologies allow for a *transparent* and *systematic* documentation and analysis. In Ciotti's words: "The creation of formal models based on an explicit conceptualization grants that all the critical discourses and analyses are firmly grounded in a common "setting" of the domain" [14]. Said that, our research in MITE is exploratory and open to the use of multiple techniques, as promising results will likely come in the next future from the development of *hybrid* systems combining multiple techniques in a coherent manner.

Second, literary scholars often adopt different critical frameworks for their interpretations, each one with its own terminology. In addition, the intended meaning of the terms they use is often left implicit or only vaguely defined [15, 16]. Hence, when analyzing texts, it might not be simple to understand whether experts use terms in the same way. With respect to what said above concerning the use of ontologies, the plurality of critical frameworks, terminologies, and concepts makes it hard to analyze the texts to understand their similarities and departing points. It is even challenging to think about the structure of an ontology for modeling literary interpretations; first, because there is no guarantee that a single ontology tuned for a scholarly text can be meaningfully generalized and applied to other texts, because of their conceptual and terminological variety; second, because it is hard to "operationalize" the concepts and theories used in literary studies for formal and computational analysis, considering, as said, that they may lack precise semantic boundaries.

Third, literary experts commonly support their interpretations through various sorts of arguments, justifying their theses through a chain of various premises. This is sometimes called the *rhetoric dimension* of scholarly debates [10]. The "logic" they adopt for their arguments is not always the same, nor it necessarily aligns with inference in classical logic, which – recall – is mostly used for knowledge representation purposes (for some logical representation of the scholarly debate, see [17, 12]). On the other hand, documenting and comparing arguments provided by different scholars is fundamental to understand how they relate to each other. For this goal, it remains challenging to provide a formal system that is uniform across multiple scholarly texts.

The report is structured as follows. Section 2 reports on the state of the art concerning the representation of observations,[1]; we compare and discuss the approaches in Sect. 2.5. Section 3 reports on existing works about the modeling of literary characters, whereas Sect. 4 reports on some projects in Computational Literary Studies (CLS). We conclude the report in Sect. 5 with some final remarks.

2 Observations

The development of formal models to represent observational data along with the methods (arguments) generating them received some attention during about the first decade of the years 2000s [17, 18, 19]; it was then left aside, while it is nowadays receiving a renewed attention with applications in scholarly areas like history [1], iconography [20], musicology [21], and literature [12, 22], among others. In this landscape, there is also an increasing attention to computational methods and applications driven by natural language processing and machine-learning aimed at the semi-/automatic analysis of texts as, e.g., what happens in the case of computational literary studies, as we will show in Sect. 4.

2.1 CRM Argumentation Model

The CRM Argumentation Model (CRMinf) [23] is an extension of the CIDOC-CRM ontology (ISO 21127; CRM for shortness) [24] designed to support the documentation of argumentation and the resulting data in domains including engineering, human and social sciences. For instance, empirical measurements are represented in the CRM Scientific Observations (CRMsci) module [25] by extending CRMinf. Both CRMinf and CRMsci are based on the work done by Doerr et al. [18].

The class *I1 Argumentation* plays a important role in the scope of the CRMinf. Following the diagram in Fig.1, an argumentation is an activity (E7 Activity) resulting in a belief (I2 Belief) which has a proposition set (I4 Proposition Set) as content (*what* the belief states) and a belief value (I6 Belief value; e.g., true, false, more or less certain, etc.). More precisely, *I1 Argumentation* "comprises the activity of making *honest* inferences or observations. An honest inference or observation is one in which the E39 Actor carrying out the I1 Argumentation justifies and believes that the I6 Belief Value associated with the resulting I2 Belief about the I4 Proposition Set is the correct value at the time that the activity was undertaken and that any I3 Inference Logic or methodology was correctly applied" [23, p.14] (emphasis is ours). In other terms, CRMinf assumes that agents carrying out argumentation activities do not cheat: "Following our principle of honest argumentation we assume that the inference maker believes him/herself at least from the end of the inference making process, which is that the inference is correct, but may be convinced at any later time that the inference is actually wrong, regardless of the truth of the conclusion" [18]. Note that according to CRMinf, a belief is a temporal entity, too, corresponding to "the period of time that an individual group holds a particular set of propositions to be true, false, or somewhere in between" [23, p.14].

In addition to the documentation of provenance (not shown in Fig. 1), CRMinf also allows documenting some aspects concerning the logic of an argumentation; in particular, that beliefs can be premises for other beliefs, as shown in Fig. 2.

The class *S4 Observation* can be used to represent, e.g., empirical measurements. In the context of CRMsci [25], this class is subsumed by *E13 Attribute Assignment*, which is the most general class in CRM for "the actions of making assertions about one property of an object or any single relation between two items or concepts" [26, p.70].

[1] We will interchangeably use the terms: statement, observation, claim.

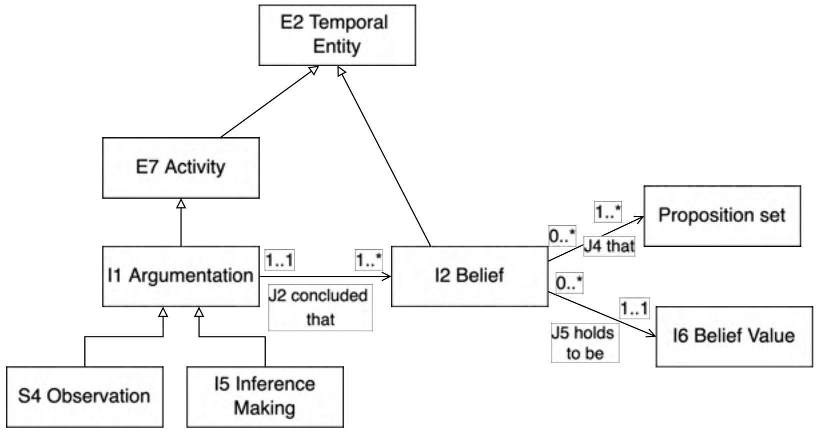

Figure 1: Partial view of CRM Argumentation Model (CRMinf)

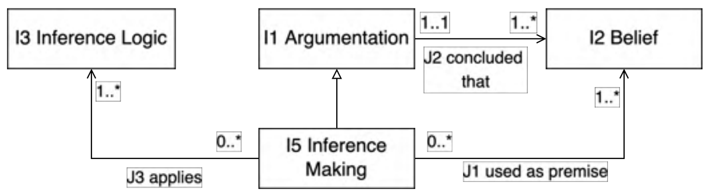

Figure 2: Partial view on the class I5 Inference making in CRMinf

The following Table 1 reports on some relations used in CRMinf.

Table 1: Some relations in CRMinf

Relation	**Description**
J1 used as premise	An inference making argumentation uses a belief as premise
J2 concluded that	An inference making argumentation uses a belief as conclusion
J3 applied	An inference making argumentation uses a certain system of logic to infer conclusions from premises
J7 is based on evidence from	It relates an event of belief adoption with an information object that is a source of evidence for the adopted belief

Remarks:

- Ontological status of beliefs: according to CRMinf, it is possible for a belief to be held by both an individual agent or by a group of agents. In the paper laying down the foundations of CRMinf, Doerr et al. [18] talk of beliefs in terms of mental entities: "We regard belief as a mental state that is determined by a human actor (crm:E39 Actor), a particular belief value, and a proposition, and that may exist for some time-span" [18, p.12]. At the theoretical level, this understanding of beliefs requires a model for belief sharing within groups of agents. The documentation of the model says: "An instance of I2 Belief comes into existence when an instance of I1 Argumentation concludes it [...]. Only one E39 Actor may hold a particular instance of I2 Belief, though the E38 Actor may, of course, be an instance of E74 Group. Such an instance of E74 Group may lose or gain members [...] without affecting the belief the group representatively maintains. The members supporting the common belief may not necessarily be individually convinced of it. This does not invalidate the belief of the Group" [23, pp.5-6].

In addition, that a belief is an agent's mental state raises two further concerns: first, how can a belief be accessed in an intersubjective manner, i.e., without "looking into agents' minds"? Second, what happens to beliefs when the agents holding them pass away?

In practice, the way out to deal with these two concerns in CRMinf is to rely on the modeling of proposition sets. Following the documentation: "This class [I4 Proposition Set] comprises the sets of formal, binary propositions that a I2 Belief is held about. It could be implemented as a named graph, a spreadsheet, or any other structured dataset. Regardless of the specific syntax employed, the effective propositions it contains should be made up of unambiguous identifiers, concepts of a formal ontology, and constructs of logic" [23, p.15]. Ontologically speaking, however, the two problems we raised still remain: if a proposition set stands for the "content" of a belief, once the belief stops existing, the proposition set stops existing as well; at least, if the proposition set is not treated as an entity that, once created, can exist independently from the belief that generated it.

