ALMA MATER STUDIORUM
UNIVERSITÀ DI BOLOGNA

DOTTORATO DI RICERCA IN

DOTTORATO DI RICERCA INCULTURE LETTERARIE E FILOLOGICHE

Ciclo 37

Settore Concorsuale: 11/A4 - SCIENZE DEL LIBRO E DEL DOCUMENTO E SCIENZE STORICO RELIGIOSE

Settore Scientifico Disciplinare: M-STO/08 - ARCHIVISTICA, BIBLIOGRAFIA E BIBLIOTECONOMIA

THE CRITICAL INQUIRY IN HUMANITIES KNOWLEDGE GRAPHS: CHALLENGES, METHODS AND INNOVATIONS

Presentata da: Valentina Pasqual

Coordinatore Dottorato

Prof.essa Valentina Garulli

Supervisore

Prof.essa Francesca Tomasi

Co-supervisore

Prof. Fabio Vitali

Esame finale anno 2025

I, , confirm that the work presented in this thesis is my own. Where information has been derived from other sources, I confirm that this has been indicated in the work.

Per la par condicio
e per un'infinità di altre ragioni,
a Bruno

Abstract

This dissertation explores the challenges of representing humanities critical inquiries in Knowledge Graphs (KGs). Current KGs and online catalogues seem to flatten such an interpretative approach in the humanities. Additionally, they typically omit superseded theories rather than representing them with a weaker logical status (WLS). After defining the domain and its theoretical framework, the dissertation surveys existing methods for introducing claims in RDF (reification methods), examines how reification can convey a weaker logical status (expressing without asserting), and explores how existing ontologies formalise interpretations. Additionally, CH items in Wikidata have been systematically surveyed to determine annotators' habits and requirements, revealing poor usage of WLS claims and difficulty retrieving claims related to critical inquiry. Considering the strengths and weaknesses of existing reification methods, a new method, called Conjectures, which extends the RDF 1.1 syntax of Named Graphs, is proposed as a potential solution to include critical inquiries in CH KGs. The effectiveness and efficiency of Conjectures is assessed, demonstrating that it is competitive with existing reification methods. The representation is applied to a case study focusing on scholarly investigations into document authenticity. Specifically, the study examines the "Index of Medieval Documents Concerning the Upper Austrian Region that are Damaged, Tampered with, or Altered" by Siegfried Haider, which indexes the critical assessment of 153 historical documents. A data model and KG are produced and evaluated through a web application that integrates the claims into a prototypical online catalogue. The dissertation concludes by addressing the main research questions of the work and proposing some further developments on the work.

Table of Contents

List of Abbreviations

Concept	**Abbreviation**
Cultural Heritage	CH
Knowledge Graph(s)	KG(s)
Weaker Logical Status	*WLS*
Resource Description Framework	RDF
Expressing Without Asserting	EWA

List of Figures

List of Tables

List of Code Listings

Introduction

Linked Open Data (LOD) is nowadays considered the standard for encoding and sharing Cultural Heritage (CH) data through the World Wide Web. By adhering to LOD principles, CH institutions can publish their catalogues in a structured, machine-readable format that promotes accessibility, interoperability and data exchange across diverse platforms and applications. Online catalogues publish cultural objects descriptive metadata in a subject-predicate-object structure using the Resource Description Framework (RDF) (e.g. author, date of creation, dimensions).

However, studying and interpreting cultural artefacts, historical events, and literary texts is often marked by evolving understandings and shifting perspectives. New discoveries, technological advancements, and the continuous re-evaluation of existing knowledge drive these changes.

This Dissertation explores the representation of the evolution of critical inquiry reshaping humanities data. This representation aims to supersede the flat, single-perspective representation in cultural object metadata within LOD catalogues. In particular, this work focuses on how historical evolution can reshape scholarly narratives and explores the benefits of including this information within a LOD environment.

Chapter 1 provides the background for this work. Section 1.1 establishes the theoretical framework, introducing the constructivist perspective on representing critical inquiry. This perspective positions humanities research within interpretive methodologies due to its irreducible historiographical uncertainty [86], contrasting it with the positivist, data-centric approaches commonly found in STEM fields.

Section 1.2 presents two guiding examples to illustrate the nature of scholarly interpretations as they shift as new data or methods become available: Vermeer's "Girl Reading a Letter at an Open Window" as an instance of evolving knowledge in art history, and the "Donation of Constantine" as an instance of scholarly challenge on documents' authenticity.

Section 1.3 identifies a gap in how CH collections handle the representation of scholarly inquiries. By examining the cataloguing practices for the guiding examples, LOD catalogues and related current metadata standards fail to represent such information. Reticence, flattening, coercion, and dumping in critical inquiry representation are often the rules than the exceptions.

The research statement (Section 1.4) advocates for a more nuanced approach to digital records, proposing Knowledge Graphs (KGs) that distinguish and integrate scholarly opinions, explicitly representing their logical status concerning their position in the scholarly debate (e.g., disputed, undisputed, settled). This representation aims to improve the precision and accessibility of data for future consultation and reuse. This approach introduces the concept of *Weaker Logical Status* (WLS), recognising that such opinions are not strictly true or false but may, for example, be widely accepted yet not universally agreed upon or previously considered true but later discarded. The section provides essential definitions, hypotheses, assumptions, and research questions.

Chapter 2 introduces the current state of RDF representation of critical inquiry. In particular, Section 2.1 surveys KGs and data models within the CH domain, specifically examining how these handle critical inquiry and its potentially conflicting conclusions.

Section 2.2, introduces reification methods to articulate *statements about statements* and, therefore, to express claims' provenance and contextual information, primarily distinguishing reification methods between ontology-dependent and ontology-independent solutions. Additionally, the concept of *Expressing Without Asserting* (EWA) is presented as the capacity of certain ontology-independent reification methods to represent claims or statements without implying their absolute truth or acceptance. EWA is defined as a possible solution for a common approach to document critical inquiry and their uncertain or concurring conclusions.

Findings are discussed in Section 2.3, in particular noting that named graphs are used in CH KGs to represent provenance [26, 66] and encode contrasting opinions when n-ary relations fall short [36]. However, named graphs ended up in RDF 1.1 but without clear-cut semantics about the logical status of their content and requiring additional methods to achieve EWA with several impractical drawbacks (e.g., N3 rules not backwards compatible with RDF 1.1 syntax).

Chapter 3 analyses over 3 million artworks in Wikidata, exploring several patterns for representing questionable statements. This includes different types of WLS claims, such as uncertain information, competing hypotheses, and temporally evolving knowledge. The chapter examines how these representations are utilised in actual collections across both the humanities and, for comparison, the hard sciences. The adoption of Wikidata has become widespread within the CH community.

Section 3.1 presents the approaches provided in Wikidata to encode WLS claims, in particular on three main families of approaches to express the WLS claims, namely *ranked statements*, *unknown objects*, and *qualified statements*.

Section 3.2 describes the data acquisition process. Two large sets of topics were accessed and downloaded from Wikidata: one belonging to CH domain (including visual works of art such as paintings and statues, text documents, and audiovisual entities) and another from astronomy (celestial bodies such as stars and galaxies). Both sets use multiple fuzzy assertions and hypotheses and, therefore, require assertions with weaker status (e.g., attribution uncertainties or physical locations moving over time for paintings vs. spectral class or radial velocity for stars).

Section 3.3 presents the Wikidata sample dataset analysis. The findings of this survey show that the amount of WLS statements in Wikidata seems suspiciously low, as only 0,4% of visual artworks report attribution debates, a fairly low figure compared to, e.g., a more reasonable 8,5% coming from the RKD images collection[1], a difference that could be attributed to the difficulty and ambiguities in the procedures to report such complex information. The poor usage of these methods is registered, besides an over-generalisation in the terminology used to capture contextual information about such opinions. A broad spectrum of terms is required to represent diverse and specific contexts accurately, evidenced by the 90 terms related to reasons for deprecation presented in Figure 3.2. However, these terms are often subsumed under overly general properties, such as sourcing circumstances (`P1480`) and the nature of statements (`P5102`), which fail to adequately address methodological aspects or the specifics of the evidence collected by annotators.

In Section 3.4, a way to simplify, streamline, and homogenise such complexity is proposed, hoping to increase the abundance, richness, and correctness of the representation of such phenomena in Wikidata.

Finally, in Section 3.5, the findings are summarised, and the conclusions about the work are discussed.

Chapter 4 examines and discusses the need and usefulness of EWA arbitrary or questionable claims in RDF. EWA is analysed as a possible solution to represent the evolution of critical inquiry to represent and retrieve statements whose logical status cannot or does not want to assert.

In Section 4.1, existing ontology-independent methods to express without asserting are listed and compared by using the guiding examples presented in Section 1.2. A short list includes reification [56], n-ary relationships [75], Wiki-

[1] `https://rkd.nl/en/explore/images`

data (employing special Statements), which can receive any kind of contextual additions and also a ranking to assert priority between competing Statements), named graphs [51], RDF-star [54] (expressing triples that are not actually asserted) and N3 [11]. While all these approaches can be used for the purpose of EWA, they can do so partially, with additional complexities and/or side effects. In particular, the non-asserted status of triples seems more like a side effect than a true design issue from an effectiveness point of view (Section 4.2.6).

Section 4.2 describes a possible solution, called *Conjectures*, and its effectiveness in achieving EWA is discussed. Conjectures is an extension of RDF 1.1 named graphs syntax, which come in two non-alternative solutions, one compliant with RDF 1.1 notation (*weak form*) and one extending RDF 1.1 syntax (*strong form*). In particular, Conjectures make use of three main types of graphs: plain named graphs (named graphs original syntax), conjectural graphs (non-asserted graphs which are deemed questionable or debated), and settled conjectures, which record both the dispute, as well as its subsequent resolution, intentionally distinct from a mere re-assertion of disputed claims, as it neither acknowledges nor mentions the existence of the dispute.

Section 4.3 presents and discusses the efficiency evaluation of EWA surveyed methods, including the Conjectures solution. In particular, Section 4.3.1 present the data acquisition, scaling and conversion of the dataset used to perform the experiment. The test surveyed several reification methods such as named graphs, RDF-star, Singleton properties, Wikidata statements, and Conjectures (weak and strong form) on four major metrics: number of triples in endpoint, loading time, dataset weight in the triplestore, and query execution time. The results, presented in Section 4.3 and discussed in Section 4.3.3. In particular, they demonstrate that, among the most efficient methods, such as RDF-star and Wikidata statements, the strong form of Conjectures exhibits notable performance gains, particularly in retrieving claims about questioned or debated knowledge.

Chapter 5 moves beyond the general framework adopted in the previous chapters and focuses on a specific case study. The chapter introduces EWA to represent scholarly opinions challenging historical documents' authenticity, considering how knowledge surrounding these documents may evolve over the centuries. The sections of this chapter outline the steps involved by the mythLOD methodology [81] adopted to address the formal representation of the case study in RDF. This methodology adopts a bottom-up approach for knowledge production, beginning with a detailed analysis of the existing data and focusing on understanding and structuring data at a granular level before integrating it into a broader, more comprehensive data model.

Section 5.1 addresses the source data analysis performed on the *Index of Medieval Documents Concerning the Upper Austrian Region that are Damaged, Tampered with, or Altered* [52], a catalogue compiled by Siegfried Haider (2022) which reports a set of inquiries over 153 known forgeries collected in the *Urkundenbuch des Landes ob der Enns* aiming to define the representation requirements of the data model.

Sections 5.2 and 5.3 cover the data management by applying a set of iterative steps to manipulate, organise and structure the source data. Specifically, Section 5.2 presents the data model that introduces EWA (in particular, using Conjectures in strong form) to represent questioned opinions, existing domain ontologies to represent scholars' conclusions and related knowledge, and an on-purpose ontology, called SEBI, to represent evidence-based analysis and authenticity assessment. Section 5.3 describes the steps taken to convert the source material into a KG based on the data model.

Section 5.4 addresses the testing of the produced KG by developing a web application named Broast, which stands as a proof-of-concept for the work. Broast has been used to test the representativeness of the data model via SPARQL queries, particularly the inclusion of the EWA approach on the data.

Finally, Chapter 6 presents the overall conclusions of the work, summarises findings concerning the research questions and addresses future works.

Chapter 1

Background

1.1 Theoretical Framework

What drives the scholar is not only the pursuit of what is true against what is false, but, among what is true, to determine what is interesting, what is intriguing, what tells the most compelling story, and what holds the most significant implications for the scholarly work. This implies that the best scholarly stories are not simply the final results of studies but, most importantly, the narration of how those results were reached. Representing the complexity of scholarly activity is essential to narrate these stories, including the points of view and opinions that were worked on, discussed, objected to, found lacking or false, or possibly rescued from obscurity and discredit.

Scholarly inquiry in the humanities heavily relies on interpretation and the uncertainty surrounding their sources' history, which is often the norm rather than the exception. Notably, Chiesa [18] asserts that the "Mona Lisa" is the only painting by Leonardo da Vinci whose attribution has never been seriously questioned.

This inherent uncertainty inevitably influences how humanities data collection is approached, requiring special attention to how information is actively constructed. In this context, Checkland and Holwell [17] distinguish between two approaches to data collection, named data and capta.

Data refers to the mass of facts given and observed. The term comes from the Latin word "dare," meaning "to give." Data is often collected through observation and measurement of external reality (e.g., the material, the dimensions, and the current location of a painting). While individual data points may not be fascinating on their own, their data collection and analysis are necessary. This approach to data collection is often associated with a realist perspective, where data is seen as an objective representation of the world.

This distinction neglects the acceptance that all knowledge is provisional. Provisionality implies that knowledge is not absolute or final and that information is always within a particular context. This perspective has been applied in various contexts such as information science [16], historical sciences [67], and physics [103].

On the other hand, capta refers to the tiny fraction of the available data that scholars actively take, filter, select, and interpret. The term comes from the Latin word "capere," meaning "to take." Capta is often associated with a constructivist perspective, where knowledge is viewed as constructed and subjective (e.g., attributing the same painting to a specific artist or artistic school). Capta involves searching for relevant information, filtering out irrelevant data, selecting the most critical information, and interpreting it based on our situated partial and constitutive knowledge. This process of knowledge construction is often seen as a more humanistic approach to research, as it acknowledges the role of the researcher in shaping the knowledge produced.

Both data and capta are dynamic, evolving with new interpretations, perspectives, and technologies [17]. This constant evolution requires scholars to remain adaptable and open to revising their understandings in light of new evidence or methodologies.

Drucker [39] further expands on this distinction, arguing that data and capta have different ontological implications. Data represents pre-existing facts, while capta represents situated, partial, and constitutive knowledge of a constructed nature involving selections, interpretations, and expressions of opinions and points of view. This distinction has important implications for knowledge production, as it challenges the notion of data as an objective representation of the world and highlights the importance of context and interpretation in shaping our understanding of the world.

Critical discourse in the humanities is characterised by the variability and uncertainty of interpretations, most frequently leading to concurring and incompatible statements about the same concepts or facts due to different viewpoints and sources. Complexities abound: on the one hand, uncertainty spans across a wide range of disciplines (from iconography to art, from philology to palaeography, etc.); on the other hand, it is essential to consider an equally broad range of analyses performed by humanities scholars, from the hermeneutic examination of a painting to the construction of a critical apparatus for a manuscript. [86].

Piotrowski [87] points our that it is essential to recognize that big data and a "digital historical positivism" do not suffice as methods for understanding history. Historical narratives are not simply derived from an aggregation of facts. We must differentiate between two types of uncertainty: historical uncertainty, which pertains to the facts of the past, and historiographical uncertainty, which involves the causal models historians construct. Historical knowledge is inherently fragmentary, as history often involves interpreting sources not intended to convey messages to future historians. Historians know that sources can be misleading, containing inaccuracies, errors, fabrications, and forgeries. Even when convinced of correctly identifying a person, place, or time of an event, they acknowledge that absolute certainty is unattainable. This awareness is reflected in their narratives, which often acknowledge relevant uncertainties.

Piotrowski [87] further clarifies that uncertainty arises not merely from data itself (consider, for example, the data extracted from the digitization of a birth certificate) but from the interpretative links established by scholars on such data (for instance, the identification of a name in the birth certificate with an actual historical actor). This definition seems parallel to the distinction between data and capta proposed by Checkland and Holwell [17], in particular, recognising in historiographical uncertainty the notion of capta.

Historical interpretation relies heavily on a bibliography, which forms the backbone of scholarly research and critical discourse. Critical engagement with primary and secondary sources traces the evolution of ideas, ensures that arguments are well-supported, and fosters the historical, cultural, and philosophical contextualisation of the research. Suppose a source or a certain feature is no longer available. In that case, scholars must rely on testimonies and secondary accounts, making their discourse heavily dependent on what they read or see.

Gadamer's seminal work "Truth and Method" [45] deeply analyses the nature of understanding and interpretation. Gadamer challenges the positivist notion that truth can be objectively measured and verified through empirical data alone. Instead, he argues that truth emerges through understanding deeply rooted in historical and cultural contexts. This hermeneutic approach aligns with the concept of capta, where knowledge is constructed through interpretation rather than merely collected as objective facts.

Gadamer claims that every act of understanding involves a dialogue between the present and the past, where the interpreter's preconceptions and the historical context of the subject matter come into play. This process is not about eliminating subjectivity but about recognising and incorporating it into the act of interpretation. This perspective aligns with Piotrowski's distinction between historical and historiographical uncertainty, pointing out that our understanding of history is always influenced by the historian's interpretive framework and the socio-cultural context in which they operate.

Moreover, Gadamer's notion of "effective historical consciousness" underscores the dynamic nature of understanding, where the meaning of a text or an artefact is continuously shaped and reshaped by its interaction with interpreters over time. This idea complements Drucker's argument that capta represent situated and constitutive knowledge, emphasizing that what is considered significant or true is often a product of historical and cultural positioning. Thus, the scholarly pursuit is discovering pre-existing truths and engaging in an ongoing interpretive process that acknowledges the evolving nature of meaning and understanding.

Gadamer's definition of hermeneutics offers a framework that embraces the subjectivity and interpretive nature of humanistic inquiry, which seems to be in contrast to the positivist approach adopted by data-centric methodologies in STEM (Science, Technology, Engineering and Mathematics) fields. This does not diminish scholarly work's rigour but enriches it by incorporating a deeper awareness of the role of context, perspective, and historical contingency in shaping our knowledge. Through this lens, the true complexity of scholarly activity becomes evident, where the narrative of how scholars arrive at certain conclusions is as important as the conclusions themselves.

More broadly, the practice of *observing* transcends disciplinary boundaries is fundamental to all scholarly work. However, intrinsic characteristics may differentiate the methodologies of the humanities and STEM fields, particularly in how scholars approach observation, uncertainty, and hypothesis testing. To illustrate these differences, observational methods employed in astronomy, a representative STEM discipline, will be compared to those used in the humanities.

Both Piotrowski and Gadamer compare humanities methodology to STEM disciplines (respectively with meteorology [45], natural sciences [86]). Later in this Dissertation, this comparison will be further explained, comparing CH data with astronomical data. The choice motivating this comparison is that both fields involve studying unique objects, such as stars or books. Yet, the way data is treated differs, with astronomical observations becoming scientific data as soon as they are used as evidence of phenomena [13], while humanities rarely can go beyond learned interpretations.

Data sources also vary, with humanities researchers using historical documents, literature, art, and oral traditions, each having varying levels of reliability and introducing systemic and insurmountable uncertainty. In astronomy, uncertainty is often related to instrumental limitations and observational conditions. Methodologically, astronomy relies on empirical observation, mathematical modelling, and experimental validation. Humanities research is frequently interpretative and qualitative, and the necessary proof to obtain historical certainty is often unattainable [12, 86]. This difference leads to distinct epistemological foundations, with the humanities acknowledging subjectivity and cultural bias in interpretations [45], and astronomy seeking to minimise them through rigorous data collection and adherence to physical principles [27].

This Dissertation hypothesises that annotators in cultural heritage and astronomy (as a stance of STEM disciplines) may approach data observations, incompleteness and uncertainty differently, with cultural heritage favouring qualitative, context-rich representations of competing hypotheses and astronomy leaning towards more quantitative, data-centric representations. This difference may reflect broader epistemological stances in their respective communities. Furthermore, this study assumes that these distinct approaches to handling data incompleteness and uncertainty may impact the ease of integrating data from these fields in interdisciplinary research, with cultural heritage data potentially requiring more effort for reconciliation due to its contextual and subjective nature.

1.2 Guiding Examples

Two guiding examples illustrate the challenges and intricacies of representing the evolving interpretations and complex narratives inherent in humanities research. In particular, two main examples were selected, each involving a distinct type of cultural property subject to a different form of critical assessment, analysed from the perspective of different humanities disciplines. namely:

- The restoration of the painting "Girl Reading a Letter at an Open Window"
- The challenge of the authenticity of "The Donation of Constantine"

Figure 1.1: The painting before the restoration

Figure 1.2: The painting after the restoration

1.2.1 Evolving knowledge in Art History: Girl Reading a Letter at an Open Window

The painting "Girl Reading a Letter at an Open Window" (ca. 1657) by Jan Vermeer depicts a pensive lady near an open window. Until 2021, there was just a white wall behind the girl, as shown in Figure 1.2. Traditionally, the work was mildly associated with the idea of love. In 1979, an x-ray inspection revealed that the wall concealed a Cupid (symbol of love) trampling on a theatre mask (symbol of hypocrisy), likely whitewashed long after the painting's creation. In 2021, after a complete restoration, a decisively new interpretation emerged, claiming the painting to represent the love that overcomes treachery and hypocrisy [101, 50].

The new interpretation depends on the changes in available data, as well as on their truth value. Initially, undisputed claims were superseded (the painting symbolises love), and new meanings were produced (the painting symbolises love that overcomes treachery and hypocrisy). Yet, many reproductions of the old version of the painting exist and are still shared, and multiple scholarly works were written before discovering the hidden detail. Silently replacing the old interpretation with the new one could be misleading: viewing an old image next to the latest interpretation, or vice versa, would be at the same time wrong and surprising[1].

Moreover, restoration is not the only significant aspect of the evolving understanding of painting. Throughout its history, it has been attributed to Rembrandt, Hooch, and, most recently, Vermeer, which is currently the accepted attribution[2].

1.2.2 Challange on Documents' Authenticity: The Donation of Constanine

Historical authenticity assessment is the scholarly practice of determining the authenticity of historical documents. Scholars from different humanities and scientific disciplines (e.g. Diplomatics, Palaeography, Philology, History, Forensics) have contributed to the field [9]. Frequently, different scholars arrive at divergent and possibly contrasting conclusions due to different evidence. Inherent factors contributing to this diversity include historical uncertainty, gaps in documentary transmission, and subjectivity [12, 45].

The "Donation of Constantine" is a medieval forgery purported to be a decree by Roman Emperor Constantine the Great, transferring authority over Rome and the western part of the Roman Empire to the Pope. This document, believed to have been written in the 4th century, was widely accepted as authentic for centuries and used by the Church to justify its temporal power.

In the 15th century, Lorenzo Valla, an Italian humanist and scholar, exposed the document as a forgery through his philological analysis. Valla demonstrated that the Latin used in the Donation was not of the 4th century but contained anachronisms indicative of a much later period (8th century). Despite Valla's convincing arguments and evidence, the Church initially resisted accepting his findings due to the implications for its authority.

The "Donation of Constantine" is one of the most famous forgeries of

[1] The two images can be compared at `http://www.wikidata.org/entity/Q700251`
[2] `https://en.wikipedia.org/wiki/Girl_Reading_a_Letter_at_an_Open_Window`

all time, and many scholars and relevant figures took part in the discussion through the centuries (e.g. Martin Lutero included the Donation in the index of the prohibited books). As with the other examples, the truth about the Donation evolved as new methods and data became available. The revelation of its inauthenticity did not immediately erase the centuries of belief and scholarship built upon the false document. Instead, it highlighted how historical narratives and interpretations are constructed, challenged, and eventually revised.

In summary, while each example pertains to different fields within the humanities — art history, history, philology and diplomatics — they collectively illustrate a shared pattern. This pattern involves continuously reassessing and revising interpretations based on new data, technologies, and methodologies. Whether through technological advancements in art restoration, historical analysis, or critical philological scrutiny, these disciplines exemplify how the humanities evolve in response to discoveries and scholarly advancements. The annotation and evolution of critical inquiry—including used methodologies, collected evidence, and contrasting sources—shift the focus from cultural goods as data to cultural goods as capta.

1.3 Problem statement

For a long time, digital history has been primarily concerned with digitising sources and deriving data from sources [86]. For practical reasons, the focus has thus been primarily on historical uncertainty, typically due to missing, inexact, partial, and ambiguous information. So far, there has been little work on formalising the historiographical uncertainty of critical inquiry (capta).

Considering the two guiding examples presented in Section 1.2, "Girl Reading a Letter at an Open Window" is catalogued in various CH platforms: Bildindex[3], Netherlands Institute for Art History (RKD)[4], Google Arts and

[3] https://www.bildindex.de/document/obj00021821

[4] https://rkd.nl/images/224949

Culture[5], and The Staatliche Kunstsammlungen Dresden (SKD)[6].

These platforms predominantly feature metadata describing the painting's physical attributes, medium, authorship attribution, and creation date. However, metadata concerning the painting's restoration is typically absent. Detailed descriptions of the restoration methodologies and iconographical analysis are usually found in supplementary articles (see, for instance, the Gemäldegalerie article from The Staatliche Kunstsammlungen Dresden[7]). Additionally, the painting is listed in Wikidata[8], where both images are shown: one pre-restoration (as it was until 2017) and one post-restoration (as it since May 2020). Changes to the depiction, such as the discovery of a hidden Cupid, are annotated as having begun in 2021.

The record for the "Donation of Constantine" can be found in the Bibliothèque nationale de France (BNF)[9]. The entry provides an imprecise date ("07..") and notes in natural language that it is an "Forgery made in the 8th century to attest to the alleged gift made by Emperor Constantine to Pope Sylvester of the city of Rome and part of the West. - Numerous ms including BnF, Paris (Lat. 2777, 9th century)". The sources used to create the catalogue record are cited for credibility and trust, but there is no mention of the document's alleged 4th-century origin. In Wikidata[10], the "Donation of Constantine" is classified as a historical forgery with an unknown author and an 8th-century publication date. Similar to the BNF entry, there is no reference to the disputed metadata or historical-critical inquiries performed on the document.

Several approaches may be adopted by annotators when no appropriate representation of such information is provided:

[5] https://artsandculture.google.com/asset/donna-che-legge-una-lettera-davanti-alla-finestra/3wFQaidzxA5mqg

[6] https://skd-online-collection.skd.museum/Details/Index/415429

[7] https://gemaeldegalerie.skd.museum/ausstellungen/der-neue-vermeer/

[8] https://www.wikidata.org/wiki/Q700251

[9] https://catalogue.bnf.fr/ark:/12148/cb123807414

[10] https://www.wikidata.org/wiki/Q238476

- **Reticence**: information that was well known was not recorded due to haste, lack of skill, or, more probably, lack of policies and software support

- **Flattening**: different kinds of information were represented all together with no differentiation between different types of information. Usually, a single set of metadata describes the image, the physical support, the depiction (and even other entities in the depiction), generating an inconsistent and meaningless record

- **Coercion**: important information for which no appropriate field was found was forced into inappropriate fields, leaving the reader to make things straight and forever baffling any automatic tool tasked with indexing and searching collections by subject

- **Dumping**: important information for which no appropriate field was found was forced as plain text inside a descriptive field, easy for humans to read but forever lost to any automatic tool

Data and capta are collected in the same CH digital collections without differentiation. The online records of "Girl Reading a Letter at an Open Window" and the "Donation of Constantine" demonstrate that critical inquiries and their evolution are often excluded from structured metadata and are relegated to free-text sections. This approach neglects the representation of critical inquiries and their interconnections, limiting the analysis data only to a single factual perspective. Usually, the data-modelling activity concentrates on representing cultural entities (e.g. books, manuscripts, paintings) and objects of artistic and scholarly interest. In this context, the representation of critical inquiry seems to be a second-order priority, and they are often not documented in digital catalogues or archives (and, therefore, in their databases). In other words, no distinction is set between the metadata recording, e.g., the language of a book (instance of data) versus, e.g., its dating attempts (instance of capta).

Introducing the concept of capta in KGs can involve paying attention to scholalrs' knowledge and perspectives, recognising interpretation and context in shaping our understanding of the world and being open to alternative perspectives and interpretations. Ultimately, this approach can lead to a computable, more robust model of the world and the issues scholars seek to understand.

1.4 Research statement

The representation and study of data in the humanities present unique challenges due to the diversity and complexity of theoretical frameworks and interpretations. This research explores methodologies for effectively capturing and representing humanities inquiries as they evolve through history in LOD environments and KGs.

Following the arguments presented in this section, definitions are provided to clarify the terminology used throughout this Dissertation.

> (C1) Critical inquiry elaborates a scholarly opinion from observations and analyses.

In other words, a scholarly opinion is a coherent set of concepts or propositions developed through research, analysis, and observation to explain, predict, or understand phenomena within a particular field. Inquiries are typically grounded in evidence and tested through academic practices. This is rooted in the concept of capta, underlying that data and evidence are actively captured and interpreted to form coherent scholarly inquiries and conclusions.

> (C2) Scholarly opinions often involve a degree of uncertainty and questionability, which stems from the inherent ambiguity in historiographical interpretation

Therefore, questionability should not be reduced in KGs but recognised as an inherent aspect of the scholarly methodologies and frameworks within the humanities. Capta, for their nature, can assume what will be referred to as

Weaker Logical Status (WSL), i.e. a hypothesis or, in general, a statement whose truth value cannot be given for granted.

> (C3) Each scholarly opinion is substantiated by evidence, sources, and methodologies to support a particular conclusion.

Provenance information is a general term referring to contextual information describing a scholarly observation. Such information can regard the date of creation of an interpretation, the argumentation behind it, the motivation, the author, etc. Representing these elements is crucial for understanding the validity and context of scholarly claims or assignments.

Provenance also provides what is called *trust*: by specifying sources, timing, and reasoning, the data consumer can decide whether to adhere to a given information. This point is not trivial: observational data should be transparent to enable the so-called "interpretative chain" (e.g. an interpretation based on another interpretation). Methodologically, the process of producing observational data must occur collaboratively to allow and facilitate the exchange of opinions and interaction among interpreters. In the case of absolute truth being unattainable (e.g. the subjective interpretation of the meaning of a cultural object), recording the provenance of each interpretation allows data consumers to choose to adhere to or refuse a trend of thought.

Provenance information extends beyond methodologies and collected evidence (see C3) to consider prior studies on the subject and related materials and sources.

> (C4) Scholars often relate their studies to past inquiries over the same subject of analysis, which they may believe or disbelieve to

The introductory statements of the new CIDOC-CRMinf primer [36] resonate significantly with the theoretical visions outlined in this Dissertation [36]:

> If scientists and scholars, and in particular curators, would start documenting for each information source the provenance of its immediate sources in publicly accessible systems, this partial knowledge of provenance could be "stitched together" to more and more complete networks of provenance, similar to the way these days citations in scientific publications are processed. This is a major motivation for CRMinf, the other is to make the way transparent how knowledge was acquired for enabling justified future revisions, and who is supporting contested propositions.

In summary, CIDOC-CRM recognises the need to introduce constructive and partial knowledge of humanities discourse in its representation and management. This point is non-trivial; critical engagement with primary and secondary sources traces the evolution of scholarly ideas, ensures that arguments are well-supported, and fosters the research's historical, cultural, and philosophical contextualisation. The recording and structuring of critical discourse keeps track of the evolution of knowledge towards a specific topic and, therefore, may reconstruct the scholarly narrative developed through history, potentially enhancing the diachronic understanding of scholarly knowledge.

> (C5) Different inquiries may report divergent opinions on the same topic, which can be categorised with respect to their position within the debate

Building on Checkland and Howell's distinction between data and capta [17], and their respective epistemic value, an additional classification is introduced to assign claims a logical status based on their position in the debate. This categorisation reflects the degree of contestation a claim has undergone, ranging from uncontested assertions to those debated.

- **Undisputed statements**: no one has questioned the claim's validity as they are not part of any dispute or debate. This definition aligns with Checkland and Howell's definition of data [17].

- **Disputed statements**: Claims that have been part of a dispute (or debate) cannot be assigned a definitive truth value, as they coexist with at least one competing claim (or capta, as defined by Checkland and Howell [17]). These statements have been challenged and may include hypotheses. A disputed claim is not necessarily false; it remains questionable and holds a weaker logical status.

- **Settled statements**: while recognising disagreement, the relevant community has examined competing stances (disputed statements), has chosen one of them, and has closed the dispute.

- **Superseded statements**: claims that a settled statement has replaced. The term superseded has been chosen over *rejected* since rejection implies that a claim is definitively false. In contrast, superseded acknowledges that the claim has been displaced by a more widely accepted statement, reflecting the evolving nature of historiographical interpretation.

- **Valid statements**: statements regarded as true, either because they remain undisputed or have been settled after the debate.

- **Statements in unsettled dispute**: statements that remain part of a dispute with no resolution. The latter acknowledges historiographical uncertainty, which often precludes definitive judgments.

This dissertation aims to develop a formal representation that captures this representation while distinguishing between them with minimal complexity.

> (C6) A set of fairly common situations (e.g. ignorance, evolving knowledge, disagreement, challenge) can lead to such contrasting opinions.

Therefore, claims are considered to evolve as new evidence emerges or scholarly perspectives shift, reflecting the dynamic nature of scholarly discourse.

We could consider several usual situations in which scholars express statements they do not want to assert.

- **Ignorance.** Scholars may not know the true answer to a question or may believe that no definitive answer to some question will ever be found (e.g. *Who was Jack the Ripper?*), and they can only provide (more or less justified) hypotheses.

- **Evolving knowledge.** New arguments or data change scholars' claims that used to be considered true and now are doubtful or discredited. Vice versa, a claim considered unlikely was eventually accepted. Lasty, new discoveries can lead to new claims. For instance, the discovery of hidden parts of a painting shed new light over its intended meaning and purpose.

- **Disagreement.** Two or more points of view exist, and there is no obvious justification to choose one opinion over the others. This is a frequent occurrence in artwork attributions, e.g. the 450M$ painting *Salvator Mundi* is by different scholars attributed to Leonardo da Vinci or to one of his pupils.

- **Challenge.** A claim that scholars want to consider false or unacceptable is precisely presented so that its consequences can be challenged, i.e., examined and discussed, to point out their flaws and therefore make the original claim untenable. This is common in mathematical proof by contradiction (or *reductio ad absurdum*). This is the typical outcome of a debate about the authenticity of a cultural object: an artefact is considered authentic until a scholar questions its authenticity, e.g. the case of the "Donation of Constantine".

This Dissertation starts with a set of theoretical assumptions and hypotheses and mainly focuses on two key research questions outlined as follows.

- Assumptions

  - (A1) The distinction between data and capta made by Checkland and Howell [17] is acceptable within data collection and formalisation
  - (A2) Cultural heritage data potentially requires more effort for reconciliation due to its contextual, uncertain and subjective nature
  - (A3) The evolution of critical inquiry can be expressed in RDF

- Hypotheses
  - (H1) Annotators in CH may approach data observations, incompleteness and uncertainty in different ways, favouring qualitative, context-rich representations of competing hypotheses
  - (H2) A comprehensive representation of scholarly opinions can reflect the evolution of critical inquiry on a given topic

- Research Questions
  - RQ1 - *What are the current limitations and potential improvements in the representation of humanities critical inquiry in current Knowledge Graphs?*
  - RQ2 - *How can critical inquiry be effectively and efficiently represented with minimal complexity?*

These research questions will be further specified by a series of sub-research questions throughout the findings of this Dissertation.

In conclusion, the study and interpretation of cultural artefacts, historical events, and literary texts are driven by evolving understandings and shifting perspectives. This Dissertation explores the formal representation of humanities scholarly theories and interpretations, emphasising how information updates can reshape scholarly narratives. This chapter laid the groundwork for this Dissertation by presenting its theoretical framework and identifying the research gap, pointing out that the representation of evolving information and critical debates in structured metadata is often overlooked despite its importance.

> I believe certainty regarding that which we see and touch — it is seldom justified, if ever. Down the ages, from our remote past, what certainties survive? And yet we hurry to fashion new ones. Wanting their comfort. "Certainty"— it is the easy path, just as you said.

Billy Knapp - *The Ballad of Buster Scruggs*, Joel and Ethan Coen, 2018

Chapter 2

RDF Representation of Critical Inquiry

Over time, several approaches have been developed and adopted to address statements in RDF, particularly regarding the concepts of questionability and uncertainty, but no current standard exists. This chapter offers an initial answer to RQ1— *What are the current limitations and potential improvements in the representation of humanities critical inquiry in current Knowledge Graphs?* — by analysing the current state of such representation.

2.1 Critical inquiry in Cultural Heritage Knowledge Graphs

Although domain catalogues (as mentioned in the Problem Statement, Section 1.3) representing the CH domain struggle to incorporate support for scholarly opinions or historical perspective into their models [88], there is a growing interest emerging in the field [85, 37]. This section addresses KGs and presents data models that tackle this issue.

The Europeana KG [85] stores about 50 million heterogeneous digitised items from museums, archives and libraries all over Europe. These collections include materials in various human languages and cover a wide range of subjects, from art and history to science and literature. Data is gathered by the content providers (i.e., usually, participating cultural institutions) according to the EDM data model [34]. EDM employs "proxies" to represent different and possibly conflicting information on objects being described, allowing tracking of the provenance of such data. For instance, the title of the Mona Lisa is listed as "Portrait de Mona Lisa" according to one proxy, and "Portrait de Lisa Ghirardini" for another [61]. However, EDM proxies are aimed at handling contrasting claims from different institutions and are not specialised in the representation of interpretative contents or critical inquiries. However, concurrent statements are not displayed on the online web pages, and there is no known indication of how the accepted proxy is determined over multiple competing statements.

CIDOC CRM [37] is a conceptual model, developed and maintained by the International Council of Museums (ICOM), widely adopted by many KGs in the CH domain [28, 42] related to visual heritage and museums. Interpretations and critical inquiries can be represented in CIDOC-CRM using n-ary relations. The most general class to address the issue is `crm:E13_AttributeAssignment`[1] intended as an activity (subclass of `CRM:E7_Activity`). It is intended to express "the actions of people making propositions and statements during certain museum procedures, e.g. the person and date when a condition statement was made, an identifier was assigned, the museum object was measured, etc. [...]" [77]. Additionally, it declares that "Multiple use of instances of E13 Attribute Assignment may possibly lead to a collection of contradictory values" [77], but does not specify such incompatibilities can be handled (e.g., SPARQL queries, reasoning), nor provides a defintion on how to handle former or disputed claims.

The Linked Art application profile provides explicitly the documentation use CIDOC-CRM to represent provenance, assignment of attributes

[1] `https://cidoc-crm.org/Entity/e13-attribute-assignment/version-6.2.1`

(assertions)[2] and uncertain or former attributions[3]. In particular, it suggests attributing to an E13 Attribute Assignment a classification from the Getty Art & Architecture Thesaurus (AAT) vocabulary[4] (e.g., `aat:00404272` possibly, or `aat:300404776` former). This conveys the "former" status to the attribution but does not solve the problem of possibly incompatible statements.

Additionally, the CIDOC-CRM provides a set of official extensions designed to expand the scope of the core model, particularly for supporting critical inquiry as scholarly and scientific work. Among these, the CRM Argumentation Model (CRMinf)[5] [36] and the Scientific Observation Model (CRMsci)[6] [35] serve as formal ontologies for integrating metadata related to argumentation, inference-making, and scientific observation. CRMinf represents argumentations by formalising concepts such as propositions, justifications, and evaluations, while CRMsci extends this to observations, hypotheses, and the interpretation of scientific data.

Despite making heavy use of n-ary relations, the CIDOC-CRM primer states that "The encoding structure known as a **named graph** also falls under this class, so that each named graph is an instance of E73 Information Object." CIDOC CRM does not provide a specific guideline about how to include reification methods with the CIDOC model, and which kind of reification method should be employed is still under discussion[7]. Additionally, the CRMsci primer states about the definition of the data property `crm:O9_observed_property_type` "[...] In an RDFS encoding, this circumscription can be transformed into an explicit representation of the observed property in terms of a formal ontology either by using a reification construct

[2] `https://linked.art/model/assertion/`

[3] `https://linked.art/model/assertion/#uncertain-or-former-assignments`

[4] `https://www.getty.edu/research/tools/vocabularies/aat/`

[5] `https://www.cidoc-crm.org/crminf/home-4`

[6] `https://cidoc-crm.org/crmsci/home-1`

[7] See the discussion about Issue 526: Named Graph Usage Recommendations / Guideline Document at the link `https://cidoc-crm.org/Issue/ID-526-named-graph-usage-recommendations-guideline-document`

or a named graph containing the observed property. The latter representation allows for more formal reasoning with the model, the former is more flexible about the kinds of observations".

Several aspects of CRMinf and CRMsci are currently under discussion within the CIDOC-CRM community. As mentioned in the Research Statement (Section 1.4), many complications arise in formal models when introducing partial and constructive knowledge. Several critical gaps have been identified in how CRM inf beliefs, proposition sets, assignments, and observations are represented. Notably, the incorporation of beliefs raises complications related to the "truthfulness" of assertions based on specific beliefs[8].

Suppose different beliefs coexist on the same factual knowledge. Therefore, such beliefs are represented as instances of the class I2 Belief. Their degree of belief or disbelief with such factual proposition is given by the use of instances of I6 Belief Value, which may have the values "true", "false," and "possibly true"[9]. Still, it is not included in the CRMinf official primer v.1.0 [36][10]. Figure 2.1 shows the possible relations with proposition sets (`crminf:I2_Proposition_Set`) and related beliefs (`crminf:I2_Belief`) respectively, possibly believing (`:belief1`), completely believing (`:belief2`) and disbelieving (`:belief3`) a proposition set (`:facts`).

[8] This is documented in a recent CIDOC-CRM Issue available at `https://cidoc-crm.org/crminf/Issue/ID-614-definition-of-i4-proposition-set-and-what-an-instance-of-i2-belief-is-about`

[9] This part is currently documented in an ongoing draft of the CRMinf primer v.1.1 available in this Google Document

[10] The complete definition of I4 proposition set is reported here for clarity: "An instance of I4 Proposition Set should be regarded per se to be neutral to its relationship to reality. The relationship to reality is determined by the link using the proposition set: If an instance of I2 Belief refers to an instance of I4 Proposition Set, the belief value of "TRUE" will mean that the propositions are believed to correspond to reality; if the propositions can be related to reality (i.e., are about real-world items, in contrast to, e.g., mathematical statements). "FALSE" would mean that at least one of the propositions in the set is regarded as not corresponding to reality. Belief values expressing possibility or probability will mean "possibly real" if the propositions can be related to reality"

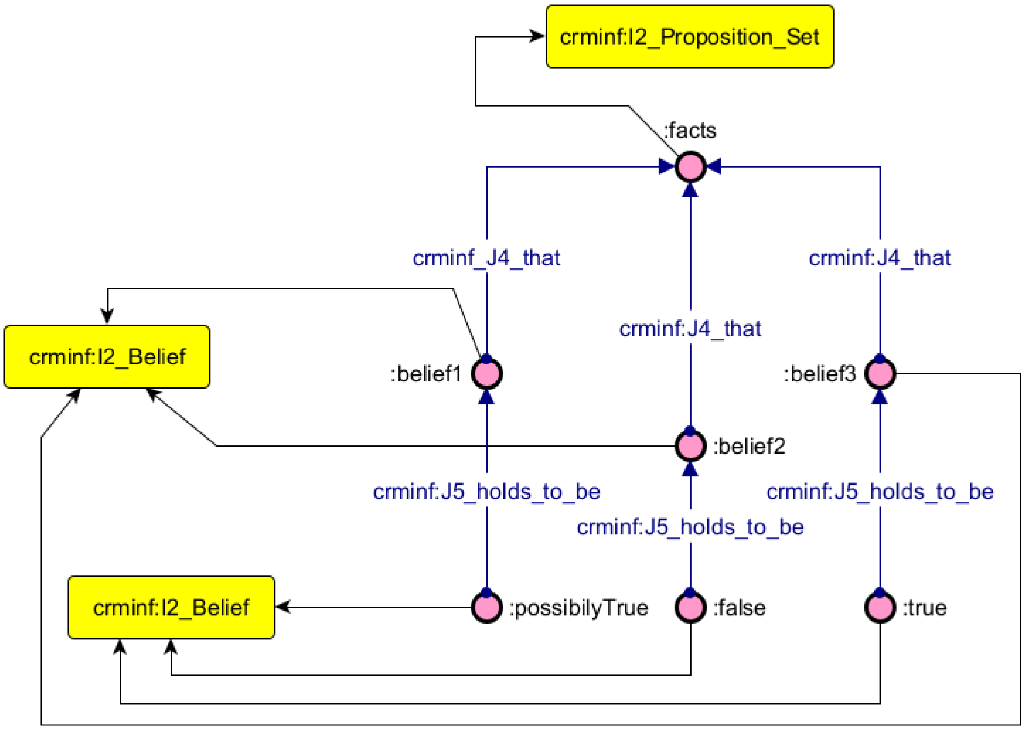

Figure 2.1: Possible relations between beliefs and disbelief of propositions in CR-Minf

The primer additionally states that instances of I4 Proposition Sets can be represented in KGs as a named graph. However, the named graph must contain factual knowledge deemed true by the KG maintainers as such contents become part of the knowledge of the entire KG. In simpler terms, propositions must be part of that system's established, accepted knowledge. Propositions that are possibly true, not confirmed, but worth tracking can also be represented as named graphs if the KG has mechanisms for filtering by provenance (where the information came from) and likelihood (how probable the information is to be true)[11].

[11]The complete description of the introduction of named graphs states as follows: "In a Knowledge Base implementation, an instance of I4 Proposition Set may be represented by the URI of a named graph, but only if the propositions are encoded in the data model of the Knowledge Base and held to be true by the maintainers of a Knowledge Base because they become part of the stated knowledge. In this case, the platform-internal relation between the URI and its content are regarded as equivalent to the property J25 is encoded by. Proposition Sets held to be possibly true by the maintainers of a Knowledge Base may also be introduced as named graphs, if the operation of the Knowledge Base foresees filtering by provenance and likelihood. In this case, named graphs are particularly effective."

The main issue with this approach is that it negates the representation of discarded but historically relevant proposition sets. For clarity, consider the case of the "Donation of Constantine" as an example: the document is now widely accepted as an 8th-century forgery (Proposition Set). Additionally, Lorenzo Valla believed this proposition (Belief). However, this is not the complete story. As claimed in the Introduction (Section 1.2.2), the Donation has been widely believed to be an authentic document made in the 4th century. According to the CRM primer, this information cannot be introduced as a concurring Proposition Set into the KG, as it's no longer considered "true".

CRMinf and CRMsci generally represent the scholars' reasoning processes underlying scholarly opinions (e.g., argumentations, observations). They enable scholars to document final judgments and trace the supporting evidence and inferential steps. However, the discussion on constructive knowledge has been recently introduced, and existing implementations of CH datasets modelled with CIDOC-CRM [93, 14, 66] primarily address the provenance of scholarly claims from a single perspective. Such modelling approaches focus on capturing provenance and critical inquiry within the contexts of museum and archaeological research [14, 93], but do not consider the integration of contrasting interpretations, disputed viewpoints or historical knowledge.

Nussbaumer and Haslhofer [76] express some concerns when the CRM comes to an actual implementation and retrieval, i.e. through SPARQL queries on RDF data. In a data integration scenario, such a global ontology abstractness of concepts (classes) can be ambiguous to any human user. Therefore, its usage leads to the risk of different incompatible conceptualisations by different institutions. Second, its verbosity (given by the event-centric approach) requires long and complicated queries. Lastly, the model lacks technical specifications on implementing mappings, representing instances, and processing data. To this extent, each institution likely applies its interpretation of the model.

Additionally, CRM-inf detailed formalisation may be excessive for simpler use cases, and the need to represent multiple argument components (inferences, evidence, beliefs) can result in redundancy, increasing data storage requirements and impacting query performance in large-scale systems. These factors may deter adoption in scenarios where more streamlined approaches could suffice. For instance, the *Index Graecorum Vocabulorum*[12] [43] is a KG which integrates ancient greek loanwords in the LiLa Knowledge Base of Linguistic Resources for Latin[13]. It reuses CRM-inf belief (`crminf:I2_Belief`), believed value (`crminf:I6_BeliefValue`) and the Open Vocabulary `ov:confidence` to introduce the notion of uncertain interpretative content of the scholarly assertion. In this case, CRM is used to express uncertainty rather than critical inquiry, and the chain of entities is excluded from the full CRM-inf model. Again, the Index tackles the representation of uncertainty from a single point of view, and no contrasting opinion is collected in the dataset.

The Netherlands Institute for Art History (RKD) catalogue[14], handles a comprehensive collection of data about Dutch works of art throughout history. By design, RKD allows and gathers contested and discarded attributions of paintings and portraits. RKD data can be downloaded in Dublin Core, EDM and LIDO[15][16]. Each record may include concurring, rejected and accepted artwork's attributions[17]. Such claims are additionally annotated with a description of the attribution source or underlying motivations. Unfortunately, these annotations are only available in natural language, and while they are

[12] `https://github.com/CIRCSE/index-graecorum-vocabulorum`
[13] `https://lila-erc.eu/#page-top`
[14] `https://www.rkd.nl/en`
[15] `https://www.rijksmuseum.nl/en/research/conduct-research/data/overview`
[16] `https://data.rijksmuseum.nl/object-metadata/download/`
[17] See for instance the painting "Portrait of a young lady as Mary Magdalene" available at `https://rkd.nl/images/51284`

available in the interface, they are not included in the XML data dump.

The Lightweight Information Describing Objects (LIDO)[18] is an XML harvesting schema. The data model is intended to deliver and use metadata in online services, from an organisation's online collections database to portals of aggregated resources. Therefore, data can be exposed, shared, and connected on the web. It is compliant with LOD standards. The RKD data dump encoded with LIDO represents attributions, essentially streamlining CIDOC-CRM E13 Attribute Assignment as `lido:attributionQualifierActor`, associated with the label "rejected attribution". However, the data dump does not correspond to the data in the catalogue interface, and a SPARQL endpoint is not currently available.

Interestingly, about 83.600 artwork descriptions from RKD[19] have been imported into Wikidata, representing ~7,5% of the total of visual artworks in Wikidata[20].

Public Knowledge Graphs such as Wikidata [40], DBpedia [6], Yago [91], and Google Knowledge Graph constitute publicly available collections that can be used for research, either expressing specialist or general knowledge. In particular, Wikidata is a collaborative public platform built and maintained by a community of contributors. The widespread adoption of Wikidata within the CH community has been well-documented, and almost 100 cultural institutions are involved [102][21]. Wikidata is seen not only as a valuable tool for data publishing, alignment and enrichment but also as a means of gaining valuable insights into CH data and the community itself [109].

Over the years, Wikidata has developed and provided a variety of representation methods that allow it to encode complex structures far beyond factual descriptive metadata. According to [71], Wikidata encompasses a mul-

[18] `https://lido-schema.org/documents/primer/latest/lido-primer.html#linked-data`

[19] `https://w.wiki/7wfW`

[20] This percentage is calculated over the Cultural Heritage Visual (CHv) dataset illustrated in Section 3.2 and discussed while comparing RKD and Wikidata items 3.5

[21] The list can be found at `http://www.wikidata.org/entity/Wikidata:GLAM`

titude of facts, including some that may be contrasting since they come from different and disagreeing sources. Additionally, time-sensitive information can also be added through the use of qualifiers and ranks. For instance, structures to represent temporally evolving information (e.g., the number of followers of a YouTube Channel that is updated year after year) or multiple coexisting (and possibly competing) claims over the same subject (e.g., maintaining both the old as well as a new theory over some topic). In many such cases, multiple information items are present. Yet, newer or better information is not replacing older or less true assertions. Still, they coexist next to each other, and one or more mechanisms are used to signal their simultaneous presence and, when appropriate, the currently adopted stance (this aspect will be tackled in detail in Section 3.1.1 in particular in the description of Wikidata ranked statements).

Wikidata data model uses its own custom reification method and the Wikibase data model as reification method [40, 31] to express statements. Statements make use of rankings[22] - Deprecated, Preferred and Normal - to annotate the logical status of claims (e.g., a former attribution can be marked with a deprecated rank to mark that it has been disputed explicitly).

Additionally, contextual information is annotated over statements via qualifiers[23], providing more detail about a statement itself. Wikidata provides a list of 356 qualifiers[24]. Some of them are designed to qualify annotators' argumentation, such as *determination method* (`P459`) and *criterion used* (`P1013`), or to explicit the reason behind a certain decision towards a statement as *reason for deprecated rank* (`P2241`) and *reason for preferred rank* (`P7452`), or express the underlying circumstances of a statement as *sourcing circumstance* (`P1480`) or to address the qualification of the truth or accuracy of a source or its interpretative nature as *nature of statement* (`P5102`). These aspects will be addressed in detail in a later section dedicated to the available patterns to

[22] `https://www.wikidata.org/wiki/Help:Ranking`
[23] `https://www.wikidata.org/wiki/Help:Qualifiers`
[24] The complete list of available qualifiers in Wikidata is available at `https://w.wiki/6TrP`

express WLS claims in Wikidata (Section 3.1). However, such qualifiers are highly generalised to fit the comprehensive knowledge in Wikidata.

Wikidata also provides sources[25] to ensure the verifiability of its statements. A reliable source, such as a book, scientific publication, or newspaper article, should support each statement. In Wikidata, references link specific sources that substantiate the data provided in a statement. However, these sources are typically marked with URLs.

Wikidata supports the representation of scholarly opinions through various modelling patterns. However, these representations often rely on structured data that may not fully represent the interpretative nature or contextual nuances of scholarly opinions in the humanities.

The Historical Context Ontology (HiCO) [26] is an extension of the Provenance Ontology (PROV-o) [63]. It represents the contextual framework surrounding scholarly claims, particularly those from hermeneutical activities, by addressing the reliability of argumentations and evaluating features such as used criteria and cited sources. It uses CITO ontology [84] to define relations with other claims (e.g., citation, agreement, disagreement).

Digital Hermeneutics [24] is a high-level, portable data model for representing hermeneutical aspects related to the cross-disciplinary analysis of archival and literary sources. Digital Hermeneutics make use of named graphs [15] to represent statements and Nanopublication [51] to express the so-called *layered approach* to knowledge representation. This term refers to the inclusion of all triples about the claim in a set of named graphs, therefore related to the Nanopublication data model to represent interpretation on four levels as shown in Figure 2.2: information not deemed to be questionable (factual data), reusing CIDOC-CRM and Fabio ontology [84] to model visual and archival sources; the subjective or questionable assertions (ontologies are not

[25] `https://www.wikidata.org/wiki/Help:Sources`

specified since they depend on the inquiry performed by cultural objects); assertion provenance information such as attribution, date, sources, and criteria supporting the interpretation, and the degree of certainty. It mainly reuses PROV-o to include provenance information, HiCo to introduce scholarly interpretation and CiTo to represent relations with sources and The Canadian Research Collaboratory (CWRC)[26] ontology to introduce annotators' certainty over each assertion.

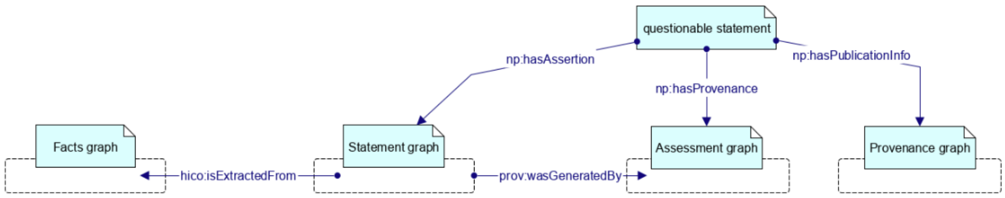

Figure 2.2: Overview of Nanopublication data model applied to the four layers of the Digital Hermeneutics data model

Three KGs are reported to reuse the model: *The lectures on Pellegrino Prisciani's Historiae Ferrari illuminated manuscript* [79] address the philological, palaeographic and art historical interpretations of scholars over the pages of an illuminated manuscript; *The scholarly digital edition of Paolo Bufalini's notebook* [23] reuse the model to represent modern philology inquiries on the excerpts of the notebook; *mAuth, a recommending system of artwork attributions* [21] integrates and stores historical contrasting attributions into the same KG and uses scholarly motivations and contextual information to rate them [21].

The Digital Hermeneutics data model aims to formalise critical inquiry in CH data at a higher level via a layered approach. It successfully formalises the distinction between data (represented in layer 0) and capta (represented in layer 1) defined by Checkland and Howell [17]. Named graphs allow the integration of several potentially concurrent scholarly theories.

The RELEVEN [3] project aims to illuminate the "short eleventh cen-

[26] `https://sparql.cwrc.ca/ontologies/cwrc.html#Certainty`

tury" (c. 1030-1095) by exploring how the Christian world was perceived on the eve of the First Crusade, particularly in regions more engaged in global trade networks. By shifting from "linked open data" to "linked open assertions," the project emphasizes the importance of contextualising data points and distinguishing between claimants and recorders, utilising the STAR model [4] to incorporate varying perspectives and ensure rigorous validation through SHACL in a Neo4J database.

The STAR model employs the CIDOC-CRM E13 Attribute Assignment to annotate assertions while integrating a bespoke ontology that introduces the notion of perspectives. Each perspective consists of one or more statements, where each statement can be either accepted or rejected. This approach models rejections as it addresses the frequent challenge of depicting information arising from a lack of evidence or statements that can only be articulated negatively. The project is currently in progress, and the dataset is inaccessible.

The ICON dataset[27] is a KG recording Panofsky's interpretations and uses the ICON ontology [96] to represent Panofsky's three levels of interpretation over artworks (preiconographical, iconographical and iconological). ICON ontology integrates several domain ontologies to address a complex network of interpretations. In the KG, data are modelled with CIDOC-CRM, while interpretations are recorded as instances of the class `icon:Recognition` (subClass of HiCo Interpretation Act) and its subclasses. CiTo [84] object properties are used to address cited evidence (`cito:citesAsEvidence`, e.g., other interpretations or bibliographic resources) and consulted bibliography (`cito:citesForInformation`). Additionally, sets of recognitions are grouped via the means of the class `icon:InterpretationDescription`. The dataset also includes concurring recognition sets from different scholars, but none of the interpretations is marked as the preferred point of view or the most authoritative.

[27] `https://raw.githubusercontent.com/SofiBar/IconologyDataset/main/data/icondataset.ttl`

The contributions discussed in this section suggest a set of requirements that would be considered to represent the evolution of critical inquiry (Section 1.3) and can summarised as follows:

- **Factual knowledge (data)**. This category includes factual and well-established information that is not questionable. It mainly described the objects studied by the scholars. Examples include formal models like CIDOC-CRM [4, 66, 79, 21], the Europeana Data Model (EDM) [34], and the Wikidata data model [40]. Additional standards like Dublin Core and the FaBiO ontology [26] are used to model bibliographic metadata and published works, respectively. Projects such as Digital Hermeneutics [24] also involve factual data graphs to represent knowledge that is widely accepted as true.

- **Scholalry opinion as a questionable entity (capta)**. "Capta" refers to data scholars actively address and discuss, which may imply subjectivity or uncertainty. In CIDOC-CRM, this can be modelled using the `E13 Attribute Assignment` class, while `CRMsci I6 Belief` can represent beliefs based on scientific evidence or observation. Similarly, Nanopublications, in combination with the HiCo ontology's `Interpretation Act` is used by Digital Hermeneutics [24] and VISU [66]. ICON `Recognition` is used in the ICON dataset.

- **Scholarly Method**: This category captures the methodological approach a scholar uses to derive a conclusion. It is often modelled using the CIDOC-CRM Argumentation model (`Argumentation`) or the HiCo `Criterion` and `Type`.

- **Studied Sources**: The studied sources refer to the bibliographic and archival materials a scholar consults in their research. CiTo relations (e.g., `cito:citesForInformation`) [84] are reused in the ICON dataset and Bufalini's Notebook [23]. Differently, these sources are typically linked as references in the Wikidata model, indicating the provenance of claims.

- **Confidence**: Confidence values indicate the degree of certainty attributed to a particular claim or belief. The CRMsci extension represents this by the `Belief` and `Belief Value` classes. Additional frameworks, such as Open Vocabulary or the CWRC ontology, can mark scholars' certainty explicitly, respectively, in the *Index Graecorum Vocabulorum* [43] and in Digital Hermeneutics [26].

- **Reification Method**: This category addresses the reification method used to represent scholarly inquiry. CIDOC-CRM employs n-ary relationships, which can be represented using blank nodes (e.g., Lila project [43]) or associating a URI to each n-ary relation (e.g., ICON dataset [96]). Named graphs and Nanopublications, as seen in Digital Hermeneutics [26], offer a way to encapsulate these relationships within self-contained assertions. Wikidata employs its own reification method [31], the Wikibase data model. This peculiar method will be addressed in the next sections of this Dissertation (in particular, 3.1 and 4.1).

- **Complex scholarly opinions**: This category addresses a set of scholarly claims that form a whole theory or perspective. This should be represented as a named graph (as in Digital Hermeneutics [24] or VISU [66]) through the use of n-ary relations (e.g., the recognition description in the ICON dataset [96] or the perspective in the STAR model [4]).