- As said, in the general CRM ontology, *E13 Attribute Assignment* is the most general class for argumentation activities. The idea is to model assertions as attribution of properties: "For example, the class describes the actions of people making propositions and statements during certain scientific/scholarly procedures, e.g., the person and date when a condition statement was made, an identifier was assigned, the museum object was measured, etc. " [26, p.70]. Given that, as CRM recognizes, "E13 Attribute Assignment may possibly lead to a collection of contradictory values", one and the same entity is literally ascribed with incompatible properties according to the model.

2.2 Historical Context Ontology (HiCO)

The Historical Context Ontology (HiCO) [4] is an OWL ontology to represent scholarly statements in the humanities [4, 7].[2] In particular, HiCO is developed "for representing the *context* of a claim. [...] it addresses features characterising hermeneutical activities performed by scholars while generating new information (i.e. an interpretation act). [...] For instance, being created by somebody, or being created at a certain time, are events related to an artefact that are claimed by an agent at a certain time, motivated with usage of primary sources, and recorded in a secondary source (e.g. a cataloguing record)" (from the online documentation, see footnote 2; emphasis is ours).

At the core of HiCO stands the class *Interpretation Act* for the "hermeneutical activity performed by an agent in order to generate new information" (from the OWL file of HiCO), see Fig. 3.

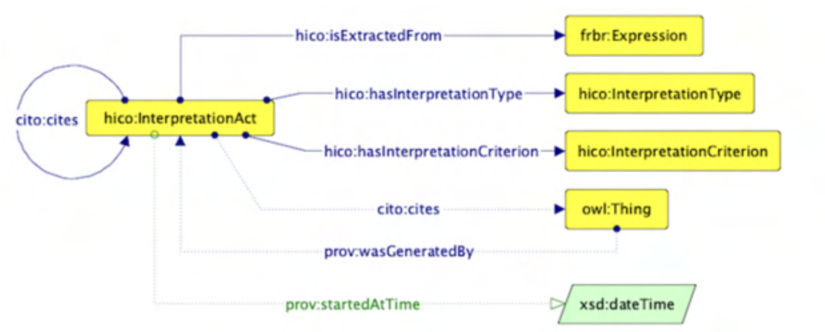

Figure 3: Class of Interpretation Act in HiCO (from the web documentation)

An interpretation act can be characterized by different elements to describe its context, e.g.:

- The classification of an interpretation, e.g., being an artwork attribution (*hico:hasInterpretationType*).
- The criteria or approach motivating an interpretation, e.g., the adoption of a bibliography concerning the research at stake (*hico:hasInterpretationCriterion*).
- The temporal extent of the interpretation, i.e. when it was claimed (*prov:startedAtTime*).

[2]Documentation: https://marilenadaquino.github.io/hico/.

- Cited sources of information, e.g. a bibliographic source (*cito:citesAsEvidence*).
- The source from which an interpreatation is extracted, e.g. a cataloguing record (*hico:isExtractedFrom*).
- Relations between interpretations, e.g., to tell that an interpretation refutes or agrees with another one (*cito:refutes*, *cito:agreesWith*, *cito:disagreesWith*).

As it can be seen from some of the modeling elements mentioned above, a number of existing ontologies are used in HiCO, including ontologies from the SPAR suite.[3]

We report in the following an example from the web documentation of HiCO.

Example: "The artwork called *The three Graces* has been attributed to Perruzzi Baldassare by cataloguers of the Federico Zeri Foundation around 1990. Bibliographic references support the claim. Another attribution ascribes the artwork to Luino Bernardino's school, and this is supported by a claim made by Christie's auction firm in 1994."

In terms of HiCO, the example is represented through the RDF triples in Fig. 4. In particular, the triples represent two instances of *hico:InterpretationAct*, i.e., *:39794-authorship-attribution-1* and *:39794-authorship-attribution-2*, along with some data to characterize them, e.g., the record source supporting the claim. The interpretation acts are also related through the relationship *cito:disagreesWith* to express a form of disagreement between them. Also, each interpretation act generates a creation event (as instance of the CRM's class *crm:E65_Creation*). In this manner, in the context of this example, the ontology represents the attribution of authorship for the artwork at stake to two different agents (*:baldassarre* and *:bernardino-school*).

```
@prefix cito: <http://purl.org/spar/cito/> .
@prefix crm: <http://www.cidoc-crm.org/cidoc-crm/>.
@prefix hico: <http://purl.org/emmedi/hico/> .
@prefix prov: <http://www.w3.org/ns/prov#> .
@prefix rdf: <http://www.w3.org/1999/02/22-rdf-syntax-ns#> .

:39794-creation-1 a crm:E65_Creation ;
    crm:P14_carried_out_by :baldassarre ;
    prov:wasGeneratedBy :39794-authorship-attribution-1 .

:39794-creation-2 a crm:E65_Creation ;
    crm:P14_carried_out_by :bernardino-school ;
    prov:wasGeneratedBy :39794-authorship-attribution-2 .

:39794-authorship-attribution-1 a hico:InterpretationAct ;
    hico:hasInterpretationType :authorship-attribution ;
    hico:hasInterpretationCriterion :bibliography ;
    cito:citesAsEvidence :book ;
    prov:startedAtTime "1990-01-01T00:00:00.001Z"^^xsd:dateTime ;
    hico:isExtractedFrom :37495-record ;
    cito:disagreesWith :39794-authorship-attribution-2 .

:39794-authorship-attribution-2 a hico:InterpretationAct ;
    hico:hasInterpretationType :authorship-attribution ;
    hico:hasInterpretationCriterion :auction-attribution ;
    prov:startedAtTime "1994-01-01T00:00:00.001Z"^^xsd:dateTime ;
    hico:isExtractedFrom :37495-record ;
    cito:disagreesWith :39794-authorship-attribution-1 .
```

Figure 4: Example based on HiCo (from the web documentation)

Table 2 reports on some relations used in HiCO to model interpretation acts.

Remarks:

- HiCO primarily focuses on the representation of the *context* of an interpretation act, that is, on meta-data of interpretation acts like their provenance. The ontology is not meant to model *what* is claimed by an interpretation act; e.g., in the case of the example in Fig. 4, the ontology is used in tandem with CRM to represent the attribution of authorship to a certain agent.

  Note however that in the proposed example the use of elements from CRM might result controversial. The class *crm:E65_Creation* is understood in CRM to represent spatio-temporal events leading to the creation of conceptual items or immaterial products (e.g., texts, music, images,

[3]SPAR suite of ontologies: http://www.sparontologies.net/.

Table 2: Some relations in HiCO

Relation	**Source**	**Description**
agrees with	CITO	To represent agreement relations between interpretation acts.
disagrees with	CITO	To represent disagreement relations between interpretation acts.
refutes	CITO	To represent relations of refuting between interpretation acts.
cites as evidence	CITO	"The citing entity cites the cited entity as source of factual evidence for statements it contains" (quoted from CITO; see example in Fig. 4)
was influenced by	PROV-O	"Influence is the capacity of an entity [...] to have an effect on the character, development, or behavior of another [...]".

etc.) (see [26, p.97]). On the other hand, in the example based on HiCO, *:39794-creation-1* and *:39794-creation-2* have only an *hypothetical* status by being related to interpretation acts that may lack veridicity.

- As we have seen, HiCO reuses relations from CITO ontology[4] like *cito:agreesWith*, *cito:disagreesWith*, and *cito:refutes*. These are useful to model relations between interpretation acts. However, since HiCO does not address the representation of statements generated by interpretation acts, dis-/agreements between the acts can be stated only from an "external perspective." At first glance, the assumption seems to be that since modelers know what interpretation acts assert, they can model the dis-/agreement between the acts, while the subject matter of the dispute is not explicitly conveyed in the model.

2.3 ICON

The ICON ontology[5] [20] aims at representing iconographic interpretations of visual artworks. It builds on the theoretical work of Erwin Panofsky as conceptual foundation to distinguish three levels of interpretation: (*i*) pre-iconographical interpretation, corresponding to the recognition of simple artistic motifs, e.g., the presence of a person or table in an artwork; (*ii*) iconographical interpretation, corresponding to the recognition of characters, places, events, symbols, personifications, etc.; e.g., a person that is recognized as being Jesus, or a composition of items standing for an allegory (the Annunciation); (*iii*) iconological interpretation, corresponding to the layer where the items in an artworks are recognized as manifestations of the underlying principles of a cultural context; e.g., the recognition of a specific genre-driven attitude towards female figures.

Figure 5 shows some of the core modeling elements of ICON (yellow classes), in particular *Recognition* being the main class for representing iconographical interpretations. Looking at the diagram, classes in colors other than yellow are borrowed from existing ontologies.

Following Sartini et al. [20], a recognition is "an interpretation act made by an agent (or interpreter, which can be a biological or electronic being) that links works of art to something related to their content. From a conceptual perspective, it is a *mental entity reflecting the agent's subjective point of view*. From a technical viewpoint, it is an n-ary predicate [...]" ([20, p. 10]; emphasis is ours). In particular, the authors stress that "each recognition is made on exactly one artwork, [and] involves exactly one agent."