- **Competing scholarly opinions**: Competing capta refers to interpretations or inquiries that conflict with each other. This aspect seems to be overlooked by the surveyed works, neglecting concurring inquiries by different scholars. However, some exceptions exist. The ICON dataset records concurring interpretations as different recognitions performed on the same artwork by different scholars [96]. CIDOC-CRM use n-ary relations and event-centric approaches to represent such information. The mAuth project [26] uses named graphs to store contradictory attributions [21].

- **Opinions logical status**: Competing capta can be categorised based on the consensus within the academic community. Again, this aspect seems to be overlooked by the surveyed works, neglecting the natural evolution of concurring inquiries through history. However, Wikidata uses Wikibase Rankings [31] to define the logical status of such claims, while mAuth [26] does not explicitly state the logical status of such claims but rates depending on their motivations and contextual information.

This section examines the shared historical perspectives across various CH disciplines. It highlights recurring patterns in addressing critical inquiry and scholarly representation in current CH KGs and data models. Two distinct approaches to represent CH scholarly opinions emerge. One relies on CIDOC-CRM and its extensions to document scholarly work [37, 35, 36]. Additionally, CRM models suggest using named graphs with possibly contrasting or questionable statements. The other approaches [26, 66, 96] use named graphs with ontologies such as HiCo [26], PROV-O [63], and CiTo [84] to represent the context of critical inquiries. However, none explicitly addresses the logical status (or truth values) of claims except for Wikidata [40].

Although these models describe argumentations, they often overlook how interpretations interact when integrated into a single KG, particularly when considering competing or concurrent perspectives (e.g., differentiating between settled and former attributions). Despite efforts to represent contextual metadata, current CH ontologies fail to adequately address the logical status of claims (e.g. disputed, undisputed and settled) based on their context or how competing views and scholarly opinions interact. This aspect is further discussed in Chapter 3 analysing Wikidata CH records.

Additionally, CRM-inf and HiCo are the models that generally address critical inquiries in the humanities. Despite the support of representational definitions of weaker logical status claims in EDM, CIDOC-CRM and RDK data models, these weaker forms of information are often poorly reported (*reticence*) or are expressed in textual annotations rather than being modelled in the data structure (*dumping*) [8].

2.2 Expressing statements in RDF

As claimed in the Introduction of this work (Chapter 1), Linked Open Data (LOD) has been established as a standard for organising, publishing, and disseminating CH information on the Web. Plain RDF syntax allows statements (or claims) to be represented as triples (subject—predicate—object) [19]. Consider for instance, the statement "The Donation of Constantine is written in Latin" can be represented with an RDF triple as `:donationOfConstantine :language :latin`. Similarly, the statement "The Donation of Constantine has been created in the 8th century" can be represented with an RDF triple as `:DonationOfConstantine :date :8Century`. Although the two statements are represented in RDF in the same way, they represent two different types of claims: the first undergoes the category of (factual) data since the work categorisation as a decree revolves around those facts that are given and observed. The second statement is the outcome of Lorenzo Valla's philologic inquiry. The latter can be associated with the notion of capta and considered an instance of situated partial and constitutive knowledge.

However, the two triples do not convey any information on the source of such statements. To overcome this shortcoming, many approaches have been implemented to express *statements about statements*, adding triples to annotate contextual information about the statement itself. This approach is called **reification** and can be expressed via different approaches either compliant with RDF 1.1 syntax (e.g. standard reification [56] or n-ary relations [75]) or by extending RDF 1.1 syntax (e.g. RDF-star [54]).

Reification enables querying and reasoning mechanisms to integrate the following key attributes[72, 44, 73]:

- *provenance*: enables the identification and representation of the source of a fact or statement (i.e. *Who claimed this fact?*)

- *time*: communicates the fact's time-related information (i.e. *Which is the most up-to-date claim about this fact?*)
- *location*: reveals locational information about an event (i.e. *Which is the location in which this claim is applicable?*)
- *certainty*: indicates the level of confidence that is attributed to a statement (i.e. *What is the level of confidence regarding a certain event?* or *is a given statement true?*),
- *versioning*: it is helpful to keep track of RDF datasets' updating history (e.g. *What data version is currently being used?*)

With the introduction of a reification method, it is possible to represent "Lorenzo Valla claims that the Donation of Constantine was created in the 8th century". This claim is related to Lorenzo Valla's philological inquiry, published in his work "De falso credita et ementita Constantini donatione". Recording provenance data is significantly more critical when representing capta than mere data. The introduction of "claims" implies a certain degree of questionability (e.g., What happens if one disbelieves what Lorenzo Valla considers true?). Therefore, Valla's claim can be represented with its contextual information concerning its source and attribution, as shown in Listing 2.1 using standard reification. The claim `:Claim1` is a statement whose subject (`rdf:subject`) is the Donation, whose predicate (`rdf:predicate`) is "date" and whose object (`rdf:object`) is "8th century". Lorenzo Valla, who is responsible for this claim, is therefore represented in `:Claim1` via the predicate `:claimedBy` and the object `:LorenzoValla`.

```
:Claim1 a rdf:Statement ;
    rdf:subject :DonationOfConstantine ;
    rdf:predicate :date ;
    rdf:object :8century ;
```

```
    :claimedBy :LorenzoValla ;
    :source :DeFalsoCredita .
```

Listing 2.1: Representation of Valla's critical inquiry with standard reification

However, this is not the whole story. The document has been widely regarded as the 4th century, as Pope Leo IX referenced it in his arguments to uphold the papacy's temporal authority. Both claims can be introduced in the representation by adding a classification of the statements themselves (`:classifiedAs`) and the actual level of consensus within the community (`:acceptedOpinion` and `:disputedOpinion`) as shown in Listing 2.2.

```
# Valla's settled opinion on the Donation
:Claim1 a rdf:Statement ;
    rdf:subject :DonationOfConstantine ;
    rdf:predicate :date ;
    rdf:object :8century ;
    :claimedBy :LorenzoValla ;
    :source :DeFalsoCredita ;
    :classifiedAs :acceptedOpinion .

# Pope Leo IX disputed opinion on the Donation
:Claim2 a rdf:Statement ;
    rdf:subject :DonationOfConstantine ;
    rdf:predicate :date ;
    rdf:object :4century ;
    :claimedBy :LeoIX ;
    :classifiedAs :disputedOpinion .
```

Listing 2.2: Representation of Valla's and Pope Leo IX's critical inquiry with standard reification

This representation involves using a custom ontology (also referred to as an ontology-dependent solution) to represent disputed claims (e.g., `:classifiedAs`).

Alternatively, the RDF standard's reification offers a natural way to express claims and indicate their current acceptance or rejection within the community. **Expressing Without Asserting (EWA)** refers to the capacity of certain reification methods to represent claims or statements without implying their absolute truth or acceptance. EWA allows the documentation of both perspectives without the inclusion of ontology-dependent solutions. In doing so, historical views are preserved, while only the settled opinion, such as Valla's 8th-century dating, is acknowledged as valid in the current representation. Pope Leo IX's superseded claim remains included for historical context but is marked as non-asserted, as illustrated in Listing 2.3. An additional triple `:DonationofConstantine :date :8century` is added to assert the settled opinion (Valla's claim), while the other remains non-asserted (Pope Leo IX's claim).

```
# Valla’s opinion on the Donation
:Claim1 a rdf:Statement ;
    rdf:subject :DonationOfConstantine ;
    rdf:predicate :date ;
    rdf:object :8century ;
    :claimedBy :LorenzoValla .

# Pope Leo IX opinion on the Donation
:Claim2 a rdf:Statement ;
    rdf:subject :DonationOfConstantine ;
    rdf:predicate :date ;
    rdf:object :4century ;
    :claimedBy :LeoIX .

# Currently accepted opinion
:DonationofConstantine :date :8century
```

Listing 2.3: RDF Representation of concurrent datings of the "Donation of Constantine"

On the other hand, non-assertion is represented in Listing 2.3 by representing the superseded dating of the Donation via the reified statement (`:Claim2`), but without representing any asserted triple (which means that `:DonationOfConstantine :date :4century` does not exists).

A SPARQL query retrieves only asserted contents by default. For instance, the question *when has the "Donation of Constantine" been produced?* performed via SPARQL query shown in Listing 2.4, theoretically retrieves the value 8th century (`:8century`).

```
SELECT ?date

WHERE {:DonationOfConstanine :date ?date}
```

Listing 2.4: Protypical SPARQL query to retrieve settled (asserted) dates attributed of the Donation of Constantine

A SPARQL query retrieves all concurring opinions if the statement pattern is searched. For instance, the question *when has the "Donation of Constantine" thought to been produced through history?* performed via SPARQL query shown in Listing 2.5, can be theoretically retrieves both 8th century (`:8century`) and 4th century (`:4century`) values.

```
SELECT ?date

WHERE {
    ?claim a rdf:Statement .
        rdf:subject :DonationOfConstantine ;
        rdf:predicate rdf:predicate :date ;
        rdf:object ?date .
    }
```

Listing 2.5: Protypical SPARQL query to retrieve all dates attributed to the "Donation of Constantine" through history

A SPARQL query retrieves the disputed statements by negating the existence of a triple asserting its contents. For instance, the question *when was the "Donation of Constantine" thought to have been produced?* shown in Listing 2.6, theoretically retrieves the value 4th century (`:4century`).

```
SELECT ?date

WHERE {
    ?claim a rdf:Statement .
        rdf:subject :DonationOfConstantine ;
```

```
        rdf:predicate rdf:predicate :date ;
        rdf:object ?date .

    FILTER NOT EXISTS {:DonationOfConstantine :date ?date}
    }

```

Listing 2.6: Protypical SPARQL query to retrieve all disputed (now discarded) dates attributed of the "Donation of Constantine"

This structure avoids reasserting superseded opinions while preserving the historical context of the evolution of critical inquiry over time.

The study of truth values and uncertainty of web data has been addressed in several works. Knowledge representation has been identified as a key issue that affects the creation and consumption of data whose truth value is not explicit and shared and eventually prevents interoperability [10]. Precisely, representation affects all related tasks, including knowledge extraction, uncertainty measurement, and services development. This aspect has been designated as the primary focus of the present work and will be further analysed and discussed in Chapter 4.2.
Two approaches to coping with such claims in the Semantic Web can be distinguished: ontology-dependent and ontology-independent solutions.

Ontology-dependent solutions include OWL ontologies describing provenance, uncertainty, and trust-related aspects. CH KGs and data models managing interpretations and observations mainly adopt n-ary relations [75] to annotate such statements and seem to tackle the possibly concurring claims slightly and, therefore, their logical status besides the knowledge representation.

At a broader scope, the ontology proposed by the W3C Incubator Group URW3-XG [62], and a recent proposal of *mUnc* [33], allows one to describe uncertainty information according to several, non concurring, theories. However, ontologies do not prevent limitations derived from the actual assertion of claims, e.g. in settled disagreements, where superseded statements remain asserted despite having been rejected, therefore requiring human intervention to understand incompatible, coexisting, contradictory claims. Many communities have developed their own ontologies to cope with domain-related issues, e.g. to characterise multisensor data [20], reliability in the heritage science [74], historical evolution in art history[28]). Each domain relies on its definition of uncertainty, and consumers must learn different expressions for the same problem.

Ontology-independent solutions have been proposed to prevent the adoption of multiple vocabularies. These span from the usage of reification [56] to singleton properties [72], to the definition of named graphs [15], models for organising named graphs [51], serialisations like Notation3 [11], and extensions to the graph model like RDF-star [54].

These solutions have been critically analysed [105, 59, 44, 100, 95, 78]. However, existing analyses are not able to identify a one-size-fits-all effective solution to represent truth value of graphs [59, 100, 95], since often the focus is on the representation of provenance only, and they limit the evaluation of proposals to only their efficiency in graph stores [44, 59, 78].

While logical proofs and theoretical validation can demonstrate a method's formal effectiveness, assessing its efficiency requires empirical evaluation. Introducing *statements about statements* often results in a decline in data retrieval performance, such as slower query execution times and increased data upload overhead in triplestores. The choice of reification method is thus crucial, as it significantly influences these performance-related aspects. In particular, such tests are run over some empirical performance indicators like the *number of triples*, *query execution time*, *query complexity*, *dataset storage*

[28] See for instance RKD collection or Linked Art application profile for former attributions as described in the previous section (Section 2.1

consumption, *support by existing tools and implementations* and other pertinent metrics. There is a vast number of academic papers that suggest this kind of experimental setup to benchmark the aforementioned reification methods [59, 72, 78, 44, 60, 29, 104]. Additionally, all the benchmarks mentioned above rely on the same four components, namely *datasets, queries, triplestores* and *reification methods.*

2.3 Discussion

This chapter explored some relevant aspects concerning RQ1 - *What are the current limitations and potential improvements in the representation of humanities critical inquiry in current Knowledge Graphs?* Some gaps in the representation of critical inquiry have been addressed.

As introduced in the problem statement (Section 1.3), many existing KGs [34], especially online catalogues, fail to represent structured contents on critical inquiry, scholarly investigation and their evolution. This limitation restricts the ability of KGs to reflect the complexities of scholarly debates and the fluid nature of historical narratives, which often lead to incompatible views and beliefs on the same subject.

Several CH KGs and underlying data models consider some aspects of the hermeneutic nature [45] and the intrinsic historiographical uncertainty [86] of humanities discourse in their design. In summary, they adopt n-ary relations [77, 34] and named graphs [24, 66] to reify their triples and express scholarly knowledge and related contextual metadata.

While frameworks such as CIDOC CRM [77] provide valuable conceptual foundations and critical inquiry representation [35], they lack the specific guidelines necessary for addressing incompatible or conflicting interpretations which are always asserted. This shortfall highlights the opportunity to refine existing ontologies and their documentation, including EWA, to represent the evolution of scholarly knowledge in their methodologies.

These aspects are limited to distinctions between data and capta, the scholarly methods applied, the sources under study, the confidence in claims, the representation of complex capta, and the identification of settled inquiries as outlined in Section 2.1.

The complexities surrounding statements logical statuses and their implications for knowledge extraction remain inadequately addressed in many RDF-based solutions. EWA does not seem to be addressed as a design choice in many Digital Humanities projects, which limits the ability to represent and retrieve complex scholarly claims and their uncertainties.

Chapter 3

Weaker logical status claims in Wikidata cultural heritage records

The research presented in this chapter is based on a study published in [30]. It is reported with the consent of all authors.

This chapter mainly addresses RQ1 - *What are the current limitations and potential improvements in the representation of humanities critical inquiry in current Knowledge Graphs?* - by defining three sub-questions to address WLS claims in Wikidata specifically.

- RQ1a - *How widespread are WLS claims in the current state of Wikidata?*

- RQ1b - *How does the cultural domain of the Wikidata topics (and, presumably, of the individuals contributing to the data regarding the Wikidata topics) affect and reflect on the relative success and richness of some approaches over others?*

- RQ1c - *Does the actual usage of the surveyed approaches match their designed use declared by Wikidata?*

3.1 Usage patterns of WLS in Wikidata datasets

Wikidata represents WLS information using at least three different representation methods: ranked statements (Section 3.1.1), statements with specific qualifiers (Section 3.1.2), and statements with a non-existing valued object (Section 3.1.3).

Although missing values are outside the primary focus of this analysis (e.g., the author of the painting is unknown), they have been included since they represent entirely uncertain claims (unknown). Such uncertainty may require contextual information or coexist with different concurring opinions.

Several other patterns may reveal questionable knowledge, scholarly opinions and critical debate in Wikidata CH records. However, a complete overview of all ways to express capta in Wikidata CH records will be far too expensive for two reasons.

- The high number of qualifiers (356) and their associated values about general knowledge. For instance, 1048 values are available to describe the *determination method* of a claim[1]. Additionally, annotators may even choose terms which do not belong to these values, increasing the number of actual results in the data.
- More than 3 million CH records are sparsely represented among more than 4000 classes[2].

3.1.1 Ranked statements

Ranking of assertions is modelled by the Wikibase data model[3] to express different degrees of the preferability of individual claims.

Claims in Wikidata are expressed through *statements*, using n-ary relations and customised prefixes modelled via the Wikibase data model[4] to

[1]All available values associated with determination method are available with this SPAQL query `https://w.wiki/BWpF`

[2]This value is calculated over instances belonging to all-recursive subClasses of work of art `Q838948` `https://w.wiki/BWpJ`

[3]`https://www.mediawiki.org/wiki/Wikibase/DataModel#Statements`

[4]`http://www.wikidata.org/entity/Help:Statements`

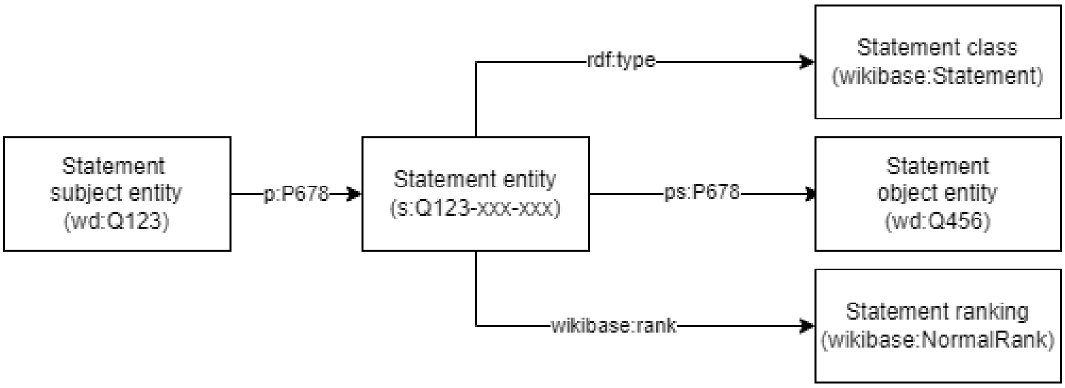

Figure 3.1: Graphical representation of basic components of an RDF statement in Wikidata data model

express contextual information (e.g. qualifiers, rankings, references) about it. As shown in Figure 3.1, a statement is formed by a claim's subject (e.g., `wd:Q123`) and a statement entity (`s:Q123-xxx-xxx`, instance of the class `wikibase:Statement`) via the claim's predicate (e.g., `p:P678`). The statement entity then refers to the claim's object (e.g., `wd:Q456`) via the same predicate but with a different prefix (`ps:P678`). Contextual information can also be associated with the statement entity, such as the claim's ranking (e.g., `wikibase:rank wikibase:NormalRank`)[5].

Statements do not assert the corresponding claim. To do so, another triple must be added that (using a different prefix) flatly relates the Statement's subject to the Statement's intended object through the Statement's predicate, thus enabling simple query support for asserted facts. The separation between Statements and their assertion is selectively provided, which allows to easily support both claims presented as facts (where both the Statement and the assertion triple exist) and claims not meant to be considered facts (the Statement exists, but no assertion triple is added).

The ranking mechanism is enriched with the representation of asserted and non-asserted statements. Rankings [57] communicate the consensus opinion for a statement as reached by the scientific community or Wikidata annotators. Disputes are separately hosted on the corresponding discussion page, in plain

[5]The identifiers in the example are fictitious and provided solely to illustrate the structure of Wikidata entities and predicates for clarity.

text. Many possible combinations of variously ranked competing statements can be found in the Wikidata collection, with various and debatable interpretations. Ranking is assigned to individual Statements using values such as *Preferred*, *Normal* and *Deprecated*).

Whether or not a statement is asserted is determined solely by the statement's rank and the absence of higher-ranked statements using the same predicate – the Wikidata engine automatically provides it and is not a conscious choice of the editors.

Normal statements

The Normal ranking is the default ranking for Statements. A Statement ranked Normal can be either asserted or not depending on the existence and intended meaning of competing Statements that are placed against it. For instance, in Listing 3.1, "The Scream" by Edvard Munch belongs to the Expressionist period[6], and this is expressed as an asserted Normal Statement, to signify that the annotator does not introduce any WLS on the claim. In Listing 3.3, on the other hand, the first Statement (lines 1-5) is ranked Normal but not asserted since the Preferred statement is present and asserted instead.

```
# "The scream" belongs to the Expressionist movement
wd:Q471379 wdt:P135 wd:Q80113 .
wd:Q471379 p:P135 s:Q471379-c3e5c17d-4730-a5dc-85cb-efc9766b7c80 .
s:Q471379-c3e5c17d-4730-a5dc-85cb-efc9766b7c80 a wikibase:Statement,
  wikibase:rank wikibase:NormalRank ;
  ps:P135 wd:Q80113 .
```

Listing 3.1: Normal rank

Deprecated statements

Deprecated statements are meant for questioned and discarded claims or do not represent a correct value in the editors' view. Deprecated statements are always automatically non-asserted independently of the ranking of the other concurring statements. For example, Listing 3.2 expresses the concept that

[6] http://www.wikidata.org/entity/Q471379

"The Lamentation"[7], a print by Albrecht Dürer, was reported to be created in 1504. The Deprecated rank and the lack of an assertion triple indicates that this date is not thought to be valid.

```
# creation date thought to be 1504
wd:Q18338462 p:P571 s:Q18338462-FDDCD91B-3919-450A-B00D-FE3ADA773A11 .
s:Q18338462-FDDCD91B-3919-450A-B00D-FE3ADA773A11 a wikibase:Statement ;
    wikibase:rank wikibase:DeprecatedRank ;
    ps:P571 wdt:P571 "1504-01-01T00:00:00Z"^^xsd:dateTime .
```

Listing 3.2: Deprecated rank

Preferred statements

Preferred statements are meant for claims with a stronger status and representing the currently presumed correct value of a predicate. They are always also asserted. For instance, as shown in Listing 3.3, a retracted attribution of the painting "Madonna with the Blue Diadem" [8] to Raphael is represented only by a Statement ranked as Normal and no assertion triple, while the attribution to Gianfrancesco Penni enjoys both a Preferred rank and the assertion triple.

Even though the first attribution is ranked Normal rather than Deprecated, it must considered a superseded claim. This example shows that the nature of Normal statements varies depending on whether they coexist or not with competing Preferred and/or Deprecated claims, and similarly may vary the presence or absence of assertion triples.

```
# attribution to Raphael (Q5597)
wd:Q738038 p:P170 s:q738038-121B92D0-E6E1-4514-960C-AE34F50054E5 .
s:q738038-121B92D0-E6E1-4514-960C-AE34F50054E5 a wikibase:Statement ;
    wikibase:rank wikibase:NormalRank ;
    ps:P170 wd:Q5597 .

# attribution to Gianfrancesco Penni (Q2327761)
```

[7] http://www.wikidata.org/entity/Q18338462

[8] http://www.wikidata.org/entity/Q738038

```
wd:Q738038 wdt:P170 wd:Q2327761 .   # assertion
wd:Q738038 p:P170 s:Q738038-7729b786-4d4f-a0ca-2ded-4ea2c6307e1c .
s:Q738038-7729b786-4d4f-a0ca-2ded-4ea2c6307e1c a wikibase:Statement;
    wikibase:rank wikibase:PreferredRank ;
    ps:P170 wd:Q2327761.
```

Listing 3.3: Preferred and Normal ranks

3.1.2 Qualifiers

Statements, independently of rank, can be decorated with additional triples annotating contextual information or specifications about the claim itself[9]. Those annotations may be *additive* when they provide additional information about the fact (e.g., to specify the character played by an actor when listing him or her as a cast member of a movie) or *contextual* when they limit the contexts in which the underlying fact is true (e.g., the claims is a hypothesis) [83].

Following the example from [2], an analysis of the 150 most frequently used qualifiers in Wikidata and their associated values reveals that `P1480`:*sourcing circumstances*[10] (47th most used one) and `P5102`:*nature of statement*[11] (134th most used one) are the most commonly used qualifiers to introduce WLS on claims.

Additionally, the Wikidata model provides the properties `P2241`:*reason for deprecated rank*[12] (42th most used qualifier) and `P7451`:*reason for preferred rank*[13] (114th most used qualifier) to annotate contextual information about superseded and preferred claims, respectively.

For example, in Listing 3.4, the painting "Abstract Speed + Sound"[14] by Giacomo Balla is described as *possibly* part of a triptych. The use of a qualifier

[9] The complete list of available qualifiers in Wikidata is available at `https://w.wiki/6TrP`

[10] The most frequently used values are: *circa, presumably, allegedly, inference, uncertainty, possibly, near, probably, conventional date, disputed*

[11] The most frequently used values are: *originally, attribution, hypothesis, often, allegedly, expected, possibly, disputed, rarely, mainly*

[12] `http://www.wikidata.org/entity/Property_talk:P2241`

[13] `http://www.wikidata.org/entity/Property_talk:P7451`

[14] `http://www.wikidata.org/entity/Q19882431`

with a Normal ranking seems to imply that the statement is considered true and therefore it is also asserted.

```
wd:Q19882431 wdt:P361 wd:Q79218 .      # part of: triptych (assertion)
wd:Q19882431 p:P361 s:Q19882431-1ac26ff2-4981-ff79-4fae-9d411ae34296 .
s:Q19882431-1ac26ff2-4981-ff79-4fae-9d411ae34296 a wikibase:Statement;
    wikibase:rank wikibase:NormalRank ;
    ps:P361 wd:Q79218 ;                # part of: triptych
    pq:P5102 wd:Q30230067 .            # circumstance: possibly
```

Listing 3.4: A qualified statement in Wikidata

Wikidata provides a list of 96 recommended values for *nature of statement* and 83 recommended values for *sourcing circumstances* in their respective *Property Talk* pages, while no list of recommended terms is provided for *reason for deprecated rank* nor *reason for preferred rank*. However, terms that were used with these properties can be retrieved via a simple SPARQL query[15], showing respectively 384 and 83 distinct terms. Even at first glance, it is possible to notice a very wide range of types and specificities (e.g., qualifiers such as *possibly*, *presumably*, and *probably* versus, say, *prosopographical phantom*, *project management estimation* or *archive footage*), and many are not connected to WLS assessments. In addition, semantic overlaps can be noticed on many of these terms, e.g. between *allegation* and *allegedly*, or between *hypothesis*, *hypothetical entity*, *hypothetically* and *scientific hypothesis*. These overlaps support arbitrariness of choice for contributors, increasing the ambiguity of the resulting annotation.

3.1.3 Missing values

There are three types of basic information structures used to describe Entities in Wikibase (called SNAKs, or *Some Notation about Knowledge* [16] in Wikidata: actual values (URIs or literals), `someValue` placeholders and `noValue` placeholders. They are used to represent that the statement is associated with an

[15] List of terms used in Wikidata with *reason for deprecated rank* *https://w.wiki/6Tpt* and with *reason for preferred rank* https://w.wiki/7VGf

[16] https://wikidata.github.io/Wikidata-Toolkit/org/wikidata/wdtk/datamodel/interfaces/Snak.html

unknown value (mapped as `someValue`) or with a non-existing value (mapped as `noValue`), which is a more precise assessment than simply not recording the statement at all. The use of the same syntactic tool is known to generate precision and correctness issues (e.g., see [58]), since the RDF standard specifically defines blank nodes with an existential semantics while SPARQL does not follow such semantics. As a result, SPARQL queries run over datasets where blank nodes are used as existentials to represent unknown values, the results can be unintuitive or arguably incorrect. Although the RDF representation of Wikidata uses blank nodes for both unknown and non-existing values, a skolemization process separates them conceptually and with a simple filtering query it is possible to distinguish them[17].

Unknown values

Unknown valued statements are claims whose object exists but is not known[18]. For instance, in "The Book of Lismore"[19] there is an unknown value for the `P195`:*collection* property, which is a positive statement that the information existed but it has not been preserved. As mentioned, in the RDF representation unknown values are represented via blank nodes as shown in Listing 3.5.

```
# unkown collection
wd:Q1371647 wdt:P195 _:15518d67963a082b352304a1ab8e016e.
wd:Q1371647 p:P195 s:Q1371647-B07F6386-A7D0-4C9D-8E77-CC2BD523354E .
s:Q1371647-B07F6386-A7D0-4C9D-8E77-CC2BD523354E ps:P195 _:0088
    bc50e53b3902bea74cc2380cbd09 ;
    pq:P3831 wd:Q768717 . # role: private collation
```

Listing 3.5: Unknown-valued statement in Wikidata

[17] https://www.mediawiki.org/wiki/Wikibase/DataModel#PropertysomeValueSnak
[18] https://www.wikidata.org/wiki/Help:Statements#Unknown_or_no_values
[19] https://www.wikidata.org/wiki/Q1371647

Non-existing values

Non-existing valued statements[20] are claims whose object is not existent (or not available in Wikidata). For example, the pilot episode of X-files[21] has a non-existing value for the *follows* (`P155`) property, considering that the pilot starts the series. Non-existing values are not central in the scope of this work, but they are listed for two main reasons: First, unknown values are inherently uncertain. Therefore, concurring hypotheses (e.g., marked with rankings) or contextualising qualifiers (e.g., nature of statement) can provide context for them. Secondly, in practice, there is some overlap between unknown valued and non-existing valued claims. For example, the "Missal for the use of the ecclesiastics of Clermont’[22], an illuminated manuscript from the 14th century, has been recorded with both a non-existing creator and author, as shown in Listing 3.6. The example is incorrect as it should use an unknown value, and this leads to confusion about the usage of missing values, further contributing to complications.

```
 # author: unknown (blank node - assertion)
wd:Q113302686 wdt:P50 _:4c60f23d697d2d89d9fe49824c8f3a01 .
wd:Q113302686 p:P50 s:Q113302686-032e3cc5-4fd6-1f20-8830-0909945ba683 .
s:Q113302686-032e3cc5-4fd6-1f20-8830-0909945ba683 a wikibase:Statement;
    wikibase:rank wikibase:NormalRank ;
    ps:P50 _:f8c6b698b13ef3dd3738e025df3a2d5d . # author: unknown

# creator: unknown (blank node - assertion)
wd:Q113302686 wdt:P170 _:759d5c5c7a58a8a286512c257514463a .
wd:Q113302686 p:P170 s:Q113302686-8d47e883-4566-bc8b-cd8f-6cffebc5414c .
s:Q113302686-8d47e883-4566-bc8b-cd8f-6cffebc5414c a wikibase:Statement;
    wikibase:rank wikibase:NormalRank ;
    ps:P170 _:28d04a432a3589d30a5c6da79d3fac50 . # creator: unknown
```

Listing 3.6: Non-existing valued statement in Wikidata

Even before checking on the actual usage patterns of these methods, the richness of annotations made possible by them can be immediately noticed, as

[20] `https://www.wikidata.org/wiki/Help:Statements`
[21] `http://www.wikidata.org/entity/Q7194381`
[22] `http://www.wikidata.org/entity/Q113302686`

the subtle nuances they afford and the variety of (potential) sources of ambiguities, overlapping connotations and representation vagueness. In particular, three specific problems that are worth further discussion:

1. Although the separate uses of Normal, Preferred and Deprecated rankings are clear and practical, there are uncertainties when they coexist on the same predicate, especially for the different representations of Normal statements when Preferred ones are also present or when all three rankings are present.

2. The sheer number of qualifiers, the differing levels of their respective specificities, and the manifest semantic overlapping of many of them make it quite hard to guarantee homogeneity and precision in their use. The use of contextualising qualifiers, be they temporal, provenance or otherwise, does not add to the base information, but changes the context within which such information is true. As [83] suggests, contextual qualifiers should not be shown to consumers. Still, basic tools (visualisers, contextualisers, reasoners) should be written to take such context into account correctly or low-level tools should remove facts that are not valid in the selected contexts.

3. The subtlety in the semantic differences between providing no statement, specifying a `noValue` blank node and providing a `someValue` blank node for a property of a Wikidata item, as well as their other types of applications makes the use of missing values complicated and ambiguous.

In a way, WLS claims can be seen simply as logical disjunctions of competing claims each of which is separately annotated with context, provenance, confidence, temporal boundaries, etc.: "according to α, $\texttt{s p o}_1$" and "according to β, $\texttt{s p o}_2$" can be seen as "$[\texttt{s p o}_1]_\alpha \vee [\texttt{s p o}_2]_\beta$" with some added annotations connecting the first branch to α and the second to β (e.g., through reification, named graph, or blank nodes). This approach has limitations both from the practical and the conceptual point of view. Practically, RDF has no real way to express disjunctions without some additional baggage to encode predicate calculus employing the systematic use of reification [68]. Conceptually, focusing on the inner statements to the exclusion of the contextualising information may miss the point that in many scholarly domains, it is not the full list of competing claims to be of interest but the very existence of the debate in the first place. Disjunctions would not help here.

Another way to formally understand WLS claims is to link them to modal statements in modal logic [48], which can be used to understand the coexistence of strong logical status claims, expressed as atomic formulas $p(s,o)$, and weak logical status ones, expressed as modal formulas $K_\alpha p(s,o)$ or $B_\beta p(s,o)$, where K_α and B_β are modal operators guided by specific modal axioms[23]. Various types of modal logics exist and have been used to introduce different operators and represent different semantics, such as possibility and necessity (the *strictu sensu* modal logic), or obligation and permission (*deontic logic*), or temporally bounded predicates (*temporal logic*), or belief (*doxastic logic*) or knowledge (*epistemic logic*). Overall, they form a complete formal mechanism to study the characteristics and principles of WLS claims that does not imply the need to proceed to a reconciliation of different world views.

However, these reflections remain empty and pointless without examining how contributors apply these methods to express actual WLS claims in their Wikidata contributions.

The representation of complex data scenarios in knowledge bases often needs to be evaluated according to multiple metrics. For instance, Piscopo and Simperl [89] survey quality metrics from 28 scientific publications on the topic and categorize quality assessments into three dimensions: intrinsic (accuracy, trustworthiness, consistency), context (relevance, completeness and timeliness) and representation (ease of understanding and interoperability). Among quality measures, evaluation of completeness, defined in [41] as the "presence of all required information in a given dataset", has been approached through various

[23]e.g., $\mathbf{T}$ ($K_\alpha\phi \rightarrow \phi$) for epistemic logic or $\mathbf{N}$ ($\vdash \phi \implies \vdash B_\alpha\phi$) for doxastic logic.

methods and assessments as comparing data for similar entities [7], measuring entity relatedness [90], evaluating thoroughness of information by determining the completeness of specific attributes of objects [46], assessing low-quality statements thought the analysis of items' discussion pages, deprecated statements and constraint violations [99], and assessing and comparing data quality across large knowledge bases [41, 1].

Overall, little or no evaluation has been conducted on the representation of weaker logical status claims to represent critical inquiry evolution in Wikidata, nor has a comprehensive analysis been carried out to assess the amount of knowledge related to WLS status in CH. The next chapter (Chapter 3) addresses these shortcomings, particularly by surveying the available representation patterns and their actual usage of WLS claims to express critical inquiry and the evolution of knowledge in CH items in Wikidata.

3.2 Data Acquisition

```
SELECT DISTINCT ?artwork ?type
WHERE {
    ?artwork wdt:P31 ?type.
    ?type (wdt:P279*) wd:Q838948.
    hint:Prior hint:rangeSafe true
    }
```

Listing 3.7: SPARQL query retrieving Wikidata entities to subclasses of work of art (Q838948)

The first dataset contains CH items, a complete snapshot of the Wikidata records of these cultural assets[24]. All Wikidata entities belonging to the class *work of art* [25]) or any of its sub-classes were collected using a SPARQL query (Listing 3.7).

Work of art has been chosen as the top-level class for retrieving entities within the Cultural Heritage (CH) domain. This selection is non-trivial, as

[24] It is assumed that discipline-oriented datasets give access to domain-specific annotation habits of scholars better than a sample of random entities

[25] http://www.wikidata.org/entity/Q838948

broader categories exist, such as cultural heritage (Q210272[26]), which includes 50,1987 entities[27]. However, this class was avoided due to the deprecation of its superclass relationship with work of art (Q838948). Instead, work of art was considered sufficiently general, representing a wide range of CH records.

The statements for all selected entities were downloaded in JSON format[28]. Data is stored in numerous JSON files, and each contains a complete representation of at most 50 Wikidata entities with their labels, descriptions and statements. This CH dataset has been semi-automatically divided into three sub-datasets due to the wide diversity of cultural properties and their associated claims:

- *Audio-Visual heritage* (CHav): This collection holds information about audio-visual materials that have cultural, historical, or artistic value. They include movies, videos, recordings of music or spoken words, and other audio-visual materials that record a particular event in a specific time or place. The dataset contains 1,251,626 entities and 17,141,394 statements organised in 25,033 JSON files.

- *Visual heritage* (CHv): This collection holds information about visual artefacts with cultural, historical, or artistic value. They include paintings, drawings, sculptures, photographs, decorative arts, etc. The dataset contains 1,078,855 entities and 12,850,825 statements organised in 21,579 JSON files.

- *Textual heritage* (CHt): This collection holds information about written and printed materials with historical or cultural significance. They include books, manuscripts, letters, and other written documents. The dataset contains 625,110 entities and 4,584,444 statements organised in 12,503 JSON files

[26] https://www.wikidata.org/wiki/Q210272

[27] The query for counting cultural heritage entities in Wikidata is available at: https://w.wiki/D2wZ

[28] via http://www.wikidata.org/wiki/Wikidata:Data_access

Wikidata entities of architecture-related classes were later discarded due to their relatively low number and the presence of many statistical ambiguities that could make their evaluation useless (e.g., many entities belonging to these classes should not be considered relevant to cultural heritage collections).

The second dataset, chosen to verify our assumptions using a different collection with a similar size, is a collection of astronomical entities organised into two datasets:

- *Stars* (ANs): This collection holds a random selection of 1,199,950 Wikidata entities (of the ~3.3 million existing) belonging to the class *Star*[29], The dataset contains 27,470,140 statements in 23,999 JSON files[30].

- *Galaxies* (ANg): This collection holds a random selection of 1,200,000 Wikidata entities (of the ~2 million existing) belonging to the class *Galaxy*[31], The dataset contains 14,439,421 statements in 24,000 JSON files.

The number of astronomical entities has been limited to 1,200,000 to approximately balance them to each other (although the CHt is about half in size with 625,110 entities), as well as the average number of statements for each entity (CHav: 13.7, CHv: 11.9, CHt: 7.3, ANs: 22.9, ANg: 12).

Both CH and AN datasets use multiple fuzzy assertions and hypotheses and, therefore, require assertions with weaker status (e.g., attribution uncertainties or physical locations moving over time for paintings vs. spectral class or radial velocity for stars).

As already introduced in the theoretical framework (Section 1.1), the decision to use a comparative dataset in this study is motivated by exploring the similarities and differences between astronomical and humanities academic

[29] `http://www.wikidata.org/entity/Q523`

[30] the ANs dataset was meant to be composed of 24,000 files with 50 entities each, but after running our tests it has been noticed that a file was corrupt. It has been chosen to discard that contribution.

[31] `http://www.wikidata.org/entity/Q318`

practices. This comparison mainly addresses the first hypothesis of this Dissertation (H2), pointing out that annotators in CH may approach data observations, incompleteness and uncertainty in different ways, favouring qualitative, context-rich representations of competing hypotheses. Three datasets of Wikidata items have been collected to generate some analysis about the actual usage of patterns to represent WLS claims and to provide an initial answer to the research questions: one about CH items (visual arts, text documents and audio-visual entities), another about Astronomical objects (galaxies and stars) and one with a selection of random entities reflecting the actual distribution of entities in classes in the whole Wikidata as discussed in the introduction (Section 1.1). The datasets were selected to be approximately comparable in size, the number of individual statements, and evidence that many types of entities rely on weaker logical status claims when they undergo re-evaluations due to new evidence or the recording of different opinions.

The third dataset is a selection of randomly chosen entities from Wikidata. This dataset was acquired to compare WLS claims in the other datasets with a randomised subset designed to mimic the overall distribution of WLS claims in Wikidata.

- *Random* (R): This dataset comprises 1,159,800 Wikidata entities (starting from a selection of 1.2 million entities from which duplicates have been removed) chosen randomly from the most numerous 100 classes to reflect the proportional distribution of entities found in Wikidata[32]. This dataset encompasses 61,798,072 statements distributed across 23,196 JSON files.

Table 3.1 summarises basic information about these collections. All these datasets can be accessed and downloaded from Zenodo[33] [82] and all Python scripts are accessible in GitHub[34].

[32] https://w.wiki/7iCR
[33] https://doi.org/10.5281/zenodo.7624783
[34] https://github.com/alessiodipasquale/Wikidata_WLS

3.3 Analysis

In the following paragraphs, WLS statements are described as all Wikidata statements showing the use of each approach described in Section 3.1, regardless of whether they have been used to make weaker logical status claims. Table 3.1 shows a tabular presentation of our analysis.

Even though critical analysis is a pivotal part of humanities discourses, plainly stated statements with no competing claims are largely the most represented information in the CH dataset: the vast majority of statements here (>99%, in particular 99.74% in CHav, 99.92% in CHv and 99.69% in CHt) are plainly asserted statements with no WLS additions. In contrast, the Astronomical datasets show a reasonably different situation, 83% overall of plainly asserted statements, specifically ANs at 72.58% and ANg at 95%. The overall distribution of the Random (R) dataset showcases a low percentage of WLS claims (1.78%), closer to the CH and the AN datasets. Yet, interestingly, almost the whole percentage is made of non-asserted statements (98.95%) matching a similar distribution in the AN dataset.

When analysing the Random (R) dataset, it can be noticed that the ranking system's simplicity leads to a clear predominance of deprecated items and, consequently, of non-asserted claims. The other approaches appear to be underutilised in a proportion closer to the AN dataset. Possibly, this is a reflection that, in the CH community, historical uncertainty and the representation of interpretation are more frequent and typical than in other disciplines.

To further explore these data, it can be observed that:

Non-asserted statements: Of the approaches previously listed (cf. Section 3.1), non-asserted statements (i.e., variously ranked statements with no corresponding asserted triples) are largely the most frequent approach for representing competing information in both AN and R. The situation is fairly different in the CH collections, non-asserted statements being the most frequently used approach in CHt (81.64%) and CHav (only 86.09%) and almost

	Cultural Heritage			Astronomy		
	Audio-visual (*CHav*)	Visual (*CHv*))	Textual (*CHt*)	Stars (*ANs*)	Galaxies (*ANg*)	Random (R)
Entities	1,251,626	1,078,855	625,110	1,199,950	1,200,000	1,159,800
Statements	17,141,394	12,850,825	4,584,444	27,470,140	14,439,421	61,798,072
Weaker Logical Status (*WLS*)	**50,193**	**227,218**	**17,216**	**7,532,169**	**721,504**	**1,101,014**
% WLS / Statements	**0.29%**	**1.77%**	**0.37%**	**27.42%**	**5.00%**	**1.78%**
Non-asserted statements	43,211	9,056	14,055	7,532,107	721,503	1,089,469
Ranked as Deprecated	7,622	3,057	1,568	2,768,829	189,691	721,870
Deprecated with a reason	4,949	769	715	2	0	8,993
Non-existing values	50,611	1,969	1,356	4	0	3,857
Unknown value	4,896	106,521	1,843	0	0	5,139
Qualified statements	2,406	114,674	1,556	532	1	7,716
WLS qualified statements	2,086	111,641	1,318	62	1	6,406
WLS qualifiers w/o *circa*	719	3,988	330	35	0	1,724

Table 3.1: Entities, statements and types of WLS statements

unused in CHv (3.99%).

Deprecated statements: Deprecated claims are visibly a small portion of the overall non-asserted statements, occurring only in 20% of the non-asserted statements of the CH entities, in 30% of the non-asserted statements of Astronomical entities and in the 66% of the non-asserted statements of Random entities. At the same time, about half of the deprecated statements were annotated with the corresponding *reason for deprecated rank* qualifier (in particular, 45.59% CHt, 25.15% CHv, 64.93% CHav – compare this with basically 0% in both AN datasets and 1.24% in R dataset), proving that scholars in the humanities have a solid interest in annotating provenance of WLS claims on CH data. Yet, only less than 1% of preferred statements have been annotated with the corresponding qualifier *reason for preferred rank*.

Unknown values: Unknown valued statements are not used at all in Astronomical data (absolute 0 in both ANg and ANs out), poorly adopted in the R dataset (0.47%), and sparsely used in the Humanities as well (9.75% in CHav and 10.71% in CHt). Higher is the result for the CHv dataset, with 46.88% of the overall WLS claims using this approach.

Non-existing values: Even if they do not represent WLS claims, they were examined in our datasets for contiguity to unknown values. Non-existing values are almost unused in Astronomical data (exactly 4 occurrences in ANs and an absolute 0 in ANg out of more than 7 million WLS claims) and very sparsely used in the Humanities and Random datasets as well: 1.969 statements in CHv, 1.356 statements in CHt and 3.857 statements in R dataset. Fairly higher is the result for the CHav dataset, with 50,611 statements using this approach. This outlier value is probably justified and will be commented on later in this section.

Qualifiers: Statements qualified with *nature of statement* and *sourcing circumstances* predicates are the least employed out of the surveyed ones, being used in 7.66% of the WLS statements in CHt, in 0.58% of the WLS statements in R and in 4.16% of the CHav statements, present in 0.0008% of the ANs statements and only in one ANg statement. Yet, they are used in 49.13% of the WLS statements of the CHv dataset. This value will be commented later on in this section.

Terms used as values for the qualifiers have been surveyed. The use of respectively 200 different values for qualifier *nature of statement* has been witnessed, 419 for *sourcing circumstances* and 588 for *reason for deprecated rank*. These values largely exceed the proposed values specified in the corresponding Wikidata property talk pages (respectively, 194 values for *nature of statement* and 175 for *sourcing circumstances*) or property constraints as for the 384 values for *reason for deprecated rank*). Furthermore, the three sets of actual terms show a considerable overlap of values between them (in our datasets, but also over all of Wikidata), as shown in Figure 3.2. This seems to imply that the semantics associated with these values, and indeed the properties themselves, may have been unclear to contributors, who sometimes selected the qualifier in non-predictable ways. Therefore, all three sets into a single category have been grouped (shown as *WLS qualified statements* in Table 3.1).

Given that the R dataset is not disciplinary, it was determined that various situations occurring across disciplinary boundaries would likely compromise any analysis beyond simple counting. Therefore, the following section will focus only on the disciplinary datasets.

The terms used as values for the qualifiers have been further surveyed.

Overall, the three sets contain a variety of terms such as generic contextual information items, e.g., provenance details, as well as domain-specific terms not relevant to our purposes (e.g., *show election*, *declared deserted*, or *text exceeds character limit*), as well as qualifiers that truly consider suggesting weaker logical statuses (e.g., *possibly*, *disputed*, *expected*, etc.).

Therefore, suggested values provided by the *Property Talk* pages have been ignored, while the actual values found in the datasets have been focused on. The list of terms has been surveyed and a subset of 101 terms that seem to

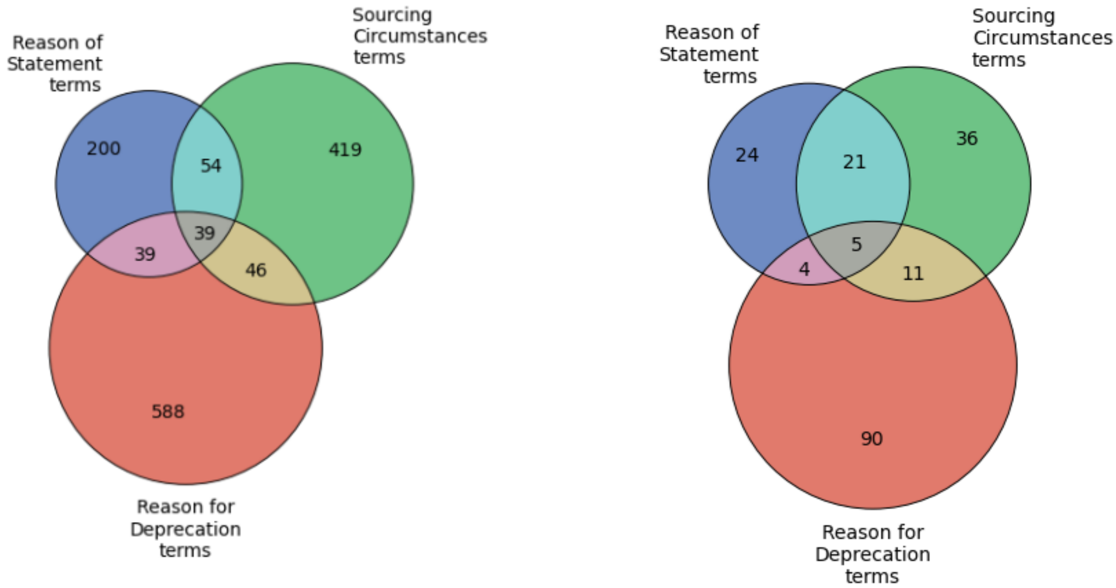

Figure 3.2: Terms used in qualifiers *nature of statement*, *sourcing circumstances* and *reason for deprecated rank* throughout Wikidata (*left*) and in the CH datasets (*right*)

refer to WLS claims concretely has been selected. This subset of WLS terms appears to be widespread in CH and Random datasets (2,086 occurrences in CHav, 111.641 occurrences in CHv, 1,318 occurrences in CHt and 6,406 occurrences in R), while almost not employed in Astronomical datasets (62 occurrences in ANs and only 1 in ANg).

The distribution of approaches to represent WLS claims in the CH dataset is not homogeneous, as unknown values and WLS-qualified statements are both highly used in the CHv dataset, while non-asserted statements for CHav and CHt. An obvious outlier is the use of one specific qualifier. Indeed, the value *circa*[35] is by far the most employed value in CHv, appearing 107,653 times in *sourcing circumstances*. This brings the overall count of this value completely out of scale concerning other values (e.g., the second most frequent WLS term in CHv is *probably*, occurring only 1.676 times). A much more homogeneous distribution of values across the three CH datasets can be seen by removing specifically the value "circa" from the others in the last line of Table 3.1. On the contrary, many other terms in the list are present only once in the whole dataset

[35]`http://www.wikidata.org/entity/Q5727902`

and contribute very little to the overall impact of the qualified statements.

In theory, the approaches to represent WLS claims are **not** meant as alternatives to each other and to be used exclusively. It would be reasonable to use them on the same statement for the same entity, e.g., to describe a deprecated qualified statement that results as non-asserted. Yet, approaches co-occurrence in the surveyed datasets is poorly represented, and datasets demonstrate very few cases of use of multiple WLS approaches for the same statements. In particular, no co-occurrence can be found in the AN dataset because almost all WLS claims are expressed via ranked statements except for a little co-occurrence of deprecated statements marked with a WLS qualifier in the ANs dataset (0.1%). Co-occurrences between approaches representing WLS information seem poorly implied in CH datasets. Almost no co-occurrence could be found between unknown and deprecated statements (0.1% in CHav, 0.04% in CHv and none (0%) in CHt), as well as the co-occurrence of deprecated and WLS-qualified statements (0.04% in CHav, 0.01% in CHv, 0.07% in CHt), as well as the co-occurrence of unknown and WLS qualified statements (in 0% in CHav, 1.14% in CHv and 0.13% in CHt).

To summarise, it becomes manifest that the prevalence of each approach is quite diverse, even between the datasets of the same domain. Specifically, in CHav the most commonly used approach representing WLS information is non-asserted (86.09%), in CHv it is the WLS Qualified statement (49.13%) followed by unknown value (46.88%), and in CHt it is non-asserted (81.64%). In the Astronomy datasets, non-asserted statements overwhelmingly represent WLS claims, but deprecated statements have a much larger impact on them than in the CH domain.

The property analysis provides valuable insights, too, as shown in Figure 3.3. The actual usage of WLS approaches was divided based on the properties where they appear. The x-axis contains, for each dataset, the ten most frequent properties in which WLS statements appear. The y-axis shows in logarithmic

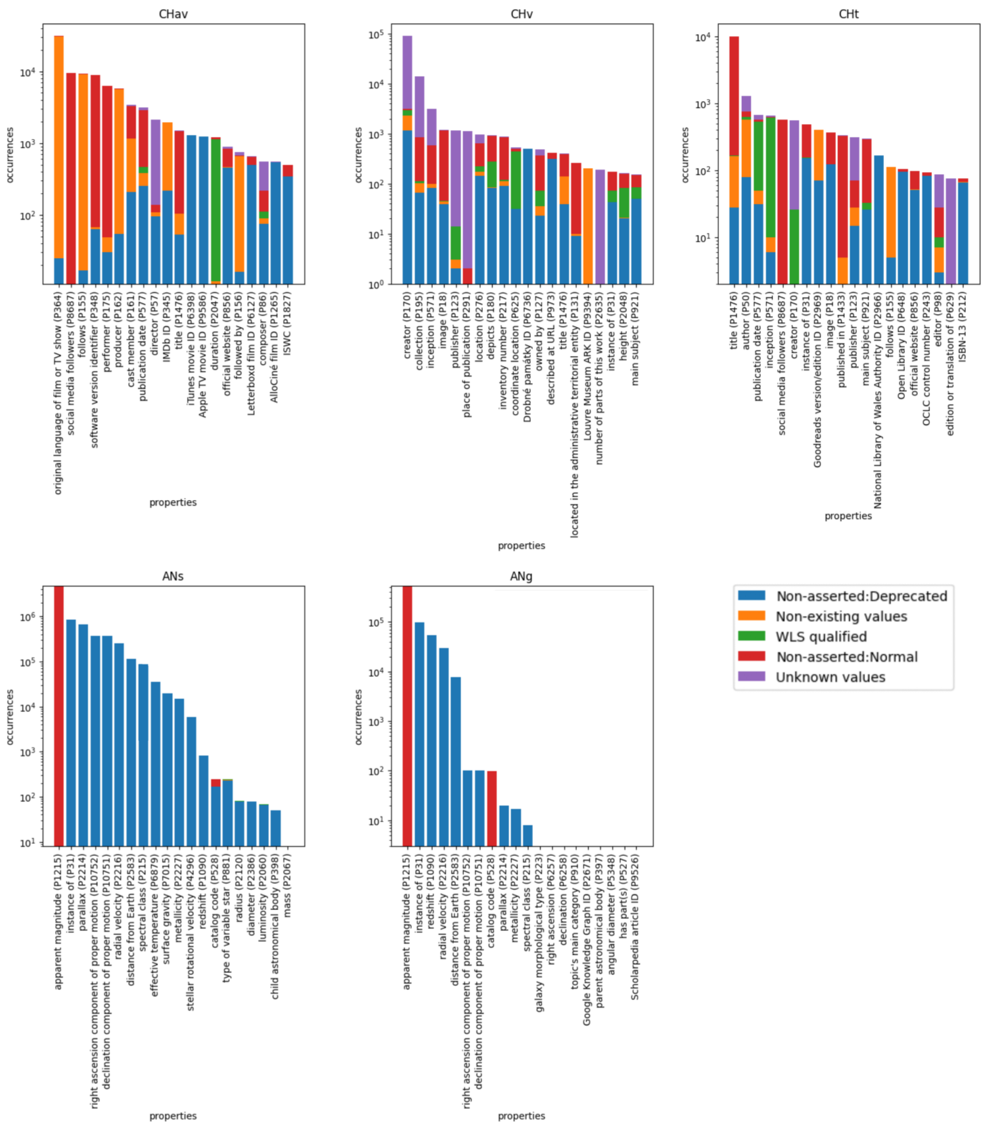

Figure 3.3: Top 10 most recurrent properties implied in WLS claims in each disciplinary dataset

scale the number of occurrences of such statements, organised by colour: non-asserted statements (with rank normal), non-asserted statements (with rank deprecated), statements with qualifiers (only WLS-related qualifiers), and non-existing valued statements.

The datasets were analysed by systematically evaluating the properties associated with the surveyed approaches. Each dataset was analysed to identify (1) the most prominent properties of each dataset and (2) the most prominent properties of each dataset with each approach.

Normal ranked, yet non-asserted statements appear in large numbers in CHav for `P8687`:*Social media followers*, `P348`:*software version identifier*, `P175`:*performer* and `P1476`:*title*. They represent peculiar uses of the non-asserted normal ranks for statements that represent multiple, independent values for the same property, none of which is "more important" than the others. Similar reflections can be made for `P18`:*image* on dataset CHv, and properties `P1433`:*published in* and `P921`:*main subject* in dataset CHt. The property `P1215`:*apparent magnitude* dominates this category for astronomical data. Most of the remaining properties employ a deprecated rank for evolving or uncertain information. Despite the different designed uses for deprecated and preferred rankings, Figure 3.3 shows that Non-asserted:Deprecated and Non-asserted:Normal claims highly co-occur with the same properties.

Qualified statements are largely present in CHv and CHt on properties `P571`:*inception*, `P577`:*publication date*, and `P625`:*coordinate location*, where, as mentioned, the *circa*[36] value for qualifier dominates the occurrences.

Unknown valued statements are primarily used in CHv and CHt datasets and only sparsely in CHav dataset. Their usage seems to be mainly implied in the description of agents in roles in all CH datasets (e.g., `P170`:*creator*, `P98`:*editor*, `P123`:*publisher*, `P50`:*author*, `P86`:*composer*, `P57`:*director*). In CHv and CHt datasets, their usage includes also locations (e.g., `P195`:*collection*, `P291`:*place of publication*), time (e.g., `P571`:*inception*) and the artworks' de-

[36] `https://www.wikidata.org/entity/Q5727902`

scription (e.g., `P2635`:*number of parts of this work*, `P629`:*edition or translation of*). The significant prevalence of unknown values when annotating agents in roles related to artworks is evident in the CHv dataset, reflecting the paramount relevance of authorship attributions given by scholars in art history.

The predominance of non-existing valued statements in CHav can also be noted (`P364`:*original language of film or TV show*, `P155`:*follows*, `P156`:*followed by*, `P162`:*producer* and `P345`:*IMDb ID*), which goes to prove the peculiarity of the use of non-existing valued statements in the CHav dataset previously described. The dataset CHt has a considerable number of non-existing valued statements, too, but only on properties `P1476`:*title* and `P50`:*author*, for untitled and/or anonymous documents.

Besides this, some co-occurrences of the use of unknown and non-existing valued statements with the same properties have been registered (e.g., `P57`:*director* in CHav, `P170`:*creator*, `P291`:*publisher*, `P180`:*depicts*, `P571`:*inception*, `P127`:*owned by* in CHv and `P50`:*author*, `P98`:*editor*, `P123`:*publisher*, `P577`:*publication date* in CHt, as shown in Figure 3.3).

To summarise, some of the complexities and ambiguities identified in both the CH and AN datasets, besides their designed use as described by Wikidata, are listed (cf. Section 3.1). The list comprehends a more fine-grained distinction of WLS situations.

- **Ranked statements**

  - *Evolving situation*: The claim is not true at the moment but was correct at some point in the past, and keeping this information is deemed interesting to maintain. For instance, the number of `P8687`:*social media followers* of artists and politicians, the change of `P276`:*location* of a movable cultural object such as a painting or a statue, or the change of its `P6216`:*copyright status*, may change over time. This change is recorded via differently ranked statements.

For instance, the print "At the Races: Anteriel"[37] recently shifted from copyright to the public domain. In this case, the deprecated statement was correct up to a given moment in time but is not correct anymore.

- *Evolving knowledge*: Because of a new observation or theory, a previous value is considered superseded. This situation is mainly connected to new observations, theories, measurements, guesses and interpretations. For instance, the introduction of a new accepted attribution of a work of art means that the previous one is now deemed as false or at least deprecated, or, in astronomy, the object "15 Orionis"[38] was previously considered an `P31`:*instance of* an *infrared source*[39], but it is now fully considered as a *star*[40]; in this case, the deprecated statement has always been incorrect, but it has been decreed as such only after a specific moment in time.

- *Less favoured versions*: Similar claims are ranked not because they are either false or true but because one of them is preferred over the others so that they are marked as preferred and asserted while the others are non-asserted. For instance, the `P1476`:*titles* of textual works are often provided in different languages, and the title in the original language is marked as the preferred version, while the translated titles in other languages are not asserted. In this case, the deprecated statement is not incorrect, but it has been demoted to prioritise another one. This is not strictly a WLS situation but uses the same ranking approach as truly WLS ones.

- **Qualified statements**

  - *Uncertainty*: For instance, the painting "Madame Antoine Ar-

[37] `http://www.wikidata.org/entity/Q79471408`
[38] `http://www.wikidata.org/entity/Q6675`
[39] `http://www.wikidata.org/entity/Q67206691`
[40] `http://www.wikidata.org/entity/Q523`

nault"[41] has `P170`:*creator* set to Jean-Baptiste Regnault[42] with a `P5102`:*nature of statement* qualifier *disputed*; Here, the statement is not certain, and competing (and incompatible) statements may be present or at least expected.

- *Caution*: For instance, the "Frontispiece to Christopher Saxton's Atlas of the Counties of England and Wales State I"[43] has the `P170`:*creator* property set to Remigius Hogenberg[44], with the contributor cautioning through a `P5102`:*nature of statement* qualifier that this is only an *attribution*[45]. Here, the statement is not certain, but it implies that the proposed value may be wrong rather than positively asserting disagreements on it.
- *Imprecision*: For instance, the hypothetical entity "IRAS 17163-3907"[46] has an observed `P2060`:*luminosity* property set to "500.000 solar luminosity" with a `P1480`:*sourcing circumstances* qualifier *circa*; similarly, the painting "Girl Reading a Letter at an Open Window"[47] by Johannes Vermeer is dated (`P571`:*inception*) 14th century with a *sourcing circumstances* qualifier *circa*. For instance, the star "Altair" (`Q12975`) has a `P1102`:*flattening* property set to 0.2 with a *nature of statement* qualifier *greater than*; Here, the statement is certain but the value is inherently loose. One may wonder if this is truly a WLS statement or a positive statement of an imprecise value.