Instances of the class *InterpretationDescription* are used to document recognitions on the same artwork. Relations like *cito:citesAsEvidence* can be used to model sources supporting a recognition, whereas *cito:givesSupportTo* relates recognitions, in particular, when one is used to support the claim made by another one, see Table 3.

[4] CITO: http://www.sparontologies.net/ontologies/cito.
[5] ICON: https://w3id.org/icon/docs/.

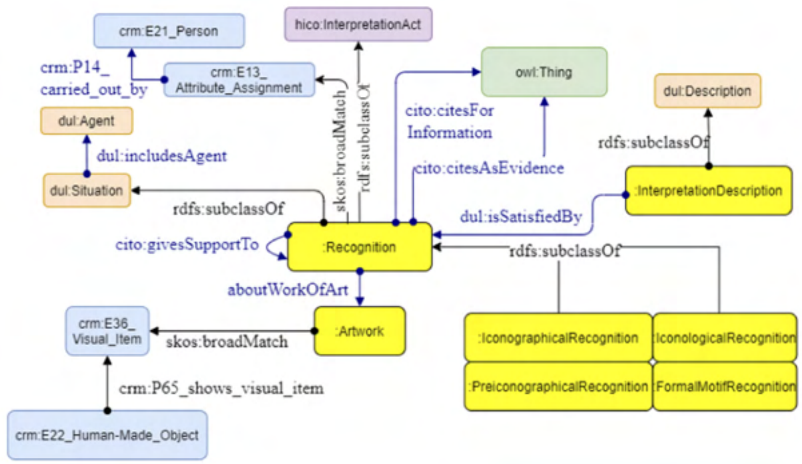

Figure 5: ICON's core modeling elements

The diagram in Figure 5 also shows the alignment of ICON with existing ontologies, including HiCO, CRM, and Dolce Ultralite (DUL).[6] According to the alignments, *Recognition* is represented as: (*i*) a subclass of *dul:Situation*, (*ii*) a subclass of *hico:InterpretationAct*, and (*iii*) a class that is more specific (*skos:broaderMatch*) than *crm:E13_Attribute_Assignment*. According to [20], SKOS' relations are used to guarantee a minimal alignment with ontologies like CRM whose conceptual framework differs from ICON under relevant extents. Also, recall that *situations* have been introduced in research related to DUL to represent n-ary relations in the expressivity of Semantic Web languages [27]. In recent works [28], situations are used to model scholarly claims, e.g., the attribution of authorship to texts.

Table 3 reports on some relations used in ICON.

Table 3: Some relations in HiCO

Relation	**Source**	**Description**
gives support to	CITO	It holds between recognitions. "The cited entity provides intellectual or factual support for the citing entity" (from CITO).
cites for information	CITO	"The citing entity cites the cited entity as a source of information on the subject under discussion (from CITO).

Remarks:

- As said, recognitions are conceived as agents' mental entities in ICON. At the same time, Sartini et al. [20] stress that recognitions can be generated by computer systems like computer vision algorithms. This would require some clarifications, in particular, in which sense recognitions as mental entities reflecting the subjectivity of interpreters can be attributed to computer systems.

  Also, in a mentalistic approach, similarly to what said with respect to CRMinf, it remains open the problem of how to access recognitions in an intersubjective manner. Following [20], if we understand it correctly, the way out of this issue is to rely on *interpretation descriptions* as documentations of "[c]oherent recognitions on the same artwork" [20, p. 10].

[6]DUL: http://ontologydesignpatterns.org/wiki/Ontology:DOLCE+DnS_Ultralite.

2.4 Wikidata

Wikidata is the knowledge-base laying nowadays in the background of Wikipedia and most of its sister projects [29]. Its data model builds on the representation of three sorts of things:[7]

- *Items*: either a class of entities (e.g., *human*, identifier Q5) or an individual entity (e.g., *Dante*, identifier Q1067);
- *Properties*: relationships connecting items, or items and values. For example, P31 stands for the relation of instantiation that can hold between *human* and *Dante*; P569 is the relation to be used to represent the birth date of *Dante* as a certain value, etc.
- *Statements*: these are entities used to make explicit the hypothetical and multi-perspectival dimension of data stored in Wikidata. This because Wikidata's purpose is not to represent *facts* involving domain entities but *claims* about such entities. For instance, representing that Dante was born in Florence is not to represent a fact of reality according to Wikidata; rather, it amounts to represent a statement about Dante, possibly along with various information about its provenance, reliability, etc.

The representation of Wikidata's statements plays a relevant role in the scope of our project, hence analyzing how they are intended and represented is relevant for our investigation.

Figure 6 represents the three statements in Wikidata about Dante's date of birth. Accordingly, the first statement attributes to Dante the birth date of 30 May 1265. This statement is supported by one reference according to which the attribution of birth date is stated in a section of a specific source, i.e., the *Philosophical Encyclopedic Dictionary*. The other two statements follow a similar structure for the documentation of references. In particular, the second statement is supported by two references, *Treccani's Enciclopedia on line* and *Croatian Encyclopedia*. Also, the third statement covers information about the time-span of the hypothetical birth date.

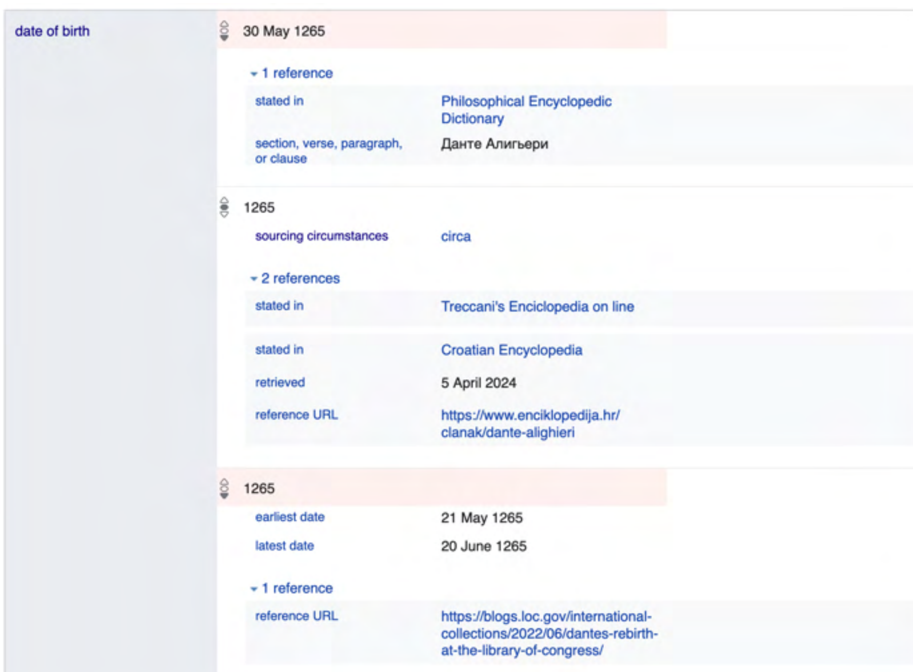

Figure 6: Example of statements in Wikidata (about Dante's birth date)

To better understand how the data is formally represented in Wikidata's statements, one can

[7]Documentation on the Wikidata's data model: https://www.mediawiki.org/wiki/Wikibase/DataModel; and primer documenation: https://www.mediawiki.org/wiki/Wikibase/DataModel/Primer. Recall that Wikidata identifies entities in its application domain through alphanumerical identifiers: Qn for items and statements, Pn for properties (relations).

explore them through the Wikidata Query Service[8] through SPARQL queries. The following query retrieves the three statements about Dante's date of birth shown in Fig 6:[9]

Query 1: SELECT * {wd:Q1067 p:P569 ?statement. ?statement ?predicate ?object }

Table 4 reports the data retrieved for the first statement in Fig 6, i.e., wds:Q1067-40D13CFD-B085-4999-B2B7-A5EB886EEB96 (we write in the table only wds:Q1067- for the sake of shortness), while Table 5 reports some IRIs used in Wikidata.[10] Recall that wd:Q1067 is the identifier for Dante in the scope of Wikidata.