- **Missing value statements**.
  - *Data entry errors*: Data include errors probably introduced during

[41] https://www.wikidata.org/entity/Q109252498
[42] https://www.wikidata.org/entity/Q453485
[43] https://www.wikidata.org/entity/Q105949375
[44] https://www.wikidata.org/entity/Q18576859
[45] https://www.wikidata.org/entity/Q230768
[46] https://www.wikidata.org/entity/Q540167
[47] https://www.wikidata.org/entity/Q700251

the annotation. For instance, the novel "Invisible Monsters"[48] is both attributed to Chuck Palahniuk (the actual author) and an unknown and probably erroneous entity. Here, there is a clear error in the dataset. Whether a `someValue` or a `noValue` is used is not important as they would both be errors.

- *Dumping from pre-existing databases*: Some non-existing values may result from an error in the conversion or an empty field of a record after importing an existing database into Wikidata. For instance, the painting "Marshy Landscape"[49] has a non-existing valued statement for the `P528`:*catalogue code* property. Again, this represents an error in the dataset, so the corresponding statement should be omitted.

- *The value does not exist*: For instance, the first and last entities of a sequence use properties `P155`:*follows* and `P156`:*followed by* with a non-existing value. For instance, the first episode of a TV series or the last song of a recording should have non-existing values for the corresponding properties. This is a correct use of a `noValue`, not a WLS claim.

- *Model fitting*: When the model does not fully support the situation to be described, some arrangements were taken, such as the use of a non-existing value for the property *original language of film or TV shows* `P364` when the entity is a silent movie. For instance see "Silent Tests"[50], whose `P364`:*original language of film or TV show* predicate is non-existing valued and additionally qualified with `P518`:*applied to part dialogue* [51]). Here, a non-existing value is correctly used for a value not felt necessary in the model (e.g., a specific property *language of dialogue* to be used in sound fields and omitted for silent

[48] `https://www.wikidata.org/entity/Q2600527`
[49] `https://www.wikidata.org/entity/Q6773948`
[50] `https://www.wikidata.org/entity/Q390207`
[51] `https://www.wikidata.org/entity/Q131395`

movies). This is, again, a correct use of the noValue claim, yet not a WLS claim.

- *The value exists but is not known*: For instance, the painting "The Welcome Home"[52] is marked to have an unknown `P170`:*creator* as a `someValue` blank node. This is probably the only true WLS use of missing value statements.

The previous list shows a series of situations where the same approaches are used for different purposes. All such purposes (except data entry errors) are legitimate. Yet, it may be difficult for users to differentiate the purpose of each use because the chosen approaches are not sufficiently precise to distinguish the specific situation clearly and unambiguously. Rather than suggest forcing all different situations into a single over-encompassing approach, Section 3.4 lists some increasingly impactful solutions to solve these ambiguities without overly revolutionising the data model.

3.3.1 Discussion

The datasets presented in the previous section and our analysis of their content allow us to reach some conclusions on the research questions specified in the introduction.

RQ1a – *How widespread are WLS claims in the current state of Wikidata?*

The current state of WLS claims in Wikidata is poor. Even though Wikidata focuses on collecting and referencing the facts claimed elsewhere[53] [107], rather than conjectural or controversial information[54], in many cases it is objective and scientifically precise to represent the complexity of uncertainty and evolving knowledge, rather than omitting information because they are not completely established. In these cases, Wikidata seems to be doing poorly, as

[52] `https://www.wikidata.org/entity/Q110041706`
[53] `https://www.wikidata.org/wiki/Wikidata:Verifiability`
[54] `http://www.wikidata.org/wiki/Help:Ranking#What_ranks_are_not`

<1% of the claims under analysis in CH datasets show weaker logical status characteristics, 5% in the ANg dataset and 27.41% in the ANs data. Of course, finding a reference that backs the uncertainty of a claim (e.g. it is disputed) can be rarer than a reference to facts that are unequivocal for their annotators. Thus, it is natural that WLS claims in Wikidata generally appear with a much lower percentage than certain facts. Nonetheless, CH datasets show a much lower figure than, e.g., the ANg and ANs datasets. Does this show an intrinsic difference in the two cultural domains, or is there something else underneath? To provide an answer to this further question, the RKD database was consulted.

RKD[55] holds detailed data about Dutch and Flemish paintings, drawings and prints throughout the ages, from XVI Centuries artworks to modern ones. Overall, more than 260,000 items belonging to the image collections are described in the database, and through an EDM-inspired data model, particular attention is given to multiple competing assertions, e.g., incompatible authorship attributions. Namely, RKD contains more than 317,000 recorded attributions, i.e., an average of 1.2 attributions per artwork. Thus, deprecated authorship attributions are present in about 8.5% of the works in the RKD image collection (e.g., about 290,000 current attributions vs. 27,000 discarded ones in the RKD images collection), a conspicuously higher figure than the meagre 1.77% WLS statements of the CHv dataset.

One may wonder that Dutch and Flemish collections are not representative of the full scale of worldwide types of artworks represented in the CHv dataset. Yet, they provide an interesting starting point for a further comparison. A sub-dataset of CHv has been created and further analysed it to improve our understanding of this issue. First of all, it should be noted that, as mentioned, about 83,600 artwork descriptions out of the 267,238 available in RKD have been linked to Wikidata[56], representing ~7.5% of the total of visual artworks in Wikidata. Thus, a dataset in Wikidata with the same artworks

[55] see `https://rkd.nl/en/`
[56] `https://w.wiki/7wfW`

	RKD Images (17th-20th c. Dutch Paintings)	Wikidata (17th-20th c. Paintings)
Paintings	267,238	501,049
Attributions	317,165	340,661
Current Attributions	289,918	340,213
Discarded Attributions	27,247	448
% Discarded	8.5%	0.13%

Table 3.2: Comparison between attributions in the RKD images collection and Wikidata for 17th-20th century paintings.

of RKD inevitably risks being polluted by RKD data itself. Since RKD is highly specialized in Dutch paintings from the 17th to the 20th century, the created Wikidata sub-dataset contains European artworks painted in the same temporal period and *explicitly excluding RKD artworks*[57]. Wikidata stores 501,049 paintings in the interval 17-20th century not present in RKD, for a total of 340,661 attributions[58]. The results of such comparison are shown in Table 3.2. Out of the total number of Wikidata statements, only 0.13% of the items are discarded attributions (448)[59]. This fact may indicate a radical under-representation of complex attributions within Wikidata entities. In conclusion, WLS statements are not particularly widespread nor successful in Wikidata collections within the CH domain, and they arguably misrepresent the complexity and variety of situations in this domain.

RQ1b – *How does the cultural domain of the Wikidata topics (and, presumably, of the individuals contributing to the data regarding the Wikidata topics) affect and reflect on the relative success of some WLS types over others?*

Our data analysis highlighted several peculiarities between the CH and As-

[57] `https://w.wiki/7VRg`

[58] The count of attributions is calculated over the number of claims having the predicate `P170`:*creator*

[59] The number of discarded attributions is calculated over the number of claims having `P170`:*creator* as predicate and not being asserted

tronomical datasets. The two families of datasets present many different representational artefacts: while the CH datasets seem to employ, with variable proportions, all the listed approaches, the astronomical datasets employ almost exclusively ranked statements. Additionally, while WLS statements in AN datasets affect a fairly small number of properties, they cover a much wider range of properties in CH, as shown in Figure 3.3. These aspects highlight key differences in what the two communities consider weaker logical status: It can be hypothesised that deprecations in astronomical data mostly reflect the result of newer and better data. In contrast, the humanities community uses WLS statements for a much larger set of uncertainties due to ignorance, scholarly interpretations and disagreements as hypothesised in Section 1.4. Thus, it may occur that the specification of the `P5102`:*nature of statement* and the `P2241`:*reason for deprecated rank* qualifiers may seem superflous in astronomical data, and a real necessity for some annotations in the humanities.

RQ1c – *Does the actual usage of the WLS claims represent critical inquiry on humanities data?* – Wikidata provides a set of designed uses for WLS claims annotation described in Section 3.1. These methods are suitable for representing critical inquiry, evolving information, and hypotheses in RDF. In addition to them, Wikidata contributors have, over time, adopted frequent annotation patterns that are only sometimes aligned with designed uses. Thus, there is much noise and ambiguity in how Wikidata contributors have used approaches provided by Wikidata to represent WLS information in the surveyed datasets. This makes it difficult to differentiate, search and retrieve WLS data. The variety of cases listed at the end of Section 3.3 summarises an incomplete yet vast collection of WLS and non-WLS situations modelled through the same WLS representation approaches. Therefore, it is difficult to search for specific data patterns over the entire dataset and even to interpret individual entities correctly. In particular, such ambiguities can be listed explicitly for the surveyed approaches:

- Ranked statements associated with asserted and non-asserted logic represent WLS information as the evolution of opinions in critical debate (e.g., evolving knowledge), historical information (e.g., evolving situation) and non-related WLS information (e.g., less preferred variant). Additionally, despite the different designed uses for preferred and deprecated statements, in practice, they frequently co-occur in the CH dataset for the same properties, showing that annotators arbitrarily choose between these two approaches to represent such information (e.g. discarded attributions are sometimes represented with a non-asserted normal rank and sometimes with a deprecated rank). Another gap has been identified in Wikidata's reification processes, which tend to oversimplify critical inquiry, failing to link statements that form part of a single theory. As noted in the description of Wikidata patters to express WLS claims (Section 3.1), Wikidata expresses claims by reifying each triple in the dataset. While no properties are available to link statements (e.g., marking that they are related to each other), they may be retrieved theoretically by addressing the same provenance information (e.g., qualifiers and references). Since no guidelines are available for this, practising it is impossible. This aspect will be addressed while discussing the reification methods and expressing without asserting in the humanities (Chapter 4.2.5).

- The selection of terms provided with *nature of statement* and *sourcing circumstance*is not exclusively related to WLS information, so that a subset of terms should be defined for this specific purpose (cf. 101 selected terms in Section 3.3). Such selection of terms has demonstrated that annotators use arbitrary terms even if declared in property Talk-related pages, especially in Humanities datasets (CH datasets). This highlights an over-fitting of expressive patterns over the two properties, representing the claim's motivation and background studies (e.g., attribution) and the uncertainty markers of such claims (e.g., possibly). Additionally, no taxonomy is provided for the types of WLS qualifiers. For this reason, automatic extraction of types of uncertainty (such as uncertainty, cautioning, and imprecision as discussed in Section 3.3) or motivations cannot be automatically performed. (3) Despite the designed use provided by Wikidata, the two types of missing values statements (noValue and someValue) present a significant co-occurrence within the same properties, indicating an unclear usage similar to the usage of ranked statements.

Furthermore, using the same approaches for WLS and non-WLS-related characterisations makes complex patterns hard to express and identify. For instance, if an artwork AW was *supposedly* moved from location X to location Y, but it is not certain, both locations X and Y must be represented as WLS, the first because of an evolving situation (AW is not at location X anymore) and the second because of uncertainty since the new location Y is only guessed. Therefore, none of these assertions can be asserted, and none can be ranked as preferred. A complete and thorough contextual annotation is required (e.g., why each claim is discarded). without this, disambiguation and a full understanding of the state and truth of the relevant predicate become impossible. In Section 3.4, a possible pattern is suggested to represent such situation (cf. point 5, in particular, *normal rank + non-asserted*).

3.4 Towards a leaner and harmonic support for WLS in Wikidata

Detailing workable solutions to improve the situation for WLS statements in a project as large and as complex as Wikidata always runs the risk of becoming an exercise in futility. This section suggests possible actions for WLS statements, starting from very conservative proposals with limited impact to more impacting changes.

Possible theoretical remediation activities for the Wikidata data model and the collection are listed to simplify and disambiguate WLS claims from others. Such suggestions are expressed as an ordered list whose first items are meant as simple cleaning-up activities of little impact and then progress to bolder and more impacting actions that sometimes require not just a modification in the data model but possibly also the systematic update of small, but still numerically relevant, selections of the current datasets.

1. Reorganise, simplify, and re-categorise the suggested values for qualifying properties `P5102`:*nature of statement*, `P1480`:*sourcing circumstances*. The lists should be clearly differentiated and with no semantic overlaps neither between lists nor within each list.

2. Require a `P7452`:*reason for preferred statement* qualifier in all preferred statements and a `P2241`:*reason for deprecated statement* qualifier in all deprecated statements. Provide simple-to-use interface widgets for their specification. Make sure such statements can only be saved with a qualifying proposition.

3. Require the specification of `P5102`:*nature of statement* and `P1480`:*sourcing circumstances* qualifiers for all WLS-related rankings: only asserted statements with normal rank are allowed to remain without qualifiers.

4. Create a new and separate *Certainty Degree* qualifier specifically for WLS statements, separating the reason for the chosen qualification from the certainty or confidence degree of the qualification. Such certainty degree should be scalar and use a limited number of values, avoiding any complexity in distinguishing between terms such as possibly, hypothetical, and dubious. A 5- or 7-item scale would suffice, e.g., *non accepted, highly unlikely, unlikely, possible, probable, almost surely*, and *accepted.* Different labels would be perfectly acceptable, even using numerical values instead of labels.

5. Reorganise the values of `P5102`:*nature of statement* and `P1480`:*sourcing circumstances* to remove values merely representing an uncertainty (replaced by the new *Certainty Degree* qualifier). To this end, an initial

list of values is being created. The current list has been generated by following a Grounded Theory approach [49]: first, labels, definitions and usage data of suggested and used qualifiers have been collected and categorised to represent different macro-themes or concepts. These concepts allowed theories to emerge and be developed from the coded data with an iterative process that continued until the theory was "grounded" in the data. The resulting list in its current state, collecting the surveyed terms from the Wikidata *Property Talk* pages and the terms used in the CH datasets, contains 150 values referring to WLS claims and organised in 18 theories and can be accessed in the GitHub folder of the project[60].

6. Restrict ranking for competing statements to just three (possibly four) different patterns and prevent any other variant:

    - *Preferred + Deprecated*: To be used whenever there are several competing statements, and some are chosen to be the best. Accepted statements are set to preferred (and asserted), while the rest are set to deprecated (and not asserted); there are no normal ranks. Both preferred and deprecated statements are fully qualified with `P5102`:*sourcing circumstances*, `P2241`:*reason for deprecated statement* and `P7452`:*reason for preferred statement* respectively, and the new Certainty qualifier. Preferred statements would be assigned an *accepted* or *almost surely* degree, while deprecated ones would be assigned a *not accepted* or *highly unlikely* certainty degree. Intermediate degrees would not be used.
    - *Normal rank + asserted*: This would be the default situation, to be used when no dispute or disagreement exists and the statement(s) are all equally accepted. All statements are also asserted. Since this is the default, no qualifier is necessary, but it is still possible to specify a `P5102`:*nature of statement* or a `P1480`:*sourcing circumstance*

[60] `https://github.com/alessiodipasquale/Wikidata_WLS`

value. No certainty degree is necessary.

- *Normal rank + non-asserted*: To be used when there are several competing statements but none of them stands above the rest as being the most likely. For instance, this would be the case of a work of art not definitely attributed to anyone but for which several competing hypotheses exist. However, none seem more convincing than the others. No statement is asserted, and `P5102`:*nature of statement* and/or a `P1480`:*sourcing circumstance* values are required. All statements would be assigned a value from the central ones, from *highly unlikely* to *probable*, excluding the extremes.

A fourth pattern could be allowed for claims for which the only reported value is wrong, but no acceptable alternatives exist. In this case, deprecated statement can be used for the reported wrong value and a non-existing valued statement with a normal rank to represent the non-existing correct value.

3.5 Conclusions

This work is the first systematic study about the representation of weaker logical status claims (WLS) over CH data in Wikidata. Through WLS claims, critical inquiry, uncertain information, competing hypotheses, temporally evolving information, etc., for which a plain and direct assertion is inappropriate, can be expressed. Three patterns used in Wikidata for WLS claims have been analysed: asserted vs. non-asserted statements achieved via ranked statements, missing values, and qualifiers.

The analysis reveals several interesting facts. First, the number of statements expressed using a weaker logical status is much lower than might have been expected by comparing similar sources. Secondly, the Wikidata data model is far from being too poor to express WLS claims; it offers users an overabundance of approaches, but their applications overlap and are also used for non-WLS applications. Finally, significant differences exist in how datasets from different domains employ these approaches for weaker logical status claims. Domain-specific non-WLS situations can be considered as a justification for much of this variety, and this contributed to the idea that WLS-specific features should be introduced in the Wikidata model to address specifically weaker logical status claims. A set of increasingly impactful modifications (points 1 to 6 in Section 3.4) to the data model was proposed, aimed at achieving a leaner and more accurate representation of these phenomena, expecting that they can improve data quality and information retrieval, specifically over uncertain, evolving and competing statements.

This analysis highlights the extensive contextualisation required for claims related to critical inquiry in the humanities, as emerged by comparing them with astronomical data. It underscores the pivotal role of context in humanities data. Wikidata's data model incorporates the representation of critical inquiry by considering its evolution by introducing its customised reification method supplemented by qualifiers to express claims and contextual information, rankings and non-asserted statements to annotate their logical status. However, this representation appears underutilised in practice, raising important questions about how critical inquiry is portrayed within knowledge graphs.

A set of increasingly impacting modifications has been defined on the data model aiming towards a leaner and more accurate representation of these phenomena with the expectation that they can improve data quality and information retrieval, specifically over questionable, evolving and competing statements.

Chapter 4

Expressing Without Asserting

The research presented in this Chapter is a reworking of the studies published in [25, 106, 80]. Articles are reported with the consent of all authors.

The Chapter examines and discusses the need and usefulness of EWA methods as a possible solution to represent the evolution of critical inquiry. It tackles RQ2 - *How can critical inquiry be effectively and efficiently represented with minimal complexity?* and redefines it into four more specific research questions.

1. RQ2a - *How do existing reification approaches express without asserting?*
2. RQ2b - *Is there evidence that existing approaches for EWA are ineffective?*
3. RQ2c - *Is there a way to represent EWA claims that can overcome such deficiencies?*
4. RQ2d - *What is the most efficient way to implement Express Without Asserting?*

4.1 Approaches to express without asserting

Reification methods are usually studied and tested for their ability to represent contextual information and allow fast retrieval effectively. EWA is a powerful tool for introducing disputed or questionable statements mimicking the historical evolution of critical inquiry. However, EWA seems to be tackled more as a side effect of reification methods. No survey assesses the effectiveness and efficiency of such methods in EWA's contents. This section addresses the effectiveness of such methods from an ontology-independent perspective.

Effectiveness is defined as the ability of a certain reification method to express without asserting claims when necessary, combining accuracy, cost-benefit, and the thoroughness of the representation. As mentioned in Section 2.2, a reification method can be said to achieve EWA if the contents of a claim are considered asserted (i.e., *true*) when a SPARQL engine evaluates the reified triple as true via an ASK query and false if otherwise.

Reification (Listing 4.1) is part of RDF 1.1 semantics. A reified triple is not asserted, therefore complying with our minimum requirement. However, since it requires an expensive addition of triples (4+) and it does not allow inference on reified triples, it has been disregarded as a solution to describe provenance and it has been proposed for deprecation[1] [108].

```

# Attributed to Rembrandt
:aa1 rdf:type rdf:Statement ;
    rdf:subject :painting-pr ;
    rdf:predicate crm:P14_carried_out_by ;
    rdf:object ulan:500011051 ;
    crm:P4_has_time-span :XVIII_cent .

# Attributed to Hooc
:aa2 rdf:type rdf:Statement ;
    rdf:subject :painting-pr ;
    rdf:predicate crm:P14_carried_out_by ;
    rdf:object ulan:500020229 ; # Hooch
    crm:P4_has_time-span :1821 .

# Attributed to Vermeer
:aa3 rdf:type rdf:Statement ;
    rdf:subject :painting-pr ;
    rdf:predicate crm:P14_carried_out_by ;
```

[1] See https://lists.w3.org/Archives/Public/public-rdf-wg/2011Apr/0164.html

```
    rdf:object ulan:500032927 ; # Vermeer
    crm:P4_has_time-span :1860 ;
    crm:P14_carried_out_by ulan:500326948 . # Thore
```

Listing 4.1: Competing attributions with Reification

N-ary relations (Listing 4.2 are adopted in for instance by CIDOC-CRM. To represent concurring statements, instances of the class `crm:E13_Attribute_Assignment` are annotated with context information. Despite CIDOC-CRM primer [36] not specifying how to record such information, Linked Art specifies the representation of attribution using blank nodes. This method shows some verbosity and redundancy of triples to add context to statements. Additionally, all statements are considered asserted, and therefore, in our example, the two concurring statements are equally asserted, despite the debate being settled towards only one.

```
:painting crm:P108i_was_produced_by [ a crm:E12_Production ;
:pr crm:P14_carried_out_by  ulan:500032927]

# Attributed to Rembrandt
:painting crm:P108i_was_produced_by [ a crm:E12_Production ;
            crm:P140i_was_attributed_by [ a crm:E13_Attribute_Assignment ;
                    crm:P141_assigned [ a crm:E12_Production ;
                            crm:P14_carried_out_by ulan:500011051 ] ;
                    crm:P177_assigned_property_of_type crm:P14_carried_out_by ;
                    crm:P4_has_time-span :XVIII_cent .
                    ] ]

# Attributed to Hooc
:painting crm:P108i_was_produced_by [ a crm:E12_Production ;
            crm:P140i_was_attributed_by [ a crm:E13_Attribute_Assignment ;
                    crm:P141_assigned [ a crm:E12_Production ;
                            crm:P14_carried_out_by ulan:500020229 ] ;
                    crm:P177_assigned_property_of_type crm:P14_carried_out_by ;
                    crm:P4_has_time-span :1821.
                    ] ] ;

# Attributed to Vermeer
:painting crm:P108i_was_produced_by [ a crm:E12_Production ;
```

```
        crm:P140i_was_attributed_by [ a crm:E13_Attribute_Assignment ;
                crm:P141_assigned [ a crm:E12_Production ;
                        crm:P14_carried_out_by ulan:500032927 ] ;
                crm:P177_assigned_property_of_type crm:P14_carried_out_by ;
                crm:P4_has_time-span :1860 ;
                crm:P14_carried_out_by :wd:Q500326948 .
                ] ] ;
```

Listing 4.2: Competing attributions with n-ary relations (CIDOC-CRM data model)

Wikidata [40] (Listing 4.3) make use of its own reification method to introduce statements[2] [40] by the means of Wikibase data model[3]. Wikidata does not reify RDF triples, but Wikidata objects [107]. Additionally, they use URIs for auxiliary resources avoiding blank nodes for objects of the data model. As mentioned in the description of Wikidata methods to represent WLS claims (Section3.1), different prefixes are used to differentiate types of information addressing the same predicate. As shown in Listing 4.3, the predicate `crm_P14_carried_out_by` is used three times in `:statement3` to indicate the main predicate of the claim (introduced by the predicate `s`), the reified predicate (introduced by the predicate `ps`) and the author of the claim (introduced by the predicate `pq`). This approach allows the same predicates to be applied in various contexts based on their prefixes. However, such mapping should be explicitly declared for all the properties via the means of the Wikibase data model [31].

```
# Attributed to Rembrandt
:painting-pr s:crm_P14_carried_out_by :statement1 .
:statement1 a wikibase:Statement ;
    ps:crm_P14_carried_out_by wd:Q500011051 ;
    pq:crm_P4_has_time-span :XVIII_cent .

# Attributed to Hooc
:painting-pr s:crm_P14_carried_out_by :statement2 .
```

[2] https://www.wikidata.org/wiki/Help:Statements

[3] The full semantics of Wikibase data model is available al https://www.mediawiki.org/w/index.php?title=Wikibase/DataModel

```
:statement2 a wikibase:Statement ;
    ps:crm_P14_carried_out_by wd:Q500020229 ;
    pq:crm_P4_has_time-span :1821 .

# Attributed to Vermeer
:painting-pr s:crm_P14_carried_out_by :statement3 .
:statement3 a wikibase:Statement ;
    ps:crm_P14_carried_out_by wd:Q500032927 ;
    pq:crm_P4_has_time-span :1860 ;
    pq:crm_P14_carried_out_by :wd:Q500326948 .
```

Listing 4.3: Competing attributions with n-ary relations (Wikidata data model)

A **singleton** (Listing 4.4) is a unique predicate that is used only once to represent a triple instead of the original predicate. As such, it can be the subject of additional triples. The property `:singletonPropertyOf` maps new predicates to the original ones. However, it is a subproperty of `rdf:type`, therefore, the new statements must be considered asserted as well.

```
# Attributed to Rembrandt
:painting-pr :P14_carried_out_by#1 ulan:500011051 ;
:P14_carried_out_by#1 rdf:singletonPropertyOf crm:P14_carried_out_by ;
  crm:P4_has_time-span :XVIII_cent.

# Attributed to Hooc
:painting-pr :P14_carried_out_by#2 ulan:500020229 ;
:P14_carried_out_by#2 rdf:singletonPropertyOf crm:P14_carried_out_by ;
  crm:P4_has_time-span :1821 .

# Attributed to Vermeer
:painting-pr :P14_carried_out_by#3 ulan:500032927 ;
:P14_carried_out_by#3 rdf:singletonPropertyOf crm:P14_carried_out_by ;
  crm:P4_has_time-span :1860 ;
  crm:P14_carried_out_by ulan:500326948. # Thore
```

Listing 4.4: Competing attributions with Singleton Properties

Named graphs (Listing 4.5) present a peculiar situation. It has been argued that graphs do not contribute in determining the truth of a dataset [110, 5, 33], which depends on the interpretation of the default graph. At least eight model-theoretic semantics [110] have been proposed to decide the truth value of graphs. No agreement exists, and several interpretations can be in place in the same dataset. In our case, all attributions may or may not be asserted.

```
# Attributed to Rembrandt
GRAPH :aa1 {
    :painting-pr crm:P14_carried_out_by ulan:500011051.
    }
:aa1 crm:P4_has_time-span :XVIII_cent .

# Attributed to Hooc
GRAPH :aa2 {
    :painting-pr crm:P14_carried_out_by ulan:500020229.
    }
:aa2 crm:P4_has_time-span :1821 .

# Attributed to Vermeer
GRAPH :aa3 {
    :painting-pr crm:P14_carried_out_by ulan:500032927 .
    }
:aa3 crm:P4_has_time-span :1860 ;
    crm:P14_carried_out_by ulan:500326948. # Thore
```

Listing 4.5: Competing attributions with named graphs

In **N3**, graphs can be quoted by other formulas, which do not assert the contents of the RDF graph (Listing 4.6). N3 syntax is compact. Since graphs in N3 do not have an identifier, the annotation must be performed in place. Yet, N3 graphs can nest at arbitrary depth levels, while RDF 1.1 graphs cannot contain other graphs. Recent works [5] have proposed the usage of **N3 rules** to bind graphs to one of the possible model-theoretic semantics. However, this would require parties to adopt only one serialization to serve and consume data, since N3 rules are not backward compatible to RDF.

```
# Attributed to Rembrandt
{[crm:P4_has_time-span :XVIII_cent] :assignes [crm:P108_has_produced :painting;
crm:P14_carried_out_by ulan:500011051]} a n3:falsehood .

```

```
# Attributed to Hooc
{[crm:P4_has_time-span :1821] :assignes [crm:P108_has_produced :painting;
crm:P14_carried_out_by ulan:500020229]} a n3:falsehood.

# Attributed to Vermeer
[crm:P4_has_time-span :1860; crm:P14_carried_out_by ulan:500326948]
:assignes [crm:P108_has_produced :painting; crm:P14_carried_out_by  ulan:500032927]
```

Listing 4.6: Competing attributions with N3

RDF-star is a recent extension of RDF that allows a concise representation of statement-level metadata (Listing 4.7). However, multiple statements, encapsulated in graphs, cannot be quoted. It is backward-compatible to RDF, requiring the specification of 6 to 9 additional triples for every quoted statement. In order to leverage all the potentialities of this proposal, bespoke languages (RDF-star and SPARQL-star) and technologies (RDF-star graph stores) must be used.

```
# Attributed to Rembrandt
<< :painting-pr crm:P14_carried_out_by ulan:500011051 >>
    crm:P4_has_time-span :XVIII_cent.

# Attributed to Hooc
<< :painting-pr crm:P14_carried_out_by ulan:500020229 >>
    crm:P4_has_time-span :1821.

# Attributed to Vermeer
<< :painting-pr crm:P14_carried_out_by ulan:500032927 >>
    crm:P4_has_time-span :1860;
    crm:P14_carried_out_by ulan:500326948.
```

Listing 4.7: Competing attributions with RDF*

Besides the previously listed methods, other less famous exists as aRDF, RDF+, PaCe.

aRDF the representation of uncertainty representation and time constraints in RDF triples by annotating predicates in the form `s, p:a, o`, where `a` represents the annotation associated to the predicates (e.g., a temporal interval of validity). Triples are not asserted. However, a RDF-compliant representation is not provided.

RDF+ uses named graphs and triple identifiers as a quintuple, and uses named graphs in place of reification, and adds triple identifiers for the explicit annotation of meta knowledge. Yet, annotated triples are asserted, and their usage requires extensions of both RDF and SPARQL.

PaCE annotates RDF triples with provenance information in an incremental, reification-like style, allowing statements on the sameness/distinctness of triples. It has three implementations, it is RDF compatible, and it requires 4 additional triples. Annotated triples are asserted.

Unfortunately, a one-size-fits-all solution to represent truth value of graphs is yet to be found [59, 100, 95], and the evaluation of proposals focuses on efficiency in graph stores only [44, 59, 78].

As mentioned at the beginning of this chapter, this section discusses the effectiveness of most of these approaches. Indeed, a complete review would require us to expand the scope to considerations of meta-knowledge, uncertainty, and fuzziness. Instead, the focus will be placed solely on the capability of methods to express without asserting, as this represents the core of the problem space. In Table 4.1, a comparison of the effectiveness of approaches to EWA is presented. To the best of my knowledge, a shared definition of effectiveness that fits the proposals of EWA has not been defined. As mentioned at the beginning of this section, effectiveness has been intuitively defined as the combination of accuracy, cost-benefit, and thoroughness of the representation. The criteria are as follows:

- *EWA*: can arbitrary content be expressed without asserting it?
- *Graph*: can non-assertion be used with graphs when appropriate?
- *Mapping*: is it possible to map data to the plain semantics of RDF?
- *Increment*: how many triples must be added to map statements to RDF?
- *Independence*: is the approach independent from bespoke technologies[4]?

[4]e.g., parsers, serializations, query languages, etc.

- *Semantics*: does it extend RDF semantics?

Being compatible with RDF 1.1 semantics ensures reusability with RDF legacy technologies, therefore being a ready-to-use, effective solution. Yet, effectiveness without efficiency seems not to be appreciated, as confirmed by existing surveys [95]. Therefore, both *mapping* and *increment* have been examined to measure cost benefits.