Table 4: Results for query

statement	**predicate**	**object**
wds:Q1067-	wikibase:rank	wikibase:DeprecatedRank
-	prov:wasDerivedFrom	wdref:17a6256c731c3556a30fbe8fd424699fbda48f76
-	psv:P569	wdv:30307d95961f155dc65877a5e3c493ae
-	ps:P569	6 June 1265

Table 5: Some IRIs used in Wikidata

Prefixes	**IRI**
p	http://www.wikidata.org/prop/
psv	http://www.wikidata.org/prop/statement/value/
ps	http://www.wikidata.org/prop/statement/
pr	http://www.wikidata.org/prop/reference/
wdt	http://www.wikidata.org/prop/direct/
wd	http://www.wikidata.org/entity/
wikibase	http://wikiba.se/ontology#

We comment on the query and its results to understand the data:

- property *p:P569*: relates an individual item like Dante to the statement about its birth date.
- property *ps:P569*: relates a statement of birth date to a certain date-value like '6 June 1265';
- property *psv:P569*: relates a statement of birth date to an instance of the class *wikibase:TimeValue*, in this case, the entity with IRI wdv:30307-. Hence, wdv:30307- is the entity corresponding to the reification of the date-value '6 June 1265'. In this manner, one can represent further information about it, e.g., to tell that a date is represented according to a calendar, etc.[11]
- property *prov:wasDerivedFrom*: is taken from the PROV-O ontology[12] and is used to represent provenance information. By exploring the entity for the provenance (wdref:17a6-),[13] one finds that it is related through *pr:P248* (label: stated in) to *wd:Q4484349* (Philosophical Encyclopedic Dictionary), and through *pr:P958* (label: section, verse, paragraph, or clause) to the section of the Dictionary, as it is shown in the first statement of Figure 6.
- *property wikibase:rank*: is used to give ranks to statements. According to the documentation,[14] "ranks are not a way of asserting your view for a disputed value, but instead are used for communicating the consensus opinion for a statement."

[8] Wikidata Query Service: https://query.wikidata.org/.
[9] The SPARQL query is accessible at: https://w.wiki/9$CX.
[10] At the time of writing this report (May 2024), there is a mismatch in Wikidata between wds:Q1067-'s date provided in the graphical interface (Fig 6) and the date accessed through the query service.
[11] See the SPARQL query on wdv:30307-: https://w.wiki/A2ti.
[12] PROV-O: https://www.w3.org/TR/prov-o/.
[13] See the SPARQL query on wdref:17a6- https://w.wiki/A2tm.
[14] https://www.wikidata.org/wiki/Help:Ranking.

There are three types of rank:

– *normal* (*wikibase:NormalRank*): assigned to all statements by default. A normal rank provides no judgment or evaluation of a value's accuracy and therefore should be considered neutral.

– *preferred* (*wikibase:PreferredRank*): assigned to the most current statement or statements that best represent consensus (be it scientific consensus or the Wikidata community consensus).

– *deprecated* (*wikibase:DeprecatedRank*): used for statements that are known to include errors (i.e. data produced by flawed measurement processes, inaccurate statements) or that represent outdated knowledge (i.e. information that was never correct, but was at some point thought to be), as in the example in Table 4.

  The Wikidata Query Service stores two versions of a statement: one for any rank with value, qualifiers and references; one for "best" rank (*wikibase:BestRank*) with the value. A statement with best rank is a statement with preferred rank; if there is no preferred rank, the best rank corresponds to the statement with normal rank.

Remarks:

- Properties relating individual items to statements convey information about the *type* of statement. For instance, as we have seen above, *p:P569* relates Dante to a statement concerning its birth date. To make another example, *p:P19* relates Dante to a statement concerning its birth place, and so on.[15] In a sense, instead of providing a taxonomy for different types of statements, Wikidata provides alternative properties (with different IRIs) to distinguish between different sorts of information in a statement.

- What seems fundamental to distinguish two statements about the same item are: the typology of the statements, and the information they state. For instance, the three statements in Fig. 6 ascribe different birth dates to Dante.

2.5 Comparison and Further Remarks

We now compare the approaches reported in the previous sections.

A first dimension for the comparison concerns the nature of observations. Both **CRMinf** and **HiCO** conceive them as temporal entities, i.e., activities carried out by agents. An important difference between the ontologies is that HiCO primarily addresses the representation of meta-data of what it calls *interpretation acts*, whereas CRMinf covers both activities of *argumentation* and the resulting *beliefs*. **ICON** sits somehow in-between CRMinf and HiCO in that instances of *Recognition* are understood as events conveying the information they generate. The resulting model remains ambiguous from this perspective because one and the same class, *Recognition*, is both an HiCO's *Interpretation Act*, i.e., something that unfolds in time, and a DUL's *Situation*, a sort of state of the world. Differently from these approaches, **Wikidata** does not represent the event of expressing a statement; rather, it focuses on the resulting information, i.e., the statement itself.

In addition to these aspects, both CRMinf and ICON adopt a **mentalistic approach** in that both CRMinf's beliefs and ICON's recognitions are agents' mental states. As said, this view raises some concerns relative to their intersubjective accessibility and persistence in time. Both ontologies have strategies to face these issues: CRMinf relies on belief's *proposition sets* to represent the content of a belief, whereas ICON uses *interpretation description* as sorts of reifications of recognitions. CRMinf's proposal does not solve the issues: a proposition set is the content of a mental entity; hence, the problems of accessibility and persistence still remain. In the case of ICON, the use of descriptions may be well suited at first glance, but they must be formally structured to "mirror" the data in a recognition. HiCO and Wikidata do not seem to face similar concerns.

A second dimension is the relation between observations and the agents expressing them. Each recognition in ICON includes a single agent, which is consistent with the understanding of recognitions

[15] Should one be interested in representing the *fact* that Dante was born at a certain date or place, the corresponding properties are *wdt:P569* and *wdt:P19*, respectively. Hence, differently from *p:P569* and *p:P19*, these are *direct* properties between items, or items and values, without any reference to statements.

as mental entities. Accordingly, representing a scenario where two (or more) agents express the same observation leads in ICON to the modeling of two different instances of *Recognition*. That is, the identity of recognitions is bounded to their creators. In the case of CRMinf, the dependency of beliefs to agents remains ambiguous, because the model allows a group of agents to share the same belief, but it lacks an in-depth modeling of belief sharing. Leaving this concern aside, the documentation says that "[o]nly one E39 Actor may hold a particular instance of I2 Belief" [23, p.5], which similarly to ICON is consistent with the understanding of beliefs as agents' mental entities. Hence, as for ICON, the representation of two agents expressing the same observation leads in CRMinf to the modeling of two argumentation activities leading to two different beliefs. Then, whether the latter can have the same proposition set, or whether proposition sets depend on specific beliefs remains underspecified in the documentation of the model. In the scope of HiCO, each interpretation act is bound to a single agent. A situation where two acts lead to the same observation can be represented with the introduction of an agreeing relation between them. Finally, differently from all these approaches, Wikidata's statements abstract from who have generated them, i.e., the same statement can be stated in multiple sources. As said, they are primarily defined with respect to their type (e.g., statement of birth date vs statement of authorship attribution), and the information they state (Dante's birth date being May 30, 1265 vs being January 1, 1265).

A third dimension concerns the representation of relations between observations, including relations of inference making. As we have previously noted, it is possible in CRMinf to document the reasoning of scholars in terms of the premises they adopt to reach certain conclusions (see Fig. 2). Both ICON and HiCO borrow some relations from CITO. Differently from HiCO, at first glance, ICON does not represent relations of dis-/agreements between observations, which can be useful to provide a general comparison between recognitions. Wikidata seems to lack relations, e.g., of dis-/agreement between statements but it includes elements (e.g., the relation *stated in*) that are used as evidences for statements.

Requirements. From a general perspective, from the analysis of the state of the art, it emerges that ontologies for observational data need at least to take into account the following requirements:

- **Accessibility:** observations as intersubjective and mind-independent public statements;
- **Provenance**: *who* expresses an observation, *when*, in *which* source;
- **Content**: *what* an observation states;
- **Justification and justification logic**: chains of premises to infer conclusions;
- **Evidence**: reference to elements used by scholars to support their claims, e.g., bibliographic sources, empirical measurements, etc.
- **Ranking:** system of values ("weights") to rank observations on the basis of certain criteria (to be established; e.g., community-based, authority-based, etc.);
- **Dynamic** of argumentation: scholars can revise their observations;
- **Elements for comparison**: a scholar can reject the observations of someone else[16] or may even express observations that contradict others. Conflicts between observations do not necessarily need to be solved; in Doerr et al's [18] words, "conflicts may be permanently accepted without resolution."

3 Literary Characters

We discuss in the following sections the representation of literary characters in ontologies used in digital humanities projects (Sect. 3.1), and large knowledge bases (Sect. 3.2), primarily Wikipedia considering its relevance in nowadays research and application for knowledge-graphs development.

[16] In a dynamic scenario, a scholar may even reject their own claims.

3.1 Characters in Ontologies for the Digital Humanities

There is a rich theoretical debate in philosophy and literary studies about the characterization of literary characters, including discussions about what sorts of things they are, criteria concerning their identity, similarity, and so on [30, 31].[17] On the other hand, there have been only few attempts to create computational ontologies for modeling characters in digital models.

Ciotti [14] presents a preliminary investigation on the ontological modeling of characters within a larger effort concerning the use of ontologies for Digital Literary Studies. Following a narratological and semiotic perspective, characters are understood as "representational devices" described through a set of properties, including characters' narrative functions in relation to Greimas' studies [32].