Method	EWA	Graph	Mapping	Increment	Independence	Semantics
Reification [56]	yes	yes	yes	4	yes	no
N-ary [75]	no	no	yes	N	yes	no
N. graphs [15]	depends	depends	yes	0	yes	yes (RDFs)
N3 [11]	yes	yes	no	N.A.	no	no
RDF+ [98, 32]	depends	depends	yes	0	no	yes
aRDF [105]	yes	no	no	N.A.	no	yes
PaCE [94]	no	no	yes	4	yes	yes
Singleton [72]	yes	yes	yes	1	yes	yes
RDF-star [54]	yes	no	yes	6-9	no	yes
Wikidata [40]	yes	no	yes	3	yes	no

Table 4.1: Comparison of methods for expressing without asserting

Reification is part of RDF 1.1 semantics. It assumes that a reified triple is not asserted, therefore it is a good candidate for EWA. However, since its early stages [108], reification has been disregarded as a full-fledged solution to describe provenance since it does not allow inferences on reified triples. Moreover, it requires an expensive addition of triples (4+) to annotate a statement. For all these reasons, it has been often proposed for deprecation[5].

N-ary relations significantly increase the number of RDF triples required to represent uncertainty through additional assertions. Recent studies (e.g., [74]) have proposed extending ontologies to accommodate conjectural attributions in CH. Nevertheless, these proposals have not seen widespread adoption.

Wikidata statements are not asserted, while re-asserted triples are simply added. The data model employed by Wikidata has not been integrated into other frameworks addressing similar issues. It reifies Wikidata objects and relies on prefix tuning to annotate such statements. However, tuned predicates are mapped via the means of the Wikibase data model. This may be impractical in SPARQL queries where the predicate of the claim is not specified[6].

Named graphs present a peculiar situation. It has been argued that graphs do not contribute to determining the truth of a dataset [110, 5, 33], which depends on the interpretation of the default graph. The truth value of graphs can be decided case by case, and at least eight model-theoretic semantics [110] have been proposed. No agreement was found on shared semantics, and several interpretations can be made on the same dataset. Despite many works encouraging named graphs as a possible straightforward solution [107, 77] to introduce statements in RDF, they always raise the issue that they do

[5]e.g., `https://lists.w3.org/Archives/Public/public-rdf-wg/2011Apr/0164.html`

[6]See for instance the Wikidata SPARQL query to retrieve all statements of an item containing another item `https://w.wiki/BVto`

not come with shared semantics.

N3 introduces the concept of *quoted graphs*. A quoted graph is a graph that other formulas can quote, and that does not assert the contents of the RDF graph as being true. Differently from RDF 1.1, graphs in N3 do not have an identifier and, therefore, must be annotated in place. Yet, N3 graphs can nest at arbitrary depth levels, while RDF 1.1 graphs cannot contain other graphs. Recent works [5] have proposed the usage of **N3 rules** to allow data providers to bind graphs to one of the eight possible model-theoretic semantics. However, this would require parties to adopt only one serialisation to serve and consume data, since N3 rules are not backward compatible to RDF.

Singleton properties have been designed to cope with reification and n-ary relations issues. A singleton is a unique predicate used only once to represent a triple instead of the original predicate. Since it is associated with a single triple, this new predicate can become the subject of additional triples. RDF is (hypothetically) extended with the predicate `rdf:singletonPropertyOf` to map newly created predicates to the original ones. The original triple cannot be asserted, however, the property is a subproperty of `rdf:type`, and therefore the original predicates become classes, which can be seen as inconvenient since this changes the semantics of external ontologies and can negatively affect reasoning tasks.

RDF-star is an extension of RDF that allows a concise representation of statement-level metadata. Quoted statements are not asserted, therefore complying with the requirements of EWA. However, multiple statements encapsulated in graphs cannot be quoted. It is backwards-compatible with RDF, requiring the specification of 6 to 9 additional triples for every quoted statement. To leverage all the potentialities of this proposal, bespoke languages (RDF-star and SPARQL-star) and technologies (RDF-star graph stores) have been implemented.

In summary, achieving EWA with or without extending RDF semantics (respectively with RDF-star, singleton, and reification) is currently possible. However, this comes at the cost of adding triples (between 4 and 9), the impossibility of annotating graphs (RDF-star), the change of vocabularies semantics (singleton), and the complexity of tracking the re-assertion of EWA triples, which hamper our initial requirement of *evolving knowledge*. Named graphs alone cannot supply EWA because of the diverging semantics that data providers may adopt. Combining named graphs with N3 rules or RDF+ is possible, but it comes with technological constraints and the impossibility of being fully backwards-compatible with RDF. Singleton properties would be an effective RDF-compliant solution, but they change the semantics associated to original predicates.

Finally, a satisfying solution to EWA is not currently available. However, this work paves the way to two potential solutions, namely: (1) extending previous works to define a strategy that is fully compliant with RDF 1.1. semantics, without dramatically increasing the number of triples to be added; (2) extending the semantics of named graphs to make explicit and shareable the adopted model theory.

This issue is addressed in further in this chapter (Section 4.2), which details a new proposal called Conjectures, and where its effectiveness (Section 4.2) and efficiency (Section 4.3) are further compared with existing methods.

4.2 Conjectures

A *Conjecture* is a special RDF 1.1 named graph specifically designed to express statements whose logical status is not explicitly stated and, therefore, not asserted. Yet, similarly to RDF-star quoted triples, Conjectures express without asserting but apply to RDF graphs rather than individual triples (this aspect will be tackled in detail in Section 4.2.5). In this framework, all triples within the conjecture are expected to be consistent.

Among the 8 available semantics assigned to named graphs [110], Conjectures adopts the fourth semantic interpretation (where each named graph defines its own context, see point 3.4 in [110]). Therefore, a Conjecture is a special named graph that fully represents the semantics of point 3.4 of [110]: “Named graphs are considered as ‘hypothetical graphs’ which bear the same consequences as their RDF graphs, but they do not participate in the truth of the dataset; this [...] allows a graph to contain contradictions without making the dataset contradictory;”. Still, Conjectures absolute truth is neither stated nor available, and it does not affect the overall truth of the dataset. This choice was made to maintain compatibility with SPARQL entailment regimes while allowing for the validation of conjectural statements. By aligning with this semantic interpretation, Conjectures retains its compatibility with the SPARQL entailment regime while preserving the distinction between asserted and non-asserted statements.

Conjectures can be implemented in two, non-alternative ways (i.e. their interpretation is the same), one based on single-use properties compliant with RDF 1.1 semantics (*weak form*, documented in Section 4.2.1), and in a more compact form as an extension of the semantics of named graphs (*strong form*, documented in Section 4.2.1). The full semantics and interpretation of Conjectures is separately documented [92].

By introducing Conjectures, claims can be conceptually specified in different epistemic categorisations (and their related logical status), namely undisputed and disputed as defined in the research statement (Section 1.4, definition C5).

- **Undisputed statements** are expressed as plain asserted ($S(A)$) RDF 1.1 named graphs. They are introduced by the keyword `GRAPH`, such as $GRAPH(S) \Rightarrow S(A)$. Such claims have never been questioned (cf. undisputed statements in definition C5 in Section 1.4), such as the width of "Girl reading a Letter at an Open Window" being 64.5 centimetres without any doubt or the "Donation of Constantine" being written in Latin. Named graphs can be used to record the claim's provenance (e.g., claims are extracted from the RKD dataset).

- **Disputed statements** are expressed as Conjectural graphs, which in the strong form is a prototypical extension of the syntax of Trig, where the keyword `GRAPH` is replaced with `CONJ` in front of a graph whose contents is expressed but not asserted NA, such as $CONJ(S) \Rightarrow S(NA)$. A conjecture **does not** imply that the claim is false, but simply that it is questionable and has a weaker logical status (cf. disputed statements in definition C5 in Section 1.4). Therefore, a Conjecture on its own **does not** contribute to the truth of the dataset. For instance, "Girl reading a Letter at an Open Window" was said to symbolise love (before restoration) but also love that overcomes treachery (after restoration); the "Donation of Constantine" was said to have been written in the 4th century but also in the 8th century.

Additionally, a third type of graph, settled conjecture, explicitly represents claims while acknowledging prior disagreement. In this case, the relevant community has assessed competing stances (disputed claims), selected one as valid, and formally closed the debate (cf. settled statements in definition C5 in Section 1.4).

- **Settled statements** are a third type of named graph that is simultaneously conjectured and asserted and introduced in the strong form by the keyword `SETT` is introduced to handle settled disputes, such as $SETT(S) \Rightarrow \{S(A), S'(NA)\}$. Conjectural graphs can be re-asserted using this supplementary graph, which creates a dual state of claims as both conjectured and asserted simultaneously. This approach is specifically and intentionally distinct from a mere re-assertion of disputed claims, as it neither acknowledges nor mentions the existence of the dispute. The settled graphs represent both the conjectural triples (inside the usual conjectural graph) and the same triples but are completely asserted (inside the settled graph). For instance, "Girl Reading a Letter at an Open Window" symbolises the love that overcomes treachery, and the "Donation of Constantine" is acknowledged as a forgery from the 8th century. Settled conjectures represent the chosen point of view among concurring ones. Therefore, it equally represents a single theory which supersedes concurring stances (i.e., the case of Vermeer's painting, where a new interpretation supersedes the existing one) and community consensus (i.e., the case of the Donation, where the relevant community acknowledges the work as a forgery after the famous first philological work performed by Lorenzo Valla). In addition, the `conj:settles` relation connects the conjecture and its settlement, simplifying the task of exploring the relationships between disputes and their settlements. The rationale behind settled conjectures is two-fold: on the one hand, to stress the difference between claims that have not been challenged and claims that emerged as winning among competing and incompatible hypotheses and on the other to represent the dual nature of settled claims as both conjectures and assertions.

4.2.1 Conjectures in weak form

A conjectural graph in weak form is defined as follows:

Definition 1 A triple (s, p, o) that needs to be expressed without being asserted is represented with a unique newly minted conjectural predicate cp, and is described with two triples: (s, cp, o) and $(cp, conj{:}isAConjecturalFormOf, p)$.

By way of an example, consider Vermeer's painting: the interpretations before and after the restoration can be represented as two conjectural graphs whose predicates describe evolving claims (Listing 4.8). Thus, both the graphs (before and after) are stated as disputed and therefore conjectural (in the specific meaning of *non-asserted*).

```

# The painting was believed to represent love until 2022
GRAPH :paintingBefore {
    :painting conj0001:P65_shows_visual_item :girl;
              conj0001:P128_carries :love.
    conj0001:P65_shows_visual_item conj:isAConjecturalFormOf
            crm:P65_shows_visual_item .
    conj0001:P128_carries conj:isAConjecturalFormOf crm:P128_carries  . }
:paintingBefore crm:P4_has_time-span :until2022.

# The painting is nowadays believed (from 2022) to represent love that overcomes treachery
GRAPH :paintingAfter {
    :painting conj0002:P65_shows_visual_item :girl, :cupid;
              conj0002:P128_carries :loveOvercomesTreachery.
    conj0002:P65_shows_visual_item conj:isAConjecturalFormOf
            crm:P65_shows_visual_item  .
    conj0002:P128_carries conj:isAConjecturalFormOf crm:P128_carries . }
:paintingAfter crm:P4_has_time-span :from2022.
```

Listing 4.8: Vermeer's painting related claims in the weak form of Conjectures

In the weak form, conjectures adopt new predicates used exactly once, mapped to their original predicate via `conj:isAConjecturalFormOf`[7]. It is worth noting that the predicate `conj:isAConjecturalFormOf` is not related in any way to `rdf:type`.[8]. This aspect is one of the two main differences between conjectures and singleton properties [72]. In addition, it complies with the usual RDF semantics with a minimal increase of triples (one new RDF triple is added for every conjectured triple, rather than, e.g., four as needed with reification).

[7]Conjectures (weak form) mint new predicates by creating URIs with a different prefix for every named graph (in these examples, `conj0001`, `conj0002`, etc.) and the same local part as the original predicate (e.g., `P128_carries`). While not technically necessary, it improves readability overall.

[8]This proposal heavily borrows from singleton properties, but, crucially, it does not change the semantics of predicates of external vocabularies, and it does not affect reasoning tasks

The use of minted predicates has other advantages. First, the original triples (i.e. those using the non-conjectural predicates) are *not* stated, thus fulfilling the basic requirement of expressing without asserting. Second, contradictory statements, each adopting a conjectural predicate derived from the same original predicate, are explicitly mapped to the original one, which ensures clear identification and facilitates querying for unsettled disputes (cf. statements in unsettled dispute in definition C5 in Section 1.4).

Conjectural data can be queried with unmodified SPARQL engines. For instance, Listing 4.9 shows a SPARQL query for all disputed subjects of Vermeer's painting. The query returns all subjects, namely: `:girl` (present in both graphs) and `:cupid` (present only in one graph).

```
SELECT DISTINCT ?visualitem
WHERE {
  GRAPH ?Conj {
    ?conjpredicate conj:isAConjecturalFormOf crm:P65_shows_visual_item.
    ?work ?conjpredicate ?visualitem . }
}
```

Listing 4.9: SPARQL query to return disputed subjects with Conjectures in weak form

4.2.2 Conjectures in strong form

The weak form of conjectures can be appreciated better by making explicit the underlying semantics of conjectural graphs. To do so, a prototypical extension of the syntax of Trig to support conjectures is proposed, where `GRAPH` is replaced with `CONJ` in front of a graph whose contents are expressed but not asserted.

Listing 4.10 shows the disputed statements claiming respectively the authenticity and the fraudulence of the "Donation of Constantine".

```
CONJ :donationStatement {
  :donationOfConstantine a :Authentic ;
    dct:date :4thCentury
  }

```

```
CONJ :vallaStatement {
  :donationOfConstantine a :Forgery ;
    dct:date :8thCentury
  }
```

Listing 4.10: The disputed claims about the authenticity of the "Donation of Constantine" represented in Conjectures strong form

The strong form avoids using conjectural predicates as in Listing 4.8, now represented in their original form but introduces the conjectural graph with the keyword `CONJ`. This simplifies writing and understanding conjectural statements, while the semantics and interpretation of the two forms are *exactly* the same [92], and therefore the two forms are interchangeable. All triples inside a strong form conjecture are considered conjectural (and vice versa) to guarantee full compatibility between weak and strong forms. All statements inside a conjectural graph in the weak form must be conjectural triples with their accompanying `conj:isAConjecturalFormOf` triple. This aspect is important when evaluating the efficiency of the Conjectures approach to EWA (see Section 4.3), which would imply bespoke software dependencies only when using the strong form (similar to RDF-star and SPARQL-star). Another key aspect is that conjectures do not use reification, *n*-ary relationships or *ad hoc* classes. Therefore, they are orthogonal and fully compatible with most other approaches.

4.2.3 Re-asserting a conjecture through a settle

Settled conjectures are introduced to handle settled disputes, a third type of named graphs that is at the same time conjectured *and* asserted, which can be defined as follows:

Definition 2 In weak form, a *settled conjecture* $c1$ consists of two graphs: the first is the usual conjecture `c1` expressed as in Definition #1, and the second is a new graph `cc1` with all the triples in `c1` but with their original predicates, except for `conj:isAConjectualFormOf` triples, plus the triple (cc1, conj:settles, c1).

Following the Donation example, Listing 4.11 shows how the relevant community decision to carry Lorenzo Valla's claims can be represented as a settled claim (:vallaStatement graph), i.e., a settled conjecture in weak form.

The settled graphs allow the representation of both the conjectural triples (inside the usual conjectural graph :vallaStatement) as well as the same triples but completely asserted (inside the settled graph :SettledVallaStatement). In addition, the conj:settles relation connects the conjecture and its settlement, simplifying the exploration of relationships between disputes and their resolution.

```

GRAPH :vallaStatement {
    :donationOfConstantine conj0003:date :8thCentury;
        conj003:type :Forgery .
    conj0003:date conj:isAConjecturalFormOf dct:date .
    conj0003:type conj:isAConjecturalFormOf rdf:type .
}

GRAPH :settledVallaStatement {
    :donationOfConstantine dct:date :8thCentury ;
        a :Forgery
    :settledVallaStatement conj:settles :vallaStatement .
}
:settledVallaStatement prov:wasAttributedTo :lorenzoValla .
```

Listing 4.11: Settled conjecture in weak form

As shown in Listing 4.11, not only all statements in :vallaStatement have been re-established in their original form, but the graph :SettledVallaStatement also explicitly identifies its settlement. Listing 4.12 shows the same Listing as 4.11, but in strong form, using the new keyword SETT. The strong form, as before, carries exactly the same semantics as the weak form but shortens and clarifies the syntax, for instance, omitting the need to repeat the re-asserted triples explicitly.

```

SETT :vallaStatement {
    :donationOfConstantine dct:date :8thCentury ;
        a :Forgery .
}
:settledVallaStatement prov:wasAttributedTo :lorenzoValla .
```

Listing 4.12: Settled claims about the "Donation of Constanine" authenticity in strong form

The representation of the dual nature of settled conjectures, which are simultaneously conjectured *and* asserted, is simplified in the strong form using the keyword `SETT`. For the sake of clarity, the term `SETT` is compatible with the notation shown in Listing 4.13 where a settled conjecture is represented by a conjectural graph (`CONJ :vallaStatement`) and a plain RDF 1.1 named graph (`GRAPH :settledVallaStatement`) and their relation is maintained by the property `conj:settles`.

```

CONJ :vallaStatement {
    :donationOfConstantine dct:date :8thCentury ;
        a :Forgery
}

GRAPH :settledVallaStatement {
    :donationOfConstantine dct:date :8thCentury ;
        a :Forgery.
    :settledVallaStatement conj:settles :vallaStatement .
}
:settledVallaStatement prov:wasAttributedTo :lorenzoValla .
```

Listing 4.13: Full form of settled claims about the "Donation of Constanine" authenticity in strong form

As illustrated in Listings 4.11 and 4.13, a settled conjecture inherently settles its contents. In the case of the Valla example, his opinion is established as the preferred resolution of the dispute. Moreover, a broader community consensus can be represented as a newly settled graph, which not only settles its own contents by design but also explicitly settles other opinions (the opinions of Lorenzo Valla and Nicholas of Cusa, both of whom identified the document as a forgery). As shown in Listing 4.14, the settlement relation between the community consensus, Valla's and Cusanus' opinion is explicitly stated in the RDF notation through the property `conj:settles`. The explicit settlement requires that all triples in conjectural graphs (`:VallaStatement` and `:CusanusStatement`) are also present in the community consensus graph (`:communityConsensus`).

```
CONJ :vallaStatement {
    :donationOfConstantine dct:date :8thCentury ;
        a :Forgery
}
:vallaStatement prov:wasAttributedTo :lorenzoValla .

CONJ :cusanusStatement {
    :donationOfConstantine a :Forgery
}
:cusanusStatement prov:wasAttributedTo :nicholasOfCusa .

SETT :communityConsensus {
    :donationOfConstantine dct:date :8thCentury ;
        a :Forgery .
    :communityConsensus conj:settles :vallaStatement, :cusanusStatement
}
```

Listing 4.14: Community consensus supporting Valla's opinion about the "Donation of Constanine" in strong form

4.2.4 Conjectures and SPARQL entailment

It has been argued [110] (See Section 3.8 *Relationship with SPARQL entailment regime*) that the relation between graphs semantics and SPARQL entailment is the most important aspect to be aware of when working with the model-theoretic semantics of named graphs. A graph should be considered asserted (i.e., *true*) if a SPARQL engine recognises the graph as true using an ASK query and false otherwise. This could also be considered a practical test to verify whether Conjectures fits the current RDF environment and tools.

To demonstrate that Conjectures successfully implement the intended semantics (fourth semantics of named graphs as already introduced at the beginning of this section), verification is conducted to ensure that according to the SPARQL entailment regime [55], that (a) *statements inside conjectural graphs are not considered true*, and (b) *statements inside non-conjectural graphs (plain graphs as well as settled graphs) are considered true*.

Following the "Donation of Constantine" example, an ASK query to identify whether graphs exist stating that the Document was created in the 8th century would correctly return `True` (Listing 4.15), since this is expressed explicitly inside a settled conjecture (Listing 4.12 and 4.13).

```

ASK WHERE {
    GRAPH ?graph {
        :donationOfConstantine dct:date :8thCentury.
    }
}
```

Listing 4.15: SPARQL query to retrieve settled claims

Conversely, an ASK query (Listing 4.16) to identify whether graphs exist stating that the Document has been created in the 4th century would correctly return `False` since this statement is shielded by the conjectural graph in the strong form and conjectural predicate (e.g., `conj0001:date`.) in the weak form.

```

ASK WHERE {
    GRAPH ?graph {
        :donationOfConstantine dct:date :4thCentury
    }
}
```

Listing 4.16: SPARQL query to retrieve disputed claims

A parser of conjectures, including a converter between strong and weak forms, is available online[9] along with a number of SPARQL queries for testing purposes. Currently, any dataset, including conjectural and settled graphs in the strong form, is transformed into the corresponding plain RDF 1.1 dataset in the weak form.

4.2.5 Expressing complex knowledge with Conjectures

The representation of critical debate and the evolution of information is a dynamic process embodied in the theories proposed by scholars. These theories serve as frameworks that influence and reshape various factors within a given topic. For instance, consider the evolving information about the painting "Girl Reading a Letter at an Open Window". Following its restoration, numerous metadata associated with the painting, such as the image itself, its depictions, and its symbolic meanings, have changed. These changes stem from the same underlying reason (the restoration): sharing provenance information.

In data representation systems like Wikidata, reification occurs at the level of individual triples. Each triple is reified independently, and a unique identifier is assigned along with its provenance, which results in the repeated attachment of provenance information across all triples. This method can lead to redundancy and complexity in managing interconnected claims. Listing 4.17 shows the claims available in Wikidata, which stores information about the painting before and after the restoration. In total, 36 triples are counted to represent 6 statements. Wikidata reification ideally allows retrieving all triples sharing the same provenance information; it falls short in real-case data. For example, in Listing 4.17, the provenance information about the image (annotated with media legend `pq:P2096` "after 2021 restoration revealing Cupid painting") does not match the provenance information about the depiction of Cupid (annotated with cause restoration `pq:P828 wd:Q217102`). Besides this, most claims do not have any source[10]. Therefore, retrieving such information

[9] `http://conjectures.altervista.org/`
[10] `https://www.wikidata.org/wiki/Help:Sources`

using the same provenance triples is impossible.

```
# The image of the painting before the restoration
wd:Q700251 p:P18 s:Q700251-0CCDC965-AEA0-4E13-8FCD-53CAE5555E8D .
s:Q700251-0CCDC965-AEA0-4E13-8FCD-53CAE5555E8D a wikibase:Statement ;
  wikibase:rank wikibase:NormalRank ;
  ps:P18 <image_link_before_restoration> ; # image url
  pq:P582 "2017-01-01T00:00:00Z"^^xsd:dateTime . #end time 2017

# The image of the painting after the restoration
wd:Q700251 p:P18 s:Q700251-07948323-4b2b-4c44-3c3d-6694e1c15857 .
s:Q700251-07948323-4b2b-4c44-3c3d-6694e1c15857 a wikibase:Statement, wikibase:
    BestRank ;
  wikibase:rank wikibase:PreferredRank ;
  ps:P18 <image_link_after_restoration> ; # image url
  pq:P2096 "after 2021 restoration revealing Cupid painting"@en ; # media legend
  pq:P580 "2021-09-10T00:00:00Z"^^xsd:dateTime . #start time 2021

# The painting depicts a girl
wd:Q700251 wdt:P180 wd:Q3031.
wd:Q700251 p:P180 s:Q700251-a974521e-4faa-308c-c437-e89ff671d427 .
s:Q700251-a974521e-4faa-308c-c437-e89ff671d427 a wikibase:Statement, wikibase:
    BestRank ;
  wikibase:rank wikibase:NormalRank ;
  ps:P180 wd:Q3031 . # depicts a girl

# The painting depicts a letter
wd:Q700251 wdt:P180 wd:Q133492.
wd:Q700251 p:P180 s:Q700251-d608244e-44fe-bf7f-2826-bab6f5b94e44 .
s:Q700251-d608244e-44fe-bf7f-2826-bab6f5b94e44 a wikibase:Statement,
    wikibase:BestRank ;
  wikibase:rank wikibase:NormalRank ;
  ps:P180 wd:Q133492 . # depicts a letter

# The painting depicts a window
wd:Q700251 wdt:P180 wd:Q35473.
wd:Q700251 p:P180 s:Q700251-c4d34e6d-4a8c-a65d-ccf6-02f518af04bf .
s:Q700251-c4d34e6d-4a8c-a65d-ccf6-02f518af04bf a wikibase:Statement,
    wikibase:BestRank ;
  wikibase:rank wikibase:NormalRank ;
  ps:P180 wd:Q35473 . # depicts a window
```

```

# The painting depicts a Cupid (from 2021 and with cause conservation)
wd:Q700251 p:P180 s:Q700251-83546b41-4ffe-a4bf-2061-ffbe9e06cf92 .
s:Q700251-83546b41-4ffe-a4bf-2061-ffbe9e06cf92 a wikibase:Statement, wikibase:BestRank ;
  wikibase:rank wikibase:NormalRank ;
  ps:P180 wd:Q5011 ; # depicts a cupid
  pq:P580 "2021-01-01T00:00:00Z"^^xsd:dateTime ; # start time 2021
  pq:P828 wd:Q217102 . # has cause conservation
```

Listing 4.17: Represenation of evolving information of "Girl reading a letter at an open Window" with Wikidata RDF statements

Contrarily, with Conjectures, the entire graph — stores interrelated claims originating from a single source — possesses a single identifier. This allows the provenance to be associated with the interrelated claims, streamlining the representation process. In Listing 4.18, the same information about the painting is presented using Conjectures. This representation implies 9 quadruples, having a triple count reduction of 75%[11]. Using a SPARQL query is therefore easy to retrieve the whole graph both querying provenance information or the all triples stored in the graph.

```
GRAPH :factual_data {
  wd:Q700251 wdt:P180 wd:Q3031. # The painting depicts a girl
  wd:Q700251 wdt:P180 wd:Q133492. # The painting depicts a letter
  wd:Q700251 wdt:P180 wd:Q35473. # The painting depicts a window
}

# Claim about the painting before the restoration
CONJ :beforeRestoration {
  wd:Q700251 wdt:P18 <image_link_before_restoration>. # image url
}
:beforeRestoration pq:P582 "2017-01-01T00:00:00Z"^^xsd:dateTime. # end time 2017

```

[11] The following calculations are performed to determine the percentage of reduction in the triple count:

- Reduction $=$ First count $-$ Second count
- Percentage reduction $= \frac{\text{Reduction}}{\text{First count}} \times 100$

```
# Claim about the painting after the restoration
SETT :afterRestoration {
  wd:Q700251 wdt:P18 <image_link_after_restoration>;
  wdt:P180 wdt:Q217102 # painting depicts a Cupid
}
:afterRestoration pq:P580 "2021-01-01T00:00:00Z"^^xsd:dateTime ; # start time 2021
  pq:P828 wd:Q217102 . # cause conservation
```

Listing 4.18: Represenation of evolving information of "Girl reading a letter at an open Window" with Conjectures

Moreover, RDF-star introduces techniques to encapsulate quoted triples that align with such representation, as using RDF collections, as illustrated in Listing 4.19. However, it is impossible to assign an identifier to these collections[12].

```
(
    # painting depicts a Cupid
    << wd:Q700251 wdt:P180 wdt:Q217102 >>
    << wd:Q700251 wdt:P18 <image_link_after_restoration> >>
)   pq:P828 wd:Q217102 . # cause conservation
    pq:P580 "2021-01-01T00:00:00Z"^^xsd:dateTime ; # start 2021
```

Listing 4.19: Represenation of evolving information of "Girl reading a letter at an open Windows" with Trig-star

Alternatively, Trig syntax can be combined with RDF-star syntax to express the content of named graphs without asserting it, making it suitable for representing interrelated claims, as demonstrated in Listing 4.20[13]. While all asserted triples in `:factual_data` remain the same as Listing 4.18 (and are therefore not repeated here), non-asserted triples (found in the graph `:beforeRestoration`) are expressed using nested RDF-star quoted triples within the named graph syntax. As for the triples stored in `:afterRestoration`, they are asserted but also represented as quoted triples, each associated with a set of contextual information, repeated for every quoted

[12] This serialisation has not been included in the efficiency evaluations, as it is not supported by the GraphDB triplestore for SPARQL queries

[13] Additionally, several example on how to use Trig-star are available at https://w3c.github.io/rdf-star/tests/trig/syntax/manifest.html#trig-star-1

statement.

```

# Claim about the painting before the restoration
GRAPH :beforeRestoration {
  << wd:Q700251 wdt:P18 <image_link_before_restoration> >>
  :beforeRestoration pq:P582 "2017-01-01T00:00:00Z"^^xsd:dateTime. # end time 2017

# Claim about the painting after the restoration
GRAPH :afterRestoration {
  << wd:Q700251 wdt:P18 <image_link_after_restoration> >>
  :afterRestoration pq:P580 "2021-01-01T00:00:00Z"^^xsd:dateTime ; # start time 2021
    .
  << wdt:P180 wdt:Q217102 >> :afterRestoration pq:P580 "2021-01-01T00:00:00Z"^^xsd:
    dateTime ; # start time 2021, painting depicts a Cupid
}
```

Listing 4.20: Represenation of evolving information of "Girl reading a letter at an open Windows" with Trig-star

This process successfully groups statements and addresses EWA with named graphs. However, this approach involves reification of reified data (RDF-star over named graphs syntax), adding layers of complexity and does not solve the redundancy in contextual data addition at a statements level (in this case, for each quoted triples). Named graphs with Conjectures present a more elegant solution by design to represent such complex claims within the RDF framework. Encapsulating all claims within named graphs enables effective management and streamlining of the provenance and evolution of these interconnected claims in scholarly debates and knowledge evolution.

4.2.6 Discussion

Analysing existing solutions has revealed limitations that hamper a correct and shareable approach (Section 4.1). A real-world scenarios that motivate the necessity of clear mechanisms for EWA has been proposed. On top of this, this section discusses, RQ2a - *How do existing reification approaches express without asserting?* and RQ2b - *Is there evidence that existing approaches for EWA are ineffective?*

The comparison of existing ontology-independent solutions to achieve EWA (Section 4.1) highlighted how such works do not provide a clear answer to the problem, leveraging results of recent works on uncertainty, provenance, and meta-knowledge with RDF technologies [5, 44, 100, 95, 78]. Established solutions were reviewed, showcasing their merits and limitations. Building upon these frameworks, a new approach was developed to address the problem of EWA more effectively.

First, ontology-dependent solutions [33, 74, 62], were excluded as a one-size-fits-all solution since these do not allow to express without asserting. On the other hand, the emphasis of most ontology-independent solutions is on provenance and meta-knowledge [44, 95, 100], regardless of the logical status of statements annotated with provenance, possibly generating EWA as a side-effect.

Secondly, while efficiency is often discussed as the main criterion for comparison [44, 78], their effectiveness in expressing claims that depend on a specific context has been overlooked. An intuitive definition of effectiveness has been proposed, and existing proposals have been compared in this light. This aspect has been addressed by considering diverse real-world scenarios concerning the humanities, which prevents seeking an over-engineered solution.

Third, this comparison does not neglect practical considerations like compatibility with RDF 1.1 semantics, dependency on bespoke technologies, and the potential expensiveness (in terms of additional triples) derived from mapping solutions back to plain RDF. Current proposals for EWA require either technological constraints [5] or expensive mappings for additional triples [56, 54].

Taking these aspects into account, RQ2c poses the question: *Is there a way to represent EWA claims that can overcome such deficiencies?* Conjectures provide a solution compliant with RDF 1.1 semantics (weak form) and is more or as concise as existing proposals in their compatible version.

Existing solutions were extended to preserve RDF 1.1 semantics as a priority, hence the proposal of a *strong form* to express without asserting. Nonetheless, a path to overcome the ambiguity in the semantics of named graphs should be introduced. Therefore, a concise solution is proposed to consistently represent (in the same dataset) graphs with three different semantics: plain graphs (with the same semantic ambiguity as before), definitely non-asserted graphs (conjectures), and definitely asserted graphs (settled conjectures). This also proposes a scholarly interpretation of these semantics to match undisputed claims, disputed claims and settled disputes, which mimic critical inquiry evolution. Concerning methods discussed in Section 3.3, the effectiveness of Conjectures is the most representative. In particular, conjectures are the only method that satisfies all the following conditions at the same time: (1) successfully achieves EWA, to express without asserting contents of named graphs, (2) to be mapped to RDF with the lowest impact on the number of additional triples (+1), (3) to make explicit the semantics associated to named graphs.

A noteworthy difference between conjectures and most other approaches listed in Section 3.3 is that conjectures do not use additional entities created only to associate provenance and other types of meta-knowledge to RDF triples. Whether they are some type of reified statements (e.g., reification or Wikidata [107]), newly minted blank nodes (e.g., rdf-star and N3) or a variety of typed entities listed in an ontology (most of the n-ary relationships such as in CIDOC-CRM [37]), these approaches all create some new entity that becomes the subject of any subsequent meta-knowledge triples such as provenance, confidence, etc. Exception to this approach are basically RDF+ (which provides a generated identifier for every triple, expanding the quad to a quintuple and getting outside of the basic RDF model), and singleton properties, using newly minted one-time predicates for this purpose.

Like singleton properties, conjectures in their weak form mint one-time predicates associated via a special predicate (`isAConjecturalFormOf` vs. `singletonPropertyOf`) to the old ones, but there are two relevant differences that set them apart:

1. `singletonPropertyOf` is a subproperty of `rdf:type`, which leads to represent original predicates as classes and newly minted predicates as instances of these classes, creating a complicated pattern of original predicates being at the same time *both* predicates *and* classes. On the contrary, the predicate `isAConjecturalFormOf` makes no such assumption, and new predicates have the same ontological characteristics of the original ones (same domain, range, constraints) plus the constraint of being used exactly once for the specific subjects and objects of the triples they are replacing the predicates of.

2. With singleton properties, the basic purpose of the newly minted predicate is to generate a new URI to be used as the subject for meta-knowledge triples such as provenance. On the contrary, conjectures use plain named graphs for associating meta-knowledge to triples, and reserve newly minted predicates to only tell apart EWA claims from actual assertions. This approach restores named graphs, an integral part of the basic RDF 1.1 model, as the duly proxy for meta-knowledge annotations, and create a clear and sharp distinction between conceptual tools for meta-annotating triples, and those for expressing their assertedness.

The conjectures current formal model assumes that all triples are to be interpreted as conjectural in a conjectural graph. Has been previously mentioned (in particular, in the description of RDF statements in Section 2.2, Chapter 2) that domain-dependent constraints and uncertainty characteristics affect the amount of statements that can be expressed as conjectures. Uncertainty models [62, 33], argumentation mining [97, 22], and probabilistic models applied to ontologies [64] are different examples of the study of logical status and uncertainty of web data, which can be helped by shareable strategies and design choices on which named graphs to create and how to characterise their logical status case by case.

Finally, Conjectures allow for effective representations and queries when evolving knowledge is present. Retrieving claims based on their logical status could be difficult if no specific EWA methodology has been adopted. Namely, SPARQL queries could rapidly become expensive if they had first to retrieve all claims and then filter out those of the wrong logical status. Consider, for instance, the task of finding all current or obsolete interpretations of works of art with multiple interpretations. When no EWA approaches have been adopted, a query for the latest one could be as complex as the ones in Listing 4.21 and 4.22, where the efficiency of the custom FILTER functions is crucial to the speed of the whole query.

```
SELECT DISTINCT ?painting ?current
WHERE {
  ?painting crm:P128_carries ?current .
  ?event a crm:E11_Modification ;
    crm:P123_resulted_in ?painting ;
    crm:P4_has_time-span ?span .
  FILTER (my:includes(?span, now())
}
```

Listing 4.21: Querying for current interpretations with plain CIDOC-CRM

```
SELECT DISTINCT ?painting, ?obsolete
WHERE {
  ?painting crm:P128_carries ?obsolete.
  ?event a crm:E11_Modification;
    crm:P123_resulted_in ?painting;
    crm:P4_has_time-span ?span.
  FILTER (
  my:doesnotinclude(?span, now()) }
```

Listing 4.22: Querying for obsoleted interpretations with plain CIDOC-CRM

Using conjectures as in Listing 4.8, on the other hand, allows much simpler and faster queries such as the ones in Listings 4.23 and 4.24, since The logical status of claims is directly accessed, distinguishing between current and obsolete attributions instead of making numerical operations over the values of the `crm:P4_has_time-span` predicate.

```
SELECT DISTINCT ?painting ?current
WHERE {
  GRAPH ?painting {
    ?painting crm:P128_carries ?current .
    }
}
```

Listing 4.23: Querying current interpretations over conjectures in weak form

```
SELECT DISTINCT ?painting ?obsolete
WHERE {
  GRAPH ?painting {
    ?painting ?carries ?obsolete .
    ?carries conj:isAConjecturalFormOf crm:P128_carries . }
  FILTER NOT EXISTS {
    ?painting crm:P128_carries ?obsolete.
    }
}
```

Listing 4.24: Querying obsolete interpretations over conjectures in weak form

The approach presented here only searches asserted or not asserted triples (using, respectively, the original predicate and the conjectural form of the original one), which can be applied in a number of different situations.

4.3 Efficiency assessment of EWA approaches

As mentioned at the beginning of this chapter, efficiency is a central metric when establishing a new RDF solution. To my knowledge, no prior study has evaluated EWA approaches to express critical debate and knowledge evolution. This section will showcase the efficiency evaluation of a selection of ontology-independent reification methods to achieve EWA with RDF claims [40, 72, 54, 108] and compares them with Conjectures in both strong and weak form.

4.3.1 Experiment setup

This section outlines the experimental setup, focusing on the data acquisition, scaling, and conversion of a JSON-formatted Wikidata dump into RDF. It details the hardware and software configurations employed for storing and retrieving the produced datasets via triplestore. Moreover, it describes the customisation of the SPARQL endpoint to facilitate support for the Conjectures strong form framework. Therefore, a set of metrics has been adopted to measure the efficiency of reification methods.

Data Acquisition, scaling and conversion. The dataset on which these experiments have been run is composed as follows and has been named D3[14]:

- **Art**: A thematic set of claims about 300k artwork entities in Wikidata (i.e., painting, manuscripts, books). This corresponds to about 10% of all artwork entities currently present inside Wikidata.
- **Random**: After considerable deliberation, it was concluded that introducing a form of entropy into the dataset would enhance its representativeness. This dataset contains the claims of 300k Wikidata random entities.
- **Dummy**: A selection of dummy statements regarding the artwork attributions (represented by the property `wdt:P50` and `wdt:P170` and including from 1 to 4 authors in each claim and the source of the claim) and artworks locations (represented by the property `wdt:P276`, including 1 possible location, time constraints and source) has been created[15]. Those new statements contain dummy arbitrary information ranked as deprecated and therefore non-asserted to represent alternative or historical claims to those contained in Art dataset. This design choice was made to increase the number of conjectural statements in the final dataset.

An excellent way to evaluate an algorithm's performance is to observe how it responds to variations in input size [78]. The whole subset of artwork entities, related individuals (attributed authors) and locations have been downloaded

[14]The dataset [82] is available at `https://zenodo.org/records/10044574`

[15]The choice of adding the dummy claims is that of non-asserted statements in the Wikidata dump was circa 1%, a low figure for this experiment

from Wikidata. This dataset, called D4, comprises about 3,5 million artwork entities and 188 thousand related entities (humans and locations). It has not been used for this comparison due to the excessive number of timeouts in many of the queries and methods used in the test. Thus, the dataset has been scaled logarithmically in three further sizes:

- Dataset D3 : D3 is obtained by extracting one tenth of the data in D4 (D3 = D4/10).
- Dataset D2 : D2 is obtained by extracting one tenth of the data in D3 (D2 = D3/10).
- Dataset D1 : D1 is obtained by extracting one tenth of the data in D2 (D1 = D2/10).

Given state of the art regarding reification methods to express without asserting (See Section 4.1) a set of methods for our analysis has been selected, namely Singleton properties [72], named graphs [15] (using Wikidata rankings to decide whether a triple is asserted or not), Wikidata [40] and the recent RDF-star [54] approach. The Wikidata JSON files have been converted into the six selected reification methods through automatic scripts[16], while the dataset documentation is available at `https://github.com/conjectures-rdf/expressing-without-asserting-efficiency-datasets`.

Wikidata designed three rankings (described in Section 3.1.1) to mark statements (S), namely **Preferred** (S_P) to mark explicit consensus, **Deprecated** (S_D) to mark superseded claims and **Normal** (S_N) given as the default value whose logical status may change depending on competing claims. Rankings decide the logical status of claims and whether they are asserted ($S(A)$) or not ($S(NA)$). Specifically, for these reification methods, the representation

[16] In particular, the data conversion has been performed using a web-app developed in *node.js* for this purpose. The web-app make use of *HandleBars.js* templates. The source code and the application documentation is available at `https://github.com/conjectures-rdf/wikidata-converter-json-to-rdf`

of ($S(NA)$) allows to achieve EWA. Therefore, the datasets have been constructed satisfying the following logical rules through the use of Handlebar.js templates[17]:

- **Rule a:** If a statement (S) is **Preferred** (S_P) it is asserted ($S(A)$). Therefore, $S_P \Rightarrow S(A)$.

- **Rule b:** If a statement (S) is **Deprecated** (S_D), it is **non-asserted** ($S(NA)$). Therefore, $S_D \Rightarrow S(NA)$.

- **Rule c:** If a statement (S) is **Normal** (S_N), it is **asserted** ($S(A)$)[18].

- **Rule d:** Given at least two concurring statements S and S'[19] where a statement is **Normal** (S_N) and the other is **Normal** (S'_N), they are both **asserted**. Therefore, $\{S_N, S'_N\} \Rightarrow \{S(A), S'(A)\}$.

- **Rule e:** Given at least two concurring statements S and S' where a statement is **Normal** (S_N) and the other is **Preferred** (S'_P), they are respectively **non-asserted** statement ($S(NA)$) and a **asserted** statement ($S(A)$). Therefore, $\{S_N, S'_P\} \Rightarrow \{S(NA), S'(A)\}$.

- **Rule f:** Given at least two concurring statements S and S' where a statement is **Normal** (S_N) and the other is **Deprecated** (S'_D), they are respectively **asserted** ($S(A)$) and **non-asserted** ($S(NA)$). Therefore, $\{S_N, S'_D\} \Rightarrow \{S(A), S'(NA)\}$.

With methods such as Wikidata statements [40], RDF-star [54], and Singleton Properties [72] a reified statement (S(s, p, o)) is asserted ($S(A)$) if the same triple (s,p,o) exists in the dataset. Therefore, $S(s,p,o) \wedge \exists(s,p,o) \Rightarrow S(A)$. Contrarily, a reified statement (S(s, p, o)) is non-asserted ($S(NA)$) if the same triple (s,p,o) does not exist in the dataset. Therefore, $S(s,p,o) \wedge \neg\exists(s,p,o) \Rightarrow$

[17] `https://handlebarsjs.com/`

[18] This rule is applied only if no concurring statement exists. Otherwise, rules d, e and f are used

[19] Note that the definition of concurring statement is given by the presence of at least two reified statements ($S(s,p,o)$) and ($S'(s',p',o')$) where they address the same subject ($s = s'$) and predicate ($p = p'$)

$S(NA)$. Then, rules[a:f] were applied to convert the dataset to the top of this representation pattern, adding corresponding triples to assert statement contents when needed.

The representation of such rules is slightly different for named graphs and Conjectures. Named graphs do not make use of rules[a:f], therefore Wikibase rankings will be used in SPARQL queries. Conjectures framework does not use additional triples to assert the claim contents. Still, it represents undisputed, disputed and settled with the keyword introducing the graph ($CONJ(s,p,o) \Rightarrow S(NA)$, $GRAPH(s,p,o) \Rightarrow S(A)$, $SETT(s,p,o) \Rightarrow \{S(A), S'(NA)\}$), therefore a slightly different mapping has been performed so that all conditions expressed by the rules[a:f] were met. They are summarised as follows:

- **Rule a':** As provided by Wikidata: "The preferred rank is assigned to the most current statement or statements that best represent consensus", matching the definition of settled conjecture `SETT`, meeting the condition $S_P \Rightarrow S(A)$[20].

- **Rule b':** As provided by Wikidata: "The deprecated rank is used for statements that are known to include errors [...] or that represent outdated knowledge [...]", which can be represented with a conjectural graph `CONJ` and meeting the condition $S_D \Rightarrow S(NA)$.

- **Rule c':** As provided by Wikidata: "The normal rank is assigned to all statements by default. A normal rank provides no judgement or evaluation of a value's accuracy and currency and therefore should be considered neutral", which matches the definition of named graphs in the Conjectures framework `GRAPH` and meeting the condition $S_N \Rightarrow S(A)$.

- **Rule d':** In Conjectures, both (S'_N) and (S_N) are simply asserted graphs `GRAPH`, meeting the condition $\{S_N, S'_N\} \Rightarrow \{S(A), S'(A)\}$.

- **Rule e':** In Conjectures, (S_N) is a Conjecture `CONJ`, while (S'_P) is a settled conjecture `SETT` meeting the condition $\{S_N, S'_P\} \Rightarrow \{S(NA), S'(A)\}$

[20]Since a settled conjecture is both asserted and not asserted as defined in Section 4.2.3

- **Rule f':** In Conjectures, (S_N) is a settled conjectures `SETT`, while (S'_D) is a conjectural graph `CONJ`, meeting the condition $\{S_N, S'_D\} \Rightarrow \{S(A), S'(NA)\}$

The rules[a:f] and rules[a':f'] ensure that the logical status of claims is aligned between different reification methods. Table 4.2 provides some data about the datasets. At the end of this process, 18 new method-specific datasets are obtained. In other words, for each dataset Dn, $n \in [1,3]$, the following datasets are constructed:

Name	Serialisation	Reification	EWA	D3 stmts
Dn-Wikidata	Turtle	Wikidata statements	yes	66,768,937
Dn-rdfStar	TurtleX	RDF-star	yes	29,779,850
Dn-conjStrong	TriG	Conjectures strong form	yes	29,058,944
Dn-nGraphs	TriG	Named graphs	via ranking	28,896,268
Dn-conjWeak	TriG	Conjectures weak form	yes	29,199,650
Dn-Singleton	Turtle	Singleton properties	yes	55,325,270

Table 4.2: Datasets created ($n \in [1,3]$)

Hardware and software configuration. Tests have been run on a computer with processor Intel Core i5-8259U CPU @ 2.30GHz, RAM 32 GB, Windows 10 pro 64 bits, 1T SSD. The TriG and SPARQL parsers of the GraphDB engine[21] were modified to parse Conjectures in strong form[22]. The GraphDB configuration uses 28G Ram allocated to the application,[23][24], and

[21] `https://graphdb.ontotext.com/documentation/10.7/index.html`

[22] This customisation is documented in [65], which is the Bachelor thesis of Gerald Manzano who developed the tuning of GraphDB and which I co-supervised during this PhD, in particular see Section "Estensione di GraphDB" (4.2.3)

[23] `https://graphdb.ontotext.com/documentation/10.1/configuring-graphdb-memory.html`

[24] `https://graphdb.ontotext.com/documentation/10.2/getting-started.html#:~:`

10G cache size. A repository has been created for each dataset with inferences off, no rule set assigned, predicates list index enabled and (when possible) contexts enabled. All other parameters are left in their default values. Repositories are already running before their performance tests are executed.

Metrics. Reification methods has been compared on four major metrics. As introduced in the description of reification methods to introduce statements in RDF (see Section 2.2), these metrics are well-established regarding RDF quantitative analysis. The performance-related features of the reification methods under consideration should all be covered by those criteria, which should also give a clear picture of the benefits and drawbacks of each method.

- *Total number of triples in endpoint*: This value is particularly interesting since it makes it possible to assess the verbosity of each method.
- *Loading time*: Time consumed by each dataset to be uploaded in the SPARQL endpoint.
- *Dataset weight in triplestore*: The storage size of the dataset after it has been uploaded and stored in the triplestore.
- *Query execution time*: Response time on a selected set of queries. Each query is executed automatically ten times. The average value is then computed.

In particular, two sets of SPARQL queries (GQn and FQn) have been designed to get the query execution time. While GQn queries do not include any filter, FQn queries restrict the results only to paintings (`Q3305213`). Each query set comprises 6 queries assessing statements' logical statuses defined in definition C5 in Section 1.4. In particular, the queries retrieve the following topics: valid claims (Q1), disputed claims (Q2), disputed claims with

`text=the%20aforementioned%20icon.-,Configuring%20the%20JVM,Contents%2Fapp%2FGraphDB%20Desktop` their provenance/time (Q3), superseded claims (Q4), settled claims (Q5), and undisputed claims (Q6)[25].

Considering that authors' and locations' attributions provide a simple, yet effective use case to test RDF representation of EWA over our dataset, GQn and FQn have been then customised on retrieving authorship attributions (GQn-P170 and FQn-P170) and artworks' locations (GQn-P276 and FQn-P276) respectively by the use of Wikidata properties `P170` and `P276`. Each query set has been automatically run 10 times and the average times have been calculated.

Table 4.3 summarises the nature of the queries[26]. For the sake of the example, Listings 4.25, 4.26, 4.27, 4.28, 4.29 and 4.30 compare Wikidata statements and Conjectures strong form on general queries addressing attributions (GQn-P170[1:6])[27].

[25]The latter categorisation defined in definition C5 in Section 1.4, namely *Statements in unsettled dispute*, has not been considered in this work since Wikidata do not include open discussions in its data dump (cf. the discussion of Wikidata ranked statements in Section 3.1.1)

[26]All actual queries are available at `https://github.com/conjectures-rdf/EWA-efficiency` as well as the full set of results

[27]Note that queries syntax may seem redundant, but such choices are non-trivial as they have been tested to be the most efficient alternative

Query	**Predicate**	**Data selected by query**
GQ1	P170	All attributions of artworks that are currently considered valid
GQ1	P276	All locations of artworks that are currently considered valid
GQ2	P170	All attributions of artworks that have been disputed
GQ2	P276	All past and disputed locations of artworks
GQ3	P170	All attributions of artworks that have been disputed, *with provenance*
GQ3	P276	All past and disputed locations of artworks, *with date of move*
GQ4	P170	All superseded attributions of artworks
GQ4	P276	All past and superseded locations of artworks
GQ5	P170	All settled attributions of artworks
GQ5	P276	All current locations of artworks that were moved
GQ6	P170	All undisputed attributions of artworks
GQ6	P276	All locations of artworks that never moved
FQ1	P170	All attributions of paintings (Q3305213) that currently are considered valid
FQ1	P276	All locations of paintings (Q3305213) that are currently considered valid
FQ2	P170	All attributions of paintings (Q3305213) that have been disputed
FQ2	P276	All past and disputed locations of paintings (Q3305213)
FQ3	P170	All disputed attributions of paintings (Q3305213), *with provenance*
FQ3	P276	All past and disputed locations of paintings (Q3305213), *with date of move*
FQ4	P170	All superseded attributions of paintings (Q3305213)
FQ4	P276	All past and superseded locations of painting (Q3305213)
FQ5	P170	All settled attributions of paintings (Q3305213)
FQ5	P276	All current locations of paintings (Q3305213) that were moved
FQ6	P170	All undisputed attributions of paintings (Q3305213)
FQ6	P276	All locations of paintings (Q3305213) that never moved

Table 4.3: Overview of query types (GQn and FQn) for artworks' attributions and locations

This experiment evaluates the efficiency of ontology-independent reification methods in achieving EWA to represent knowledge evolution and critical debate. Specifically, the SPARQL queries are used to assess retrieval response times. For the purpose of query design, rankings were excluded (except for named graphs, where it was the only viable solution to mimic the achievement of EWA) since rankings are inherently ontology-dependent (part of the Wikibase datamodel). Instead, the queries focus on retrieving concurring statements and their logical status (asserted or not asserted) to derive their results when needed.

Q1 (Listing 4.25) retrieves all attributions considered valid and therefore asserted. In Wikidata, it is represented as a plainly asserted triple (`?artwork wdt:P170 ?artist`) and no reified statement is needed. The same result is retrieved with Conjectures by querying the same triple pattern within the basic named graph pattern (`GRAPH`).

```
SELECT DISTINCT ?artwork ?artist
WHERE {
    ?artwork wdt:P170 ?artist
    }
```

```
SELECT DISTINCT ?artwork ?artist
WHERE {
    GRAPH ?g {
    ?artwork wdt:P170 ?artist
    }
}
```

Listing 4.25: Q1 query addressing all valid claims with Wikidata statements (left) and Conjectures strong form (right)

Q2 (Listing 4.26) retrieves all disputed claims, specifically all attributions that have been challenged, regardless their resolution (e.g., settled or superseded). In Wikidata, this involves retrieving attributions (`?statement`) contingent upon the existence of a concurring disputed statement (`?statement2`, non-asserted and specified by the `FILTER NOT EXISTS` condition). Disputed claims are retrieved by querying conjectural graphs (`CONJ`, therefore non-asserted), including settled conjectures, which are asserted graphs and non-asserted conjectures by definition (cf. settled conjectures in Section 4.2.3).

```
SELECT DISTINCT ?artwork ?author
WHERE {
    ?artwork p:P170 ?statement.
    ?statement ps:P170 ?author.
    ?artwork p:P170 ?statement2.
    ?statement2 ps:P170 ?author2.
    FILTER NOT EXISTS {
        ?artwork wdt:P170 ?author2
    }
}
```

```
SELECT DISTINCT ?artwork ?artist
WHERE {
    CONJ ?g {
    ?artwork wdt:P170 ?artist
    }
}
```

Listing 4.26: Q2 query addressing all disputed claims with Wikidata statements (left) and Conjectures strong form (right)

Q3 (Listing 4.27) retrieves the same information as Q2 (Listing 4.26). Additionally, it includes the sources (`?sources`) associated with the retrieved attributions (`?artwork` and `?author`), which are matched using the property `pq:P248, source`, as an instance of contextual information.

```
SELECT DISTINCT ?artwork ?author ?
    source
WHERE {
    ?artwork p:P170 ?statement.
    ?statement ps:P170 ?author.
    ?artwork p:P170 ?statement2.
    ?statement2 ps:P170 ?author2.

    FILTER NOT EXISTS {
        ?artwork wdt:P170 ?author2, ?
    author
    }
    ?statement pq:P248 ?source
}
```

```
SELECT DISTINCT ?artwork ?artist
WHERE {
    CONJ ?g {
    ?artwork wdt:P170 ?artist.
    }
    ?g pq:P248 ?source.
}
```

Listing 4.27: Q3 query addressing all disputed claims and their source with Wikidata statements (left) and Conjectures strong form (right)

Q4 (Listing 4.28) retrieves all superseded attributions. In Wikidata, this is represented as a claim (`?statement`) that is not asserted—meaning that a triple asserting the claim does not exist (`FILTER NOT EXISTS`). In Conjectures, superseded claims are retrieved via conjectural graphs (`CONJ`) that are not settled, also accomplished through the function `FILTER NOT EXISTS` in conjunction with the property shortcut `conj:settles`.

```
SELECT DISTINCT ?artwork ?artist
WHERE {
    ?artwork p:P170 ?statement.
    ?statement ps:P170 ?artist.
    FILTER NOT EXISTS {
        ?artwork wdt:P170 ?artist
    }
}
```

```
SELECT DISTINCT ?artwork ?artist
WHERE {
    CONJ ?g {
    ?artwork wdt:P170 ?artist.
    FILTER NOT EXISTS {?g2 conj:
    settles ?g}
    }
}
```

Listing 4.28: Q4 query addressing all superseded claims with Wikidata statements (left) and Conjectures strong form (right)

Q5 (Listing 4.29) retrieves settled attributions. In Wikidata, those are identified by retrieving the valid attributions (`?statement`) which are asserted and therefore matching the pattern `?artwork wdt:P170 ?author`. This is contingent upon the existence of a disputed statement (`?statement2`, which is non-asserted). In Conjectures, settled attributions are retrieved by querying settled graphs (`SETT`).

```
SELECT DISTINCT ?artwork ?author
WHERE {
    ?artwork wdt:P170 ?author.
    ?artwork p:P170 ?statement.
    ?statement ps:P170 ?author.
    ?artwork p:P170 ?statement2.
    ?statement2 ps:P170 ?author2.

    FILTER NOT EXISTS {
        ?artwork wdt:P170 ?author2
    }
}
```

```
SELECT DISTINCT ?artwork ?artist
WHERE {
    SETT ?g {
    ?artwork wdt:P170 ?artist.
    }
}
```

Listing 4.29: Q5 query addressing settled claims with Wikidata statements (left) and Conjectures strong form (right)

Q6 (Listing 4.30) retrieves undisputed claims. In Wikidata, this is retrieved by matching all valid attributions (`?artwork wdt:P170 ?author`) which never occur with a disputed concurring statement (`?statement2`, which is non-asserted as it matches the pattern `FILTER NOT EXISTS {?artwork wdt:P170 ?author2}`). In Conjectures undisputed claims are retrieved by querying all valid graphs (`GRAPH`) and filtering those which have been settled (`FILTER NOT EXISTS` and `conj:settles`).

```
SELECT DISTINCT ?artwork ?author
WHERE {
    ?artwork wdt:P170 ?author .
    MINUS {
        ?artwork p:P170 ?statement2.
        ?statement2 ps:P170 ?author2.
        FILTER NOT EXISTS {
            ?artwork wdt:P170 ?author2
        }
    }
}
```

```
SELECT DISTINCT ?artwork ?artist
WHERE {
    GRAPH ?g {
    ?artwork wdt:P170 ?artist.
    }
    FILTER NOT EXISTS {
        ?g conj:settles ?g2
    }
}
```

Listing 4.30: Q6 query addressing all undisputed claims with Wikidata statements (left) and Conjectures strong form (right)

4.3.2 Test results

Number of triples in endpoint. All existing reification methods add additional triples to the already existing ones to represent the necessary metadata (e.g. Singleton properties) or extend RDF 1.1 syntax (e.g. RDF-star). As shown in Figure 4.1 and in Table 4.4, named graphs are the method which uses reification with the lower number of triples, but with no explicit distinction between asserted and non-asserted graphs. While other surveyed methods (in particular, RDF-star, Wikidata statements and Singleton properties) use reification methods and assert each claim with an additional triple, Conjectures uses named graphs structure to express both statements and reification without adding additional triples resulting in the method to express without asserting with the lowest addition of triples.

	D1	D2	D3
nGraphs	325,169	3,557,846	28,896,268
rdfStar	336,220	3,706,221	29,779,850
singeton	614,788	6,813,637	55,325,270
Wikidata	753,395	8,360,954	66,768,937
conjWeak	328,769	3,598,925	29,199,650
conjStrong	326,079	3,568,341	29,058,944

Table 4.4: Dataset number of triples or quads

Loading Time. In the context of dataset D1 and D2, Conjectures in the strong form remain competitive with the most efficient methods, notably RDF-star, and outperform Wikidata statements and Singleton properties as shown in Figure 4.1 and Table 4.5. However, the loading times increase in D3. This performance discrepancy is attributed to the triplestore's parser method for recognizing conjectural data. Specifically, the process of checking each resource's presence in a collection during loading contributes to the observed delays. In essence, the loading time of the dataset increases proportionally with the quantity of non-asserted triples (conjectures).

	D1	D2	D3
Wikidata	32	347	2738
rdfStar	19	289	1516
conjStrong	22	458	24255
nGraphs	17	203	1546
conjWeak	19	207	1693

Table 4.5: Datasets loading time in seconds

Dataset weights in triplestore. The Singleton method exhibits a storage size tenfold greater than alternative approaches, with Conjectures in their weak form and Wikidata occupying intermediate positions. RDF-star, Conjectures in their strong form, and named graphs demonstrate similar sizes as shown in Figure 4.1.

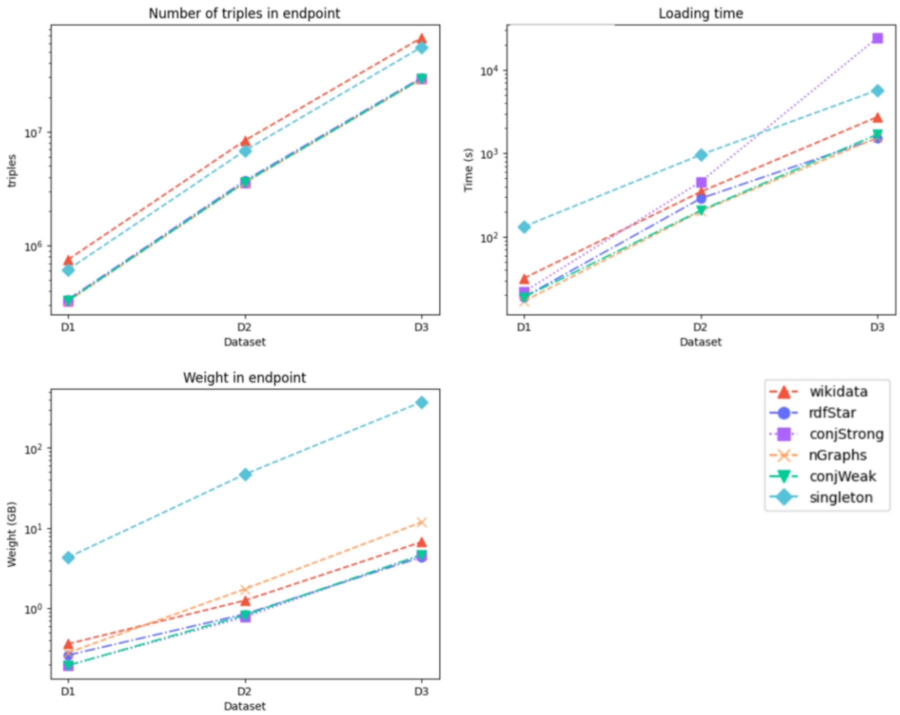

Figure 4.1: Dataset measurements for the surveyed reification methods. In particular, number of triples in endpoint (top-left), dataset loading time (top-right) and dataset weight in endpoint (bottom-left)

Query Execution Time. The time response average for each dataset seems to increase linearly for each surveyed dataset Dn, $n \in [1,3]$. For this reason, Figure 4.2 provides the snapshot of the execution time of queries GQn and FQn on attributions and locations only on dataset D3.