Hastings and Schulz [33] discuss some challenges relative to the representation of literary characters by recalling some remarks in the philosophical literature. For instance, that an ontology should distinguish between the "internal" and "external" views on characters, i.e., from the perspective of the narrative where it appears (e.g., Holmes is a detective living in England according to Doyle's texts), and outside from it (e.g., Holmes as a fictional character created by Doyle), respectively.

Damiano et al. [34] present an ontological analysis and model for the representation of *dramas*, i.e., "actions played live by characters". Instead of characters, on the basis of research work in both aesthetics and AI, they talk of *agents* stressing the representation of the *actions* that agents undertake or undergo. In this perspective, the description of agents is inspired by the BDI model, i.e., they are represented as having goals and executing plans to achieve them.

Pannach et al. [35] present the ProppOntology, a Semantic Web ontology for representing data of folktales based on the work of the Russian folklorist Vladimir Propp, in particular, his study on the *Morphology of the Folktale*. Recall that Propp, among other things, individuated 31 narrative functions recurring in folktales divided into five macro-categories (preparation, complication, functions of the donor, struggle, and dénouement). For each function, he also provided information of the characters it applies to. For instance, the *Wedding* function applies only when the *hero* character marries the *princess* or a character that fulfills the narrative role of a *princess*. Following Propp, if a wedding between two characters different from the couple hero/princess takes place, the function does not apply. The ProppOntology has been used to analyze folktales of different cultures, for instance, to understand what kinds of functions or characters – following Propp's categories – are found in the analyzed corpora. From this perspective, according to the authors, the proposed system "allows the study of the Proppian morphology interculturally and language-independently which might lead to new findings in folktale research."

Zöllner-Weber [36], following studies by Jannidis [37], among others, addresses the representation of characters from an interpretational perspective; this because information about characters can be acquired only through texts' interpretations. In this view, characters are interpreters' mental representations, described by Zöllner-Weber through different sorts of properties.

3.2 Characters in Knowledge Bases

Characters in Wikidata. Wikidata includes reference to several literary characters like Doyle's Sherlock Holmes (Q4653), Tolstoj' Anna Karenina (Q4066531), Kafka's Gregor Samsa (Q3566632), just to make some examples. For each character, data is represented following the core principle of the data model (see Section 2.4), hence through various sorts of statements. Figure 7 shows an excerpt for the representation of Gregor Samsa in Wikipedia.

Literary character (Q3658341) is the main category used for representing literary characters. A literary character is a "fictional character appearing in written works". This category is ultimately subsumed by *non-existent entity* (Q64728693), i.e., an "entity which does not exist, and has never existed." Characters' representations cover different types of data, e.g., reference to their creators, the work where they are present, the work where they firstly appeared, their narrative role, etc, see e.g., Fig. 8. In addition, for some characters, their representation in Wikidata covers reference to derivative characters, as well as entities upon which the character is meant to be inspired. For instance, Doyle's character of Sherlock Holmes (Q4653) is represented as being inspired by the real-world person Joseph Bell (Q648680), and as being the source for Leslie Bricusse's character of Sherlock Holmes (Q63892120) in *Sherlock Holmes: The Musical*.

At first glance, characters' statements in Wikidata cover at least:

[17] In the scope of MITE, the analysis of philosophical theories of literary characters falls within the scope of WP2.

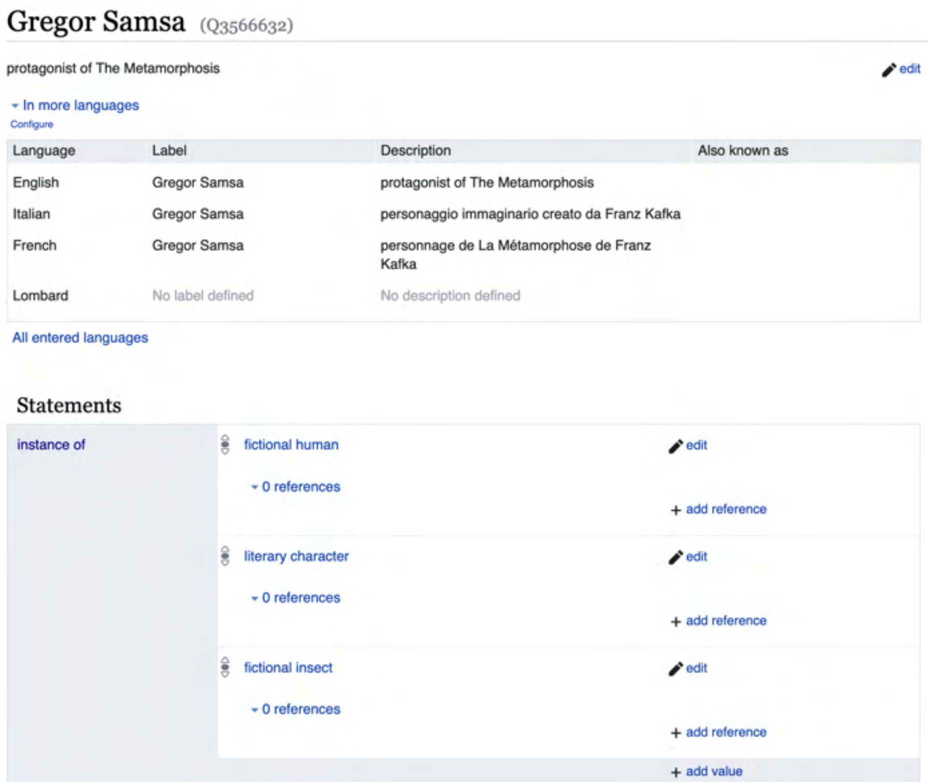

Figure 7: Gregor Samsa in Wikidata (partial view 1)

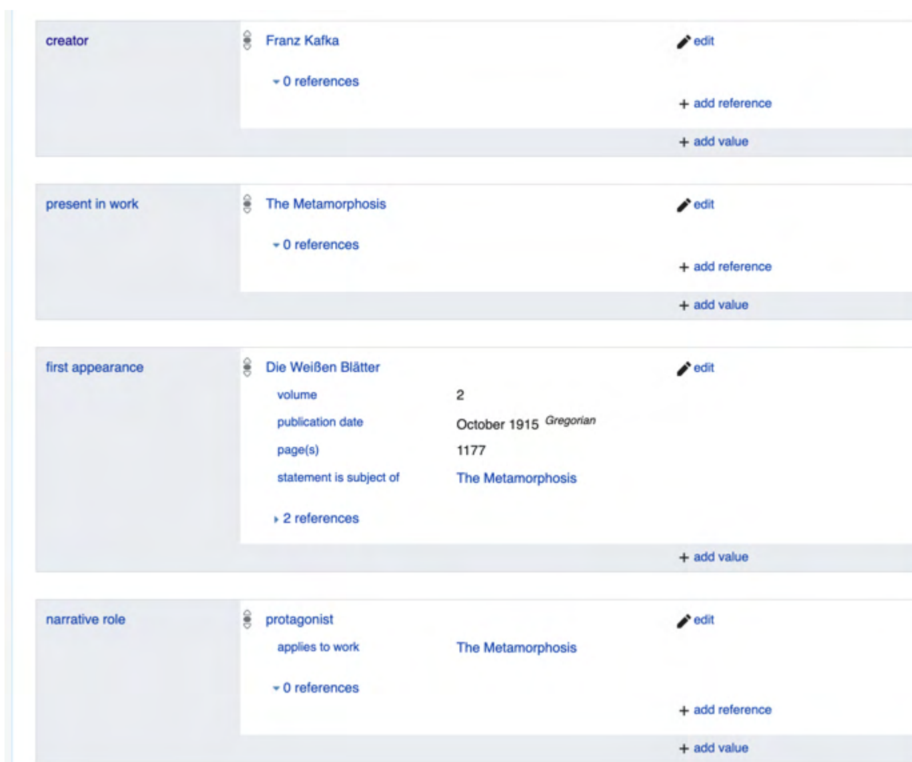

Figure 8: Gregor Samsa in Wikidata (partial view 2)

1. Data within the story perspective:
   - Anagraphic data: name, surname, gender, birth/death date/place, cause of death, residence address, etc.;
   - Family data: sibling, parental, marriage relationships, etc.;
   - Social roles: occupation, etc.;
   - Lifestyle: e.g., being a smoker.
2. Data from the outside of the story perspective:
   - Creation context: creator, first appearance, works where it appears, etc.;
   - Relations with other characters: Bricusse's Holmes based on Doyle's Holmes;
   - Relations with real-world entities: e.g., Holmes being based on Joseph Bell;
   - Declarations of identity: Goethe's Faust (Q63928429) is said to be the same as Faust (Q55000426) in the opera *Faust* by Charles Gounod;[18]
   - Elements of narrative roles: e.g., Holmes as main character of *The Hound of the Baskervilles*, as having Professor Moriarty as enemy;
   - Personality traits: e.g., eccentricity, apathy, etc.

Characters in Dbpedia and YAGO. Knowledge bases like Dbpedia[19] and YAGO[20] cover the representation of literary characters, too. They provide less data than Wikidata but align very much with the sort of things we have just seen above. An important difference with respect to Wikidata is that characters in these knowledge bases are not represented through statements but through standard predication mechanisms.

Remarks. Despite the research effort in the state of the art, further research is in our view needed concerning the way in which literary characters can be represented in ontologies, especially when both philosophical theories and research results in literary studies and criticism are taken into account. For instance, at first glance, it is only Hasting and Schulz [33] that consider philosophical approaches to fictional entities to design their model, which nevertheless remains at a preliminary development phase. Considering the amount of work done by philosophers in analyzing fictional entities, we think that there are several lessons to be learned from their work. At the same time, looking at the literary side, besides the contribution of Zöllner-Weber [36], existing approaches fail to analyze the relations between texts, characters, and interpreters. Even in Zöllner-Weber's contribution, this relation is investigated only from a preliminary standpoint; e.g., the author does not discuss what happens when interpreters come up with different and incompatible descriptions of what is meant to be a single character, not to mention that she embraces a mentalistic approach, opening the way to the sorts of issues discussed above (see Sect. 2.5).

Concerning knowledge bases and in particular the case of Wikidata, because of its bottom-up, community efforts, it is hard to grasp the principles for representing characters in a certain manner. For instance, it is hard to understand how characters are distinguished. Recalling Doyle's Holmes, on the one hand, the same character is represented as being found in multiple works, including Doyle's stories and several film adaptations, whereas on the other hand Bricusse's Holmes is modeled as a different, derivative Holmes.[21] To make an example closer to MITE's case studies, Dante's Beatrice (Q232913) is represented as a person, rather than as a literary character, that is present in Dante's works such as *The Divine Comedy* and *Vita Nuova*. This approach conceptually entails that the real Beatrice underwent all events told by Dante in the *Comedy*. To avoid this, a different strategy could consist in distinguishing between the person of Beatrice and the character created by Dante in his works, aligning to what done for the relationship between Doyle's Holmes and Joseph Bell.

[18] Recall that this is not identity strictly speaking but a statement of identity in Wikidata's sense.
[19] Dbpedia: `https://www.dbpedia.org/`.
[20] YAGO: `https://yago-knowledge.org/`.
[21] Similar considerations can be done for Dbpedia and YAGO.

4 Computational Literary Studies

Computational Literary Studies (CLS) cover a flourishing sub-field of the Digital Humanities where "different types of computer-aided methods are applied to literary texts or [are] especially developed in order to explore texts or test previous hypotheses about them" [38, p. 11] (for a broad survey on CLS' methods, see [39]). Research in CLS typically involves the use of quantitative techniques, such as text mining, machine learning, network analysis, and data visualization, to analyze large corpora of texts for topic modeling, sentiment analysis, stylometric analysis, frequency analysis, and more [15].

Approaches along these lines sometimes face criticism from the standpoint of literary scholarship (see the debate generated by [40], e.g., [41]). This because their results are not perceived as providing relevant or novel insights on texts [42] or because they seem to "reduce the meaning of literary works to their [mere] descriptive/countable features" [38, p. 13]. In a sense, CLS research is perceived as neglecting the nature of literature and literary studies, namely, that literary texts can be interpreted in more than one way, and that literary studies can offer a variety of approaches for interpretation so "that it is hard to imagine how fully formalised and/or automated procedures can do justice to the peculiarities of the discipline and its aesthetic objects" [38, pp. 13-14].

Criticisms of these sorts have brought to the emergence of other approaches within CLS, including what is sometimes called Digital Hermeneutics. "This field is focusing on the question of how digital or computational methods can assist in interpreting literary texts, or in other words: in fostering insights into literary texts that do justice to their complexity and aesthetic quality" [38, pp.14]. In particular, "[o]ne central insight in the context of Digital Hermeneutics lies in the fact that the interpretation of a literary text can hardly ever yield definitive and unanimous results — not only because literary texts are often deliberately ambiguous but also because of the theoretical and methodical plurality of Literary Studies, where different theoretical contexts define different aims and quality criteria for interpretations" [38, p.14]. In a nutshell, if we understand it correctly, Digital Hermeneutics does not seem to necessarily differ from CLS with respect to the techniques; it could be rather understood as a methodological approach to orient CLS research towards analyses that are aware of the intricacies and nuances of literary studies.

Providing a throughout analysis of the state of the art in CLS is out of our scopes here (see [15, 39] for recent surveys on CLS techniques). We limit to report on some projects where the use of ontologies is of some relevance.

MiMoText. Research and application work in the MiMoText project[22] [43] has led to the formal modeling of some aspects relative to literary novels in the French literature of the second half of the 18th century along with scholarly statements concerning the novels. One of the purposes of the project has been the development of a knowledge graph for data relevant to literary historiography. The data has been then analyzed through various approaches, e.g., to calculate stylometric-based similarity between the novels in the corpus.

The knowledge graph has been designed by following the principles of Wikidata; hence for each item (text, in this case), the graph includes various sorts of statements.[23] In particular, for each novel, the graph represents metadata such as authorship, title, publication date and place, etc, the topics it addresses (e.g., *Candido* by Voltaire is about love, migration, disaster, etc), narrative forms (e.g., heterodiegetic, autodiegetic), and reference to the scholarly texts (included in the corpus) referring to it, among other information.[24] In this manner, the knowledge graph links primary sources, scholarly literature about them, providing at the same time bibliographic data.

heureCLÉA. In the context of heureCLÉA,[25] Gius et al. [8] have provided a methodology and system to support the interpretation of literary texts via computational means.[26]

A core idea is to provide an analytic and systematic support for interpretation grounded on the *plural nature* of interpretation. The assumption behind the work of Gius et al. is that literary texts

[22] MiMoText: https://mimotext.uni-trier.de/english/.
[23] https://data.mimotext.uni-trier.de/wiki/Main_Page
[24] To make an example, see the data for Voltaire's *Candido*: https://data.mimotext.uni-trier.de/wiki/Item:Q1022.
[25] https://www.heureclea.de/index.html.
[26] The computational system developed in the project is called CATMA – Computer Assisted Textual Markup and Analysis – and is available at: https://catma.de/.

are “ambiguous or polyvalent [in meaning] by nature” (p.234); hence, different sorts of even conflicting interpretations for the same text are not only possible but can coexist.

From a methodological standpoint, the idea that annotations on texts can be used to express interpretations plays a fundamental role. The authors borrow the notion of *hermeneutic annotation* from Piez [44] meaning an annotation that “is not limited to describing aspects or features of a text that can be formally defined and objectively verified. Instead, it is devoted to recording a scholar’s or analyst’s observations and conjectures in an open-ended way. [..] Rather than being devoted primarily to supporting data interchange and reuse [...] hermeneutic markup is focused on the presentation and explication of the interpretation it expresses” (Piez, quoted in [45]).

In practice, the proposal consists in having multiple interpreters expressing their interpretations by annotating the text at stake through a shared vocabulary. In particular, interpreters first individually annotate a text, then compare their results, especially when “interpretive disagreement”, in Gius et al.’s words, emerges.

The various steps of the methodology are depicted in Fig. 9. It is important to stress that a core idea of the methodology is that interpreters must share the same concepts when annotating texts, i.e., they must use the adopted vocabulary in the same way. If disagreements among the interpreters arise because of ambiguities/inaccuracies in the vocabulary, these have to be discussed among the interpreters and solved. Genuine cases of disagreements arise when it is the text itself to entail multiple conflicting interpretations.

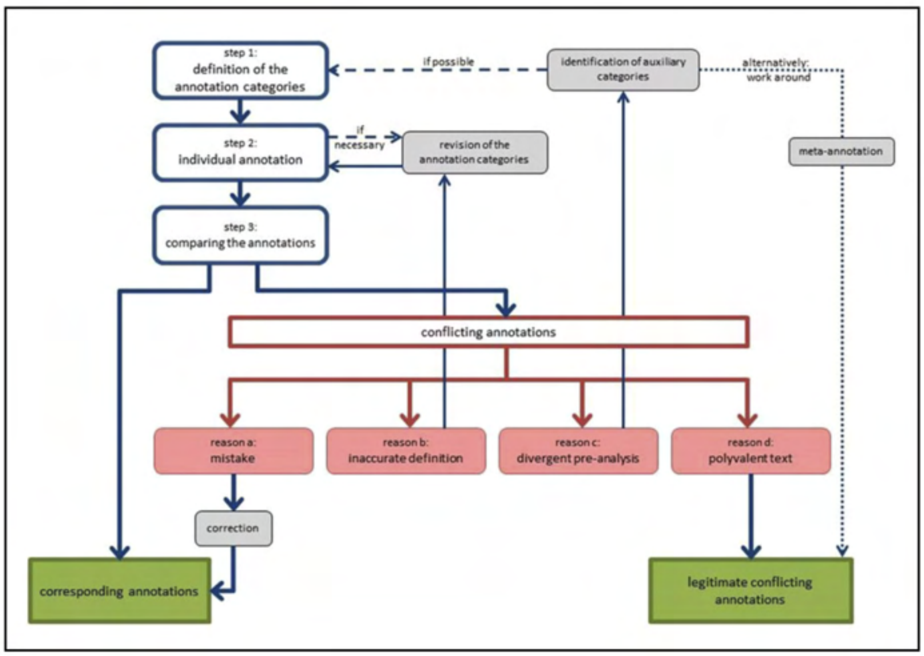

Figure 9: Methodology for annotations, from Gius et al. [8]

In addition to the definition of a methodology and vocabulary for annotation, Gius et al. have also set machine-learning algorithms to automatically annotate the texts (see [45]).

CLS Infra. The Computational Literary Studies Infrastructure[27] (H2020 EU project) is likely one of the most important initiatives in the CLS community at the current state of research and development. The purpose is “to build the shared and sustainable infrastructure needed to undertake literary studies in the digital age” (from the project’s webpage). The project builds around the concept of *programmable corpora*, which is essentially the idea of treating a literary corpus as a computational artefact that can be manipulated through computational means [46]. DraCor, standing for *Drama Corpora Platform*, is the prototype application developed for this goal [47].[28] “At its core there are

[27] CLS Infra: https://clsinfra.io/

[28] DraCor: https://dracor.org/. See https://www.dracor.org/doc/research for a list of research results employing DraCor and CLS Infra techniques.

homogenized corpora in several (European) languages. These corpora are curated in GitHub repositories and stored in an XML database as a central data store which provides a RESTful API that powers a front end and can be used individually to retrieve the raw TEI-XML data, metadata, and derived data in several formats. Attached microservices offer additional functionalities, e.g., a Metrics Service is used to calculate network metrics based on the play data and a Triple Store, which holds representations of the plays as Linked Data and provides a SPARQL endpoint" [47, p.21].

As the quote above suggests, DraCor includes a RDF serialization of data concerning the corpora that is accessible through a SPARQL endpoint.[29] It also builds connections with Wikidata by reusing its identifiers.[30] The ontology used for the serialization of the data in RDF is a simple model containing few classes, object, and data properties, with no axioms besides some domain/range declarations.[31] Current research work aims at improving the ontology possibly connecting it with existing works in the state of the art, including CIDOC-CRM.

Remarks. Differently from "standard" CLS approaches, we do not adopt in **MITE** quantitative computational approaches to generate data from (primary or secondary) texts to support their interpretation, as it is done, e.g., for stylometric analysis. From this perspective, our approach is close to what is done in projects like MimoText, and heureCLÉA, where ontologies are used to drive the extraction of data in a controlled manner, ensuring transparency and trustworthiness. We also share with Gius et al. [8] significant theoretical and methodological principles, above all that a literary text does not have a single prescribed meaning, hence, it can lead to multiple and possibly conflicting interpretations. In this sense, the role of domain experts is fundamental in MITE because they support the shaping of the vocabularies for the documentation of observational data.

From the perspective of ontological modeling, what seems to be a core distinction with MimoText, heureCLÉA, and CLS Infra is that the framework we aim at developing will be explicitly designed to address the representation of data as *observational data* satisfying (some of) the requirements identified in Sect. 2.5. Also, differently from text annotations based on semi-structured terminologies, the use of a logic-based framework can allow us to formally compare observational data in a precise manner as well as to automatically reason over the data. In addition, MITE's framework will cover the modeling of arguments and their structure, which seems a topic scarcely explored at the current state in CLS.

These differences apart, as said in the Introduction, we understand MITE and CLS approaches as being complementary, and look forward to research bringing together multiple methodologies, objectives, and techniques.

5 Conclusions

The analysis of research work presented in the report is the first step in MITE towards the development of an exploratory conceptual framework to document scholarly interpretations with a focus on literary characters. In particular, the report allows us to frame our goals with respect to similar efforts.

Concerning the representation of observations in the Digital Humanities, the approach we attempt to design will address three core dimensions of literary interpretation. First, the *textual dimension*, namely, the focus on the analysis of texts and what they mean for interpreters. Second, the *rhetorical dimension*, namely, reference to the arguments put forward by scholars and critics to ground their interpretations; third, the *intersubjective dimension*, namely, that interpretations must be accessible independently from who state them. Although these three dimensions are more or less present in the state of the art, further research work is needed to address the limits of current approaches, as highlighted in the report.

An important distinction with respect to research in CLS is that MITE is not to be conceived as a quantitative computational approach for the production of data from texts. Instead, MITE is an approach shaped by methodologies and vocabularies developed with domain experts to make certain aspects of scholarly and critical discourse on literary texts accessible through information systems. This is aimed at facilitating their documentation, retrieval, comparison, and analysis. Notably, current

[29] SPARQL Endpoint: https://www.dracor.org/sparql.

[30] See the documentation on DraCor and Linked Open Data: https://github.com/dracor-org/dracor-notebooks/blob/lod-intro/lod-intro/lod-intro.ipynb.

[31] The DraCor ontology is available at: https://dracor.org/ontology.

cultural heritage web platforms primarily manage the metadata of cultural heritage, including literature, while the representation of scholarly debates on literary texts is overlooked.[32] Consequently, critics and scholars cannot use these platforms to search for specific claims about a text or to compare different claims about the same text. In contrast, similar platforms exist in other fields, such as musicology.[33] We believe that advancing MITE's objectives will potentially underpin the conceptual models of future platforms for Digital Literary Studies, thereby making interpretations of literary texts available for further research.

For the sake of clarity, three important considerations are needed to outline our goals. First, by adopting formal modeling techniques, we must find reasonable compromises with respect to interpretation practices in literary studies. Inevitably, some dimensions of literary interpretations will be only partially represented or simplified compared to how scholars and critics express them. A significant challenge is to achieve a representation granularity that is meaningful from a scholarly standpoint without trivializing experts' work. From this perspective, we do not aim to "translate" scholars' statements into formal terms, as this would be a worthless effort given the limited expressivity of formal languages compared to the nuances of natural language used in scholarly contexts. In a sense, our framework can be understood as providing a "virtual map" (a knowledge graph) of scholarly debates such that, for each statement of interest, a scholar can trace who claims it, in which source, on the basis of which arguments, evidences, etc.

Second, as it follows from what just said, our approach does not intend to replace "close" interpretation. Instead, we aim to support this effort by providing conceptual tools that emerge from the intersection of research in literary studies, philosophy, logic, and computer science.

Third, documenting the interpretation found in a text is a *selective* task requiring decisions about what to document and what to leave aside, at what level of abstraction, etc. This can easily lead to mistakes or the documentation of claims that might not be asserted in the interpreted text. In our view, there is no straightforward solution to this situation. Methodologically, similar to what is done for textual annotations in [8], the task of documenting interpretations should be carried out systematically and collaboratively, with multiple agents discussing their results to avoid the proliferation of mistakes. In the next development phases of the project, we plan to engage with domain experts in literature and also with projects in other fields that have faced similar issues to learn from their experiences.

Acknowledgments: We wish to thank Marco de Cristofaro, Claudio Masolo, Alessandro Mosca, and Antonio Sotgiu for their valuable comments on previous versions of this document. The authors are solely responsible for any remaining misunderstandings or mistakes.

To quote this work: Sanfilippo, E.M., Ferrario, R., "D3.1 - Observations Modeling: State of the Art." Project report: MITE – Make it explicit: Documenting interpretations of literary fictions with conceptual formal, funded by European Union – Next Generation EU, May 2024.

References

[1] T. Andrews, "Modelling historical data in the releven project," *Programming and Data Infrastructure in Digital Humanities*, p. 16, 2023.

[2] G. Barabucci, F. Tomasi, and F. Vitali, "Supporting complexity and conjectures in cultural heritage descriptions," 2021.

[3] S. Cristofaro, E. M. Sanfilippo, P. Sichera, and D. Spampinato, "Towards the representation of claims in ontologies for the digital humanities.," in *SWODCH*, 2021.

[4] M. Daquino and F. Tomasi, "Historical context ontology (hico): a conceptual model for describing context information of cultural heritage objects," in *Metadata and Semantics Research: 9th Research Conference, MTSR 2015, Manchester, UK, September 9-11, 2015, Proceedings 9*, pp. 424–436, Springer, 2015.

[5] I. Willis, *Reception*. Routledge, 2017.

[32] For more on this, see the slides from the first MITE research seminar: https://www.loa.istc.cnr.it/mite/index.php/events/mite-research-seminars/.

[33] An example is CRIM: https://crimproject.org/.

[6] C. Masolo, A. Botti Benevides, and D. Porello, "The interplay between models and observations," *Applied Ontology*, vol. 13, no. 1, pp. 41–71, 2018.

[7] M. Daquino, V. Pasqual, and F. Tomasi, "Knowledge representation of digital hermeneutics of archival and literary sources," *JLIS. it*, vol. 11, no. 3, pp. 59–76, 2020.

[8] E. Gius and J. Jacke, "The hermeneutic profit of annotation: On preventing and fostering disagreement in literary analysis," *International Journal of Humanities and Arts Computing*, vol. 11, no. 2, pp. 233–254, 2017.

[9] M. Piotrowski and M. Neuwirth, "Prospects for computational hermeneutics," in *Atti del IX Convegno Annuale AIUCD*, 2020.

[10] M. Caracciolo, "Chapter gli studi letterari cognitivi e lo statuto dell'interpretazione: un tentativo di mappatura teorica," in *La narrazione come incontro*, pp. 37–58, Firenze University Press, 2022.

[11] N. Carroll, "Interpretation," in *The Routledge Companion to Philosophy of Literature*, pp. 302–312, Routledge, 2015.

[12] E. M. Sanfilippo, A. Sotgiu, G. Tomazzoli, C. Masolo, D. Porello, and R. Ferrario, "Ontological modeling of scholarly statements: A case study in literary criticism," in *Formal Ontology in Information Systems (FOIS 2023)* (N. Aussenac-Gilles, T. Hahmann, A. Galton, and M. M. Hedblom, eds.), vol. 377 of *Frontiers in Artificial Intelligence and Applications*, pp. 349–363, IOS Press, 2023.

[13] S. Pan, L. Luo, Y. Wang, C. Chen, J. Wang, and X. Wu, "Unifying large language models and knowledge graphs: A roadmap," *IEEE Transactions on Knowledge and Data Engineering*, 2024.

[14] F. Ciotti, "Toward a formal ontology for narrative," *MATLIT: Materialidades da Literatura*, vol. 4, no. 1, pp. 29–44, 2016.

[15] H. O. Hatzel, H. Stiemer, C. Biemann, and E. Gius, "Machine learning in computational literary studies," *it-Information Technology*, vol. 65, no. 4-5, pp. 200–217, 2023.

[16] A. Pichler and N. Reiter, "From concepts to texts and back: Operationalization as a core activity of digital humanities," *Journal of Cultural Analytics*, vol. 7, no. 4, 2022.

[17] N. Benn, S. B. Shum, J. Domingue, and C. Mancini, "Ontological foundations for scholarly debate mapping technology," *Frontiers in Artificial Intelligence and Applications*, vol. 172, p. 61, 2008.

[18] M. Doerr, A. Kritsotaki, and K. Boutsika, "Factual argumentation—a core model for assertions making," *Journal on Computing and Cultural Heritage (JOCCH)*, vol. 3, no. 3, pp. 1–34, 2011.

[19] S. B. Shum, J. Domingue, and E. Motta, "Scholarly discourse as computable structure," in *International Workshop on Structural Computing*, pp. 120–128, Springer, 2000.

[20] B. Sartini, S. Baroncini, M. van Erp, F. Tomasi, and A. Gangemi, "Icon: An ontology for comprehensive artistic interpretations," *ACM Journal on Computing and Cultural Heritage*, vol. 16, no. 3, pp. 1–38, 2023.

[21] E. M. Sanfilippo and R. Freedman, "Ontology for analytic claims in music," in *European Conference on Advances in Databases and Information Systems*, pp. 559–571, Springer, 2022.

[22] E. Bruno, V. Pasqual, and F. Tomasi, "Italo calvino's 'destini incrociati'. an experiment of semantic narrative modelling and visualisation," *Umanistica Digitale*, no. 17, pp. 47–69, 2024.

[23] M. Doerr, C.-E. Ore, and P. Fafalios, "Definition of the crminf. an extension of cidoc-crm to support argumentation," tech. rep., CIDOC CRM-SIG, 10 2023.

[24] G. Bruseker, N. Carboni, and A. Guillem, "Cultural heritage data management: the role of formal ontology and cidoc crm," *Heritage and archaeology in the digital age: acquisition, curation, and dissemination of spatial cultural heritage data*, pp. 93–131, 2017.

[25] M. Doerr, A. Kritsotaki, Y. Rousakis, G. Hiebel, and M. Theodoridou, "Crmsci: The scientific observation model," tech. rep., CIDOC CRM-SIG, 10 2023.

[26] C. Bekiari, G. Bruseker, E. Canning, M. Doerr, P. Michon, C.-E. Ore, S. Stead, and A. Velios, "Conceptual reference model (cidoc-crm), version 7.2.4," tech. rep., CIDOC CRM-SIG, 10 2023.

[27] A. Gangemi and P. Mika, "Understanding the semantic web through descriptions and situations," in *OTM Confederated International Conferences" On the Move to Meaningful Internet Systems"*, pp. 689–706, Springer, 2003.

[28] V. A. Carriero, A. Gangemi, M. L. Mancinelli, A. G. Nuzzolese, V. Presutti, and C. Veninata, "Pattern-based design applied to cultural heritage knowledge graphs," *Semantic Web*, vol. 12, no. 2, pp. 313–357, 2021.

[29] F. Erxleben, M. Günther, M. Krötzsch, J. Mendez, and D. Vrandečić, "Introducing wikidata to the linked data web," in *The Semantic Web–ISWC 2014: 13th International Semantic Web Conference, Riva del Garda, Italy, October 19-23, 2014. Proceedings, Part I 13*, pp. 50–65, Springer, 2014.

[30] J. Eder, F. Jannidis, and R. Schneider, *Characters in fictional worlds: Understanding imaginary beings in literature, film, and other media.* de Gruyter, 2010.

[31] F. Kroon and A. Voltolini, "Fictional Entities," in *The Stanford Encyclopedia of Philosophy* (E. N. Zalta and U. Nodelman, eds.), Metaphysics Research Lab, Stanford University, Fall 2023 ed., 2023.

[32] A. J. Greimas, *Sémantique structurale: recherche de méthode.* Puf, 2015.

[33] J. Hastings and S. Schulz, "Representing literary characters and their attributes in an ontology.," in *JOWO*, 2019.

[34] R. Damiano, V. Lombardo, and A. Pizzo, "The ontology of drama," *Applied Ontology*, vol. 14, no. 1, pp. 79–118, 2019.

[35] F. Pannach, C. Sporleder, W. May, A. Krishnan, and A. Sewchurran, "Of lions and yakshis," *Semantic Web*, vol. 12, no. 2, pp. 219–239, 2021.

[36] A. Zollner-Weber, *Noctua literaria - A Computer-Aided Approach for the Formal Description of Literary Characters Using an Ontology.* PhD thesis, Universitaet Bielefeld, 2008.

[37] F. Jannidis, *Figur und Person: Beitrag zu einer historischen Narratologie.* Walter de Gruyter, 2004.

[38] M. Flüh, J. Horstmann, J. Jacke, and M. Schumacher, "Toward undogmatic reading," 2021.

[39] C. Schöch, J. Dudar, and E. Fileva, eds., *Survey of Methods in Computational Literary Studies (= D 3.2: Series of Five Short Survey Papers on Methodological Issues).* CLS INFRA.

[40] N. Z. Da, "The computational case against computational literary studies," *Critical inquiry*, vol. 45, no. 3, pp. 601–639, 2019.

[41] F. Jannidis, "On the perceived complexity of literature. a response to nan z. da," *Journal of Cultural Analytics*, vol. 5, no. 1, 2020.

[42] F. Ciotti, "Distant reading in literary studies: a methodology in quest of theory," *Testo e Senso*, no. 23, pp. 195–213, 2021.

[43] C. Schöch, M. Hinzmann, J. Röttgermann, K. Dietz, and A. Klee, "Smart modelling for literary history," *International Journal of Humanities and Arts Computing*, vol. 16, no. 1, pp. 78–93, 2022.

[44] W. Piez, "Towards hermeneutic markup: An architectural outline.," in *DH*, pp. 202–205, 2010.

[45] T. Bögel[1], M. Gertz[1], E. Gius, J. Jacke, J. C. Meister, M. Petris, and J. Strötgen[1], "Collaborative text annotation meets machine learning: heurecléa, a digital heuristics of narrative.," *The Digital Humanities Commons*, 2015.

[46] F. Fischer, I. Börner, M. Göbel, A. Hechtl, C. Kittel, C. Milling, and P. Trilcke, “Programmable corpora: introducing DraCor, an infrastructure for the research on european drama,” *Digital Humanities*, vol. 2019, p. 5, 2019.

[47] I. Börner and P. Trilcke, “D7.1 on programmable corpora. Report and prototype (DraCor),” project report, Computational Literary Studies Infrastructure, 2023.