As illustrated in Figure 4.2, the response times obtained from the execution of general queries (GQn) on dataset D3 show that the strong form of Conjectures is less efficient than other methods when retrieving asserted data, particularly in retrieving valid claims, superseded and undisputed claims (queries GQ1, GQ4 and GQ6) for both creators and locations. However, Conjectures in weak

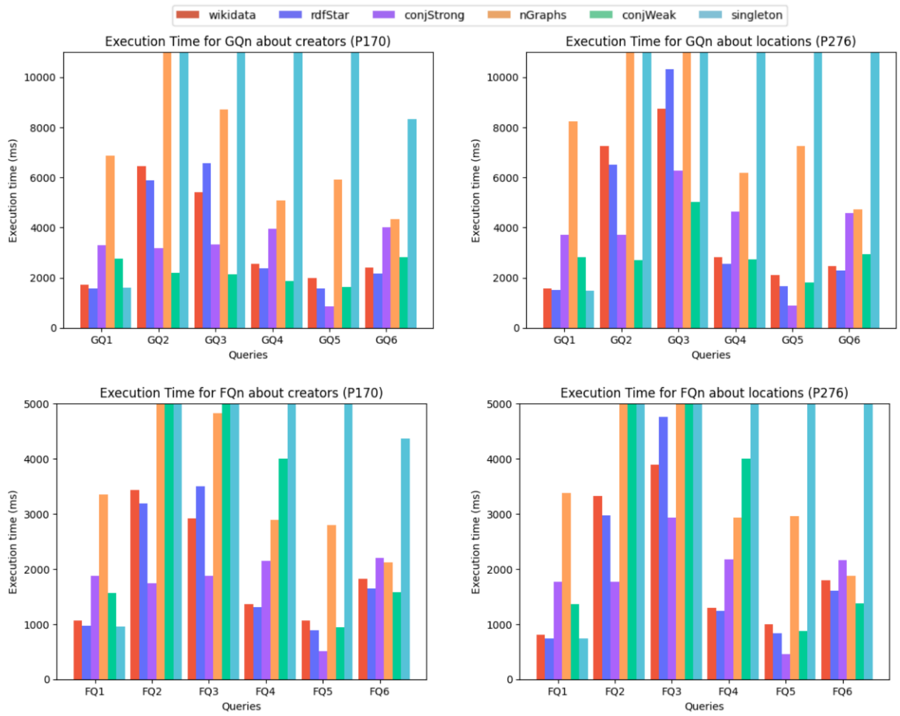

Figure 4.2: Time responses for queries sets GQn and FQn on creators and locations run against D3 dataset

and strong forms outperform other surveyed methods in retrieving disputed statements with and without provenance information (queries GQ2 and GQ3) and settled (query GQ5) for both locations and creators.

Similar to what was observed in GQn, Conjectures are less efficient in retrieving valid claims (FQ1). On the contrary, Conjectures strong form is the most efficient method in retrieving disputed claims with and without provenance (queries FQ2 and FQ3) and settled claims (FQ5). In the remaining queries, superseded (FQ4) and undisputed claims (FQ6), Conjectures still maintain competitive times with the rest of the methods. Conjectures in strong form, in particular, address the significant increase in response times for weak form queries FQn[3:8]. Essentially, a notable improvement in the performance of the strong form has been detected, proving to be the most efficient method in half of the selected queries and, in the remaining ones, a valid competitor.

4.3.3 Discussion

Section 4.3 compares the efficiency of existing well-established reification methods, evaluating their performance in EWA. In particular, the section addresses RQ2d - *What is the most efficient way to implement EWA?* The efficiency of EWA mechanisms is compared between several reification methods (Wikidata, RDF-star, named graphs, Singleton properties) and the novel Conjectures approach (weak and strong form) on four major metrics (number of triples, loading time, dataset weight and query execution time). Among the most efficient methods as RDF-star and Wikidata statements, the strong form of Conjectures exhibits notable performance gains, particularly in retrieving claims about debates (e.g., disputed claims with and without provenance information and settled claims).
Overall, several trends concerning the surveyed methods' efficiency can be seen: Singleton properties are systematically slower than the others, while named graphs and Conjectures in weak form performs at an intermediate level about the fastest methods, Wikidata, Conjectures in strong form and RDF-star. Conjectures in strong form form also outperforms RDF-star in many queries where the specifics of disputed attributions and past locations become meaningful. Conjectures in strong form is the quickest method for expressing debates (disputed claims, GQ2 and FQ2, disputed claims with provenance GQ3 and FQ3, settled claims GQ5 and FQ5), with a small loss in terms of performance for what concern asserted claims (valid claims GQn and FQn, undisputed claims GQ6 and FQ6) and superseded claims (GQ5 and FQ5).
Conjectures in strong form are also competitive regarding the number of triples and overall weight in the triplestore. It is competitive regarding loading time, but loading times show an interesting performance loss for large datasets that need further investigation.

Although this section addresses the efficiency of reification methods achieving EWA to represent evolving knowledge and critical debates, the analysis—particularly the query design—revealed several notable aspects of these methods that warrant further discussion.

As shown in the queries (Listings from 4.25 to 4.30) Wikidata, as the other reification methods like RDF-star and Singleton properties, represents valid (Q1) statements using reified triples, formulating an asserted statement as $S(s,p,o) \wedge \exists(s.p,o) \Rightarrow S(A)$[28]. In contrast, superseded claims (Q3) are defined as claims that lack an asserted triple, expressed as $S(s,p,o) \wedge \neg\exists(s,p,o) \Rightarrow S(A)$. Therefore, Q1 and Q3 are retrieved by simply assessing if the claim is asserted or not.
Disputed claims (Q2 and Q3), settled claims (Q5), and undisputed claims (Q6) are retrieved in SPARQL by comparing concurring statements as $Retrieve(S(s,p,o) \wedge S'(s,p,o'))$. Here, in this experiment, the definition of concurring claims is that statements S and S' share the same subject s and predicate p but can have different logical statuses (A or NA) and potentially different objects (o and o'). Therefore, queries Q2, Q3, Q5 and Q6 are achieved comparatively and can be retrieved if at least one concurring claim is disputed. For disputed claims (Q2 and Q3), all claims (both asserted and non-asserted) that have at least one concurring claim which is disputed (non-asserted) are retrieved. Settled claims (Q5) are identified by assessing all asserted claims when at least one concurring claim is disputed. Undisputed claims (Q6) are those asserted claims that lack any concurring claim that is disputed.
Conjectures achieve the assertion and non-assertion as the other reification methods with $GRAPH(s,p,o) \Rightarrow S(A)$ and $CONJ(s,p,o) \Rightarrow S(NA)$[29]. Additionally, the settled conjecture $SETT(s,p,o) \Rightarrow \{S(A), S'(NA)\}$ adds a third type of graph which is not directly aligned with other reification methods, but which helps in achieving the same query results for all queries in the test Q[1:6] for both the debate of attributions (P170) and evolving locations (P276).

[28]see data conversion rules in the experiment setup section (Section 4.3.1)
[29]see data conversion rules in the experiment setup section (Section 4.3.1)

This design choice reveals some intrinsic differences between Conjectures and traditional reification methods.

In cases where disputes are not explicitly recorded and only accepted statements are preserved, traditional reification approaches prove inadequate. This is because reification does not differentiate between undisputed claims (Q6) and settled (Q5); both are stored as asserted triples (see Rule *a* in Section 4.3.1). While Wikidata addresses this distinction through its use of Preferred rankings, these rankings were not incorporated into the queries used in this experiment. As previously stated, rankings are inherently ontology-dependent solutions, and although they were necessary for constructing the dataset, they fall outside the scope of this analysis. For example, the painting *Portrait of Dona Isabel de Requesens* (`Q29651096`)[30] is attributed to Giulio Romano as a settled claim, yet no alternative attribution is reported. The concept of settled conjectures captures this distinction in an ontology-independent manner, making it adaptable across any KG. Although this can be achieved by adding preferred rankings in SPARQL queries when needed (e.g., Q5 and Q6) via a UNION operation, such an approach would increase query complexity for reification methods such as Wikidata, RDF-star and Singleton Properties. Therefore, the slight discrepancy observed in the query results (less than 0.01%) has been accepted for this experiment.

Another interesting case is the occurrence of two claims reifying the same but with different qualifiers. In other words, they share the same content but with different contextual information. Consider, for instance, the case of the attribution of the painting *The Sultan Mehmet II* (`Q3937436`)[31] to Gentile Bellini. First, a normal-ranked claim states the attribution (explicitly marked as an attribution with the aid of a specific qualifier) to Gentile Bellini. This claim is not-asserted (see Rule *e* in Section 4.3.1). Besides this, a competing preferred ranked claim is reported, stating the authorship

[30] `https://www.wikidata.org/wiki/Q29651096`
[31] `https://www.wikidata.org/wiki/Q3937436`

to Gentile Bellini (and avoiding the qualifier). The historical attribution (the claim with normal rank) cannot be retrieved in a SPARQL query using traditional reification methods since they retrieve superseded claims (Q4) as $S(s,p,o) \wedge \neg\exists(s,p,o) \Rightarrow S(A)$. In this case, the result is not achievable since it does not satisfy the condition $\neg\exists(s,p,o)$ since s,p,o exists in KG as generated by the concurring claim (with preferred rank). Again, these cases are less than 0.01% in the datasets of this experiment (Dn, $n \in [1,3]$). Further analysis on this aspect deals with the influence of qualifiers on claims' logical statuses, as discussed Wikidata qualifiers (Section 3.1.2, see in particular Patel-Schneider [83] distinction between *additive* or *contextual* qualifiers), which is outside of the scope of this experiment.

Finally, this experiment does not address the representation of interrelated claims. While it is theoretically possible to extract concurrent sets of claims by querying their contextual information (e.g., retrieving all claims with the same selection of qualifiers or references), this process is not straightforward, as discussed in Section 4.2.5. Therefore, it has not been included in this experiment. Further analysis is required among this topic, which will be addressed in the next chapter (Chapter 5).

Chapter 5

Introducing EWA to represent the evolution of critical inquiry. Case study on challenging document authenticity

In this chapter, RQ2 - *How can critical inquiry be effectively and efficiently represented with minimal complexity?* - is further discussed, focusing on scholarly practices concerning the challenge of historical documents' authenticity. In particular, this chapter further refines RQ2 as:

- RQ2e - *Which are the typical patterns that scholars use to challenge a document's authenticity?*
- RQ2f - *How can EWA be introduced to represent such patterns?*

5.1 Domain analysis and source material

The *Index of Medieval Documents Concerning the Upper Austrian Region that are Damaged, Tampered with, or Altered* [52] is a catalogue compiled by Siegfried Haider and published in 2022. It contains over 150 known or suspected forgeries from Austria collected in the *Urkundenbuch des Landes ob der Enns*. Haiders' catalogue provides essential information about the authenticity of the inquiry of each document.

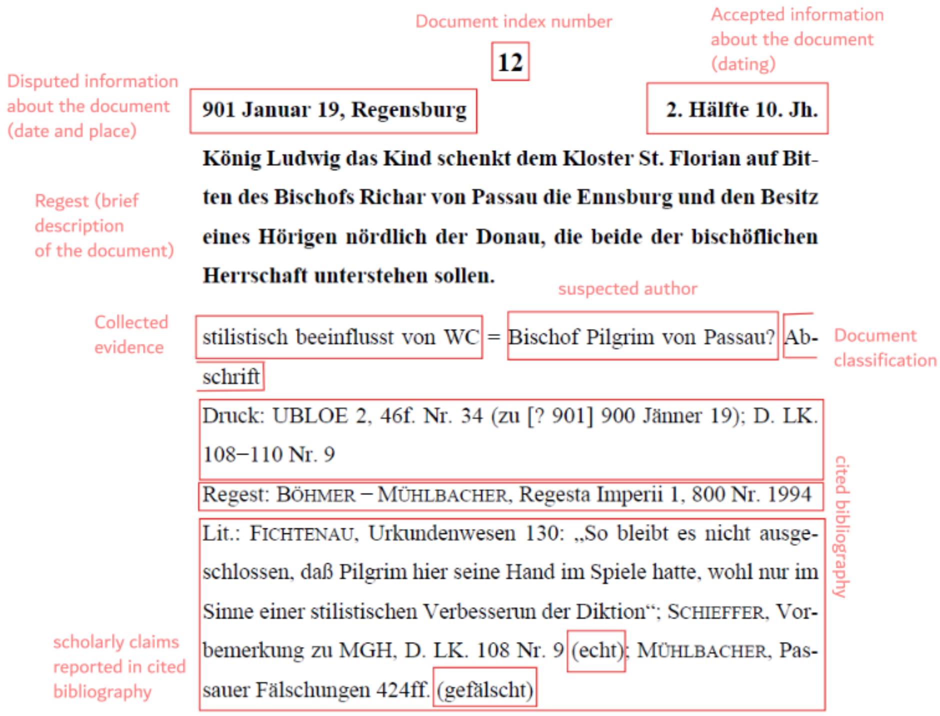

Figure 5.1: An example item of the catalogue

The catalogue is structured in the form of brief summaries. Each summary includes, on the left, the alleged date of the document and, on the right, the proposed (hypothesised) dating. Additional information follows, detailing the transmission type and manipulation, such as the presumed original, copies, insertions, interpolations, erasures, etc. A comprehensive bibliography is provided limited to publications that address criteria of authenticity or forgery, or where suspicions related to a document are discussed. These references may also include scholarly opinions on the document's authenticity, such as identifying it as a forgery. This creates a set of contrasting views regarding the document's authenticity: the information provided by the document itself (asserting its authenticity) versus the interpretations offered by scholars, particularly Haider (sometimes drawing on others' opinions).

Each catalogue entry records some relevant aspects concerning the authenticity inquiry of the described document, as illustrated in Figure 5.1 for what concerns the 12th document of the collection. Such aspects are presented as follows:

- Alleged information about the document (date and place of creation) stated in the document is now deemed manipulated by Haider. For instance, the date and place declared by charter 12 are *901 Januar 19, Regensburg* as shown in Figure 5.1.

- Proposed information about the document (notably, the dating of the document) reported by Haider. For instance, Haider dates charter 12 in the second half of the 10th century (*2. Hälfte 10. Jh.*) as shown in Figure 5.1.

- The regest of the document, which consists of a brief description of the document's contents.

- The document's classification, collected evidence to support its opinion, and the suspected author as interpreted by the scholar. For instance, Bishop Pilgrim of Passau is suspected of some stylistic influence on the document (*stilistisch beeinflusst von WC = Bischof Pilgrim of Passau?*) and the document is categorised as a copy *Abschrift* as shown in Figure 5.1.

- A comprehensive bibliography citing previous studies or mentions of the document is reported. It is categorised into regests (*Regest*), printed editions (*Druck*), and literature (*Lit.*). Figure 5.1 shows several example as *SCHIEFFER, Vorbemerkung zu MGH, D. LK. 108 Nr. 9* or *UBLOE 2, 46f. Nr. 34*.

- Cited regests, printed editions and literature may include scholarly opinions concerning the document's authenticity, (e.g. identifying it as a forgery). Figure 5.1 shows some examples as the dating attempt in *UBLOE 2* or the declared authenticity (marked as *echt*) by Schieffer and declaration of inauthenticity declared by Mühlbacher (marked as *gefälsht*).

Besides the description of each document, the catalogue includes:

- The bibliography reports all cited work. For example, the citations mentioned above *SCHIEFFER, Vorbemerkung zu MGH, D. LK. 108 Nr. 9* is included in the Bibliography section as *MGH, Die Urkunden der deutschen Karolinger, Bd. 4: Die Urkunden Zwentibolds und Ludwigs des Kindes, bearb. von Theodor SCHIEFFER, Berlin 1960.*

- A Table of abbreviations used in the text. For example, *NÖUB* stands for *Niederösterreichisches Urkundenbuch.*

- A concordance table mapping the document's index number to its collection, volume and number. For instance, the 12th document of the collection is *UBLOE 1 2/Anhang 1*, meaning the document is the first insert (*Anhang*) of the first volume of the collection *Urkundenbuch des Landes ob der Enns*

Overall, the catalogue contains 153 known manipulated documents thought to have been issued from 498 to 1363, 183 works about the documents are reported in the bibliography, 42 (out of 153 documents) are reported with their issuing locations around Europe (mainly Austria and Germany), and 58 (out of 153 documents) a motivation or evidence collected or contextual information about their authenticity inquiry.
Following the source data analysis described in mythLOD methodology [81], a set of representation requirements over the data and data manipulation instructions from free-text to RDF has been defined to bind the domain of this work. Table 5.1 reports each record category in the catalogue in relation to the representation requirements (cf. first column) to build a data model that effectively represents the key aspects relevant to this study.
In particular, each representation requirement addresses the role of specific information presented in the catalogue concerning the modelling of scholarly opinions. This includes aspects such as the opinion content, context, associated logical status in the dispute, and relationships between those opinions. In this study, Haider's interpretations are considered authoritative and des-

ignated as the settled point of view on each dispute, implying that alleged information is considered superseded. For example, the accepted information about the document (e.g., dating) deals with two main representation requirements, specifically that information is part of a scholarly opinion and that such opinion is part of the debate but is considered settled (and therefore asserted). Representation requirements ensure consistency between data sources and the modelling activity documented in the next section (Section 5.2).

Additionally, Table 5.1 outlines the expected data manipulation instructions to be performed on the dataset, transitioning from free-text to RDF as modelled by the data model. For example, regarding the accepted information about the document (e.g., dating) reported in the catalogue, several actions will be undertaken, including the extraction of years, conversion of years into computable objects, handling of fuzzy dates (timespans), and the assignment of a unique identifier to the reported dates.

Record category	**Representation requirement**	**Data managment actions**
Document index number	//	(1) Unique identifier assignment to each indexed document
Alleged information about the document (e.g., date and place)	(1) Date and place form a superseded information (2) Superseded information are WLS opinions, and they require EWA (and therefore non-asserted)	(1) Values separation (2) Dates conversion in computable objects (3) Fuzzy dates handling (4) Location names normalization and geo-referencing (5) Unique identifier assignment to the reported date and place

Accepted information about the document (e.g., dating)	(1) Haider's dating attempt is part of a scholarly opinion (2) Such opinion is part of the debate but is considered settled (and therefore asserted)	(1) Dates conversion in computable objects (2) Fuzzy dates handling (3) Unique identifier assignment to the reported date
Regest (a brief description of the document's contents)	(1) The regest is not part of the scholarly debate and therefore should be represented as asserted	(1) Regest text extraction (2) Text translation from German to English
Document's classification	(1) Haider's diplomatic classification is part of a scholarly opinion (2) Such opinion is part of the debate but is considered accepted (and therefore asserted)	(1) Information extraction from free text (2) Normalization of document category (copy, indicated as *Abschrift*) (3) Unique identifier assignment to the document category
Suspected author	(1) The forger of the document is the actual creator of the document (2) Haider's attempt to identify a forger is part of a scholarly opinion (3) Such opinion is part of the debate but is considered settled (and therefore asserted)	(1) Value separation from other record categories reported in the same catalogue field (2) Extraction and normalization of the suspected creator (3) Unique identifier assignment to the suspected creator

Collected evidence	(1) Some evidence is collected by the scholar in his study (2) Collected evidence supports the scholar's opinion (3) Evidence collection addresses some document feature (e.g., style) on which the scholars perform a certain evaluation (e.g., inconsistency) (4) Document features and evaluations may require a controlled vocabulary (5) Evidence collection is not a WLS information and therefore do not require EWA	(1) Value separation from other record categories reported in the same catalogue field (2) Extraction and normalization of document features addressed by the evidence collection (3) Extraction and normalization of the evaluation performed in an evidence collection (4) Unique identifier assignment to evidence, features, and evaluations
Bibliography (regests, printed editions, and literature) citing previous studies or mentions of the document	(1) The reported bibliography supports the scholar's opinion (2) Cited works and scholars contextualize the scholar opinion by representing its study background (3) Bibliography reporting is not a WLS information and therefore does not require EWA	(1) Separation of citations and unique identifier assignment (2) Reconciliation against complete bibliographic record (3) Unique identifier assignment to each citation

Cited bibliography reporting scholarly opinions concerning the document's authenticity	(1) The opinion is part of the scholarly debate (2) The source of the opinion is related bibliographic entry (3) Its WLS depends on its agreement or disagreement with the settled opinion	(1) Extraction of the claimed content (2) Normalisation of the claimed content (3) Unique identifier assignment to identify the claim (4) Decision of its logical status depending on the opinion contents
Bibliography	(1) Descriptive metadata (author, place, and date of publication) of bibliographic entries contextualize the scholarly claims which cite them	(1) Extraction and indexing of bibliographic entries (2) Mapping with the table of abbreviations (3) Mapping with citations in the catalogue (4) Descriptive metadata extraction
Table of abbreviations	//	(1) Mapping with catalogue contents (e.g., Bibliography and documents' titles)

Table 5.1: Summary of the representation requirements and the data management actions extracted from the source material analysis

5.2 Knowledge representation

In alignment with LOD standards and as suggested by mythLOD methodology, the reuse of existing ontologies has been prioritised and extended only when necessary. Additional classes and properties were introduced solely in cases where existing ontologies were insufficiently representative or to accommodate specific contents of the catalogue.

5.2.1 Scholarly opinions representations and categorisation depending on their logical status

Conjectures has been used as the reification method to represent the claims' logical status and reach EWA when needed. Undisputed information is structured as a distinctly named graph named `factual_data` as suggested in the Digital Hermeneutics data model [24], so their contents are asserted. Following the conjectural graphs definition (Section 4.2.1), disputed claims are represented through conjectural graphs; therefore, their contents are not asserted. Conjectural graphs include both alleged information provided by charters themselves as well as scholarly opinions reported in the cited bibliography of each document in the catalogue. Finally, accepted scholarly claims are represented as settled conjectures, re-asserting their contents through a settled as defined in Section 4.2.3.

Each conjecture (both settled and not) is associated with the set of available contextual information (e.g., the bibliography mentioned, the evidence collected).

```
@prefix dct:    <http://purl.org/dc/terms/> .
@prefix geo:    <http://www.opengis.net/ont/geosparql#> .
@prefix owl:    <http://www.w3.org/2002/07/owl#> .
@prefix riopac: <http://opac.regesta-imperii.de/id/> .
@prefix xsd:    <http://www.w3.org/2001/XMLSchema#> .
@prefix rdfs:   <http://www.w3.org/2000/01/rdf-schema#> .
@prefix prov:   <http://www.w3.org/ns/prov#> .
@prefix hico:   <https://w3id.org/hico#> .
@prefix sebi:   <https://w3id.org/sebi/> .
@prefix time:   <https://www.w3.org/TR/owl-time/> .
@prefix wd:     <http://www.wikidata.org/entity/> .
@prefix cito:   <http://purl.org/spar/cito/> .
@prefix fabio:  <http://purl.org/spar/fabio/> .
@prefix ov:     <http://open.vocab.org/terms/> .
@prefix rico:   <https://www.ica.org/standards/RiC/ontology#> .

@prefix conj:   <https://w3id.org/conjectures/> .
@prefix :     <https://w3id.org/broast/urk/> .
```

```
@prefix doc:    <https://w3id.org/broast/urk/documents/> .
@prefix people: <https://w3id.org/broast/urk/scholars/> .
@prefix pub:    <https://w3id.org/broast/urk/publications/> .
```

Listing 5.1: List of prefixes

Each claimed content attempts to classify a document's authenticity by expressing information about the document itself, which usually may be debated: authenticity classification (e.g. the document is authentic, is suspicious or is a forgery), date and place of creation, suspected author. All these elements should be expressed as a conjectural graph (despite being re-asserted by a settled conjecture or not), but none is required. In simpler terms, a scholarly opinion may propose a date (e.g., 950-1000) and a juridical categorisation (e.g., forgery) of a document but neglecting its authorship (e.g., a suspected author is unknown or simply not considered).

To my knowledge, no ontology formalises specifically critical inquiry on forged documents. For this reason, an ontology called SEBI (Scholarly Evidence Based Interpretation ontology)[1] is proposed. SEBI is a simple pattern representing the evidence scholars collect to support their interpretations. In particular, the ontology is provided with a set of named individuals and classes which characterise the use of the ontology to represent the evidence which supports scholarly interpretations about the classification of a document's authenticity (in particular by defining such document an Authentic document, a Forgery or a Suspected Forgery). The data model, therefore, integrates SEBI ontology and other existing ontologies such as Dublin Core Terms[2] to describe the document metadata and bibliographical entries descriptions, Time ontology[3] to handle fuzzy datings, as shown in Figure 5.4.

Documents categorisations are formalised as instances of one of the classes `sebi:Forgery`, `sebi:Authentic`, and `sebi:SuspectedForgery`, and they are all disjoint mimicking the need to choose a single point of view in each claimed

[1] The documentation of the ontology is available at `https://valentinapasqual.github.io/sebi/`

[2] `https://www.dublincore.org/specifications/dublin-core/dcmi-terms/`

[3] `https://www.w3.org/TR/owl-time/`

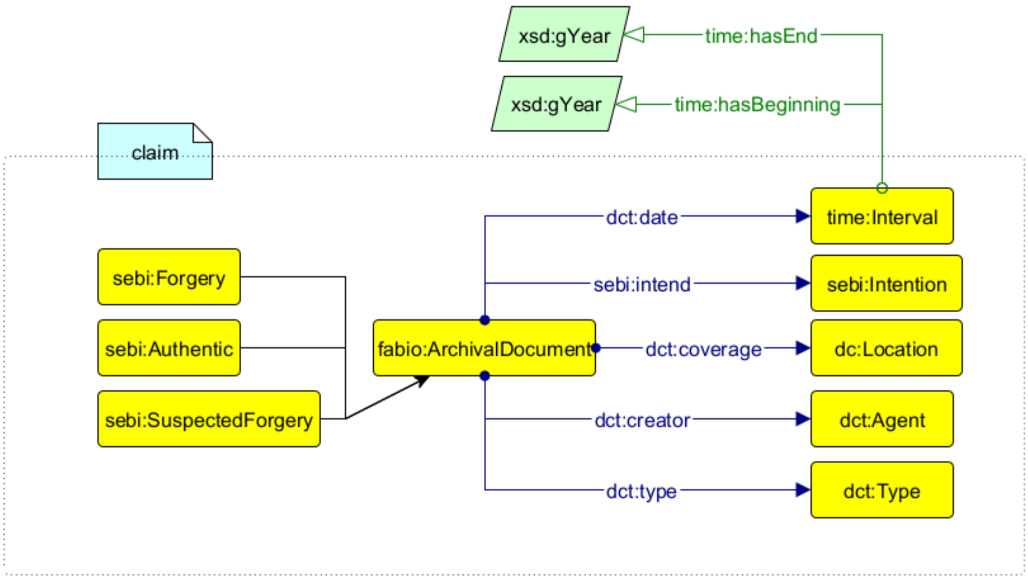

Figure 5.2: Selection of classes and properties to represent scholarly opinions tackling authenticity inquiry of a document

content (if no conclusion can be reached, other metadata as date, place and author should be registered). Each document is an instance of the class `fabio:ArchivalDocument`. The creator of the document (expressed through `dct:creator dct:Agent`), the date of creation (`dct:date time:Interval`), location of creation (`dct:coverage dct:Location`). The `dct:date` property is connected to a `time:Interval` class, which includes `time:hasBeginning` and `time:hasEnd` properties to specify the creation period and handle fuzzy time-spans. Even if not recorded in the catalogue, the intention behind the document creation is formalised via introducing a new predicate and class (`sebi:intended sebi:Intention`)[4]. Additionally, dates can be annotated with `rico:dateQualifier` and other entities with `rico:confidence` to record uncertainty related markers (e.g., circa, possibly).

Regardless of their logical status or origin, all opinions about document forgery detection can be represented using at least one of the classes and properties described and illustrated in Figure 5.2.

[4]In the catalogue, hypothesised intentions are not recorded, therefore this information is kind of information is not present in this work KG

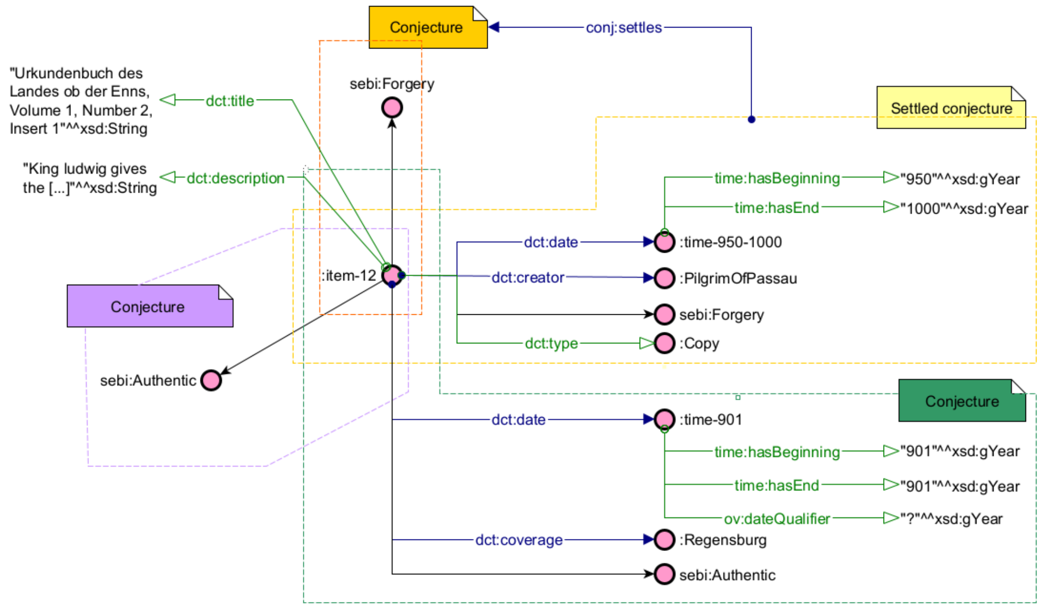

Figure 5.3: Graphical representation of four opinions and their categorisation concerning their logical status about the authenticity of the 12th document of the collection

For instance, Figure 5.3 illustrates four of the available opinions involved in the authenticity inquiry of the 12th document of the collection (`doc:12`) modelled with the data model just introduced.

All opinions, indicated with a surrounding coloured dashed line, share the same subject, the document under analysis (`doc:12`). Haider's opinion, coloured in yellow, is marked as a settled conjecture and comprehends a dating of the document (950-1000), suspected authorship (Pilgrim of Passau). The superseded opinion, indicated in green and categorised as a conjecture, describes the document with date 901 (`:time-901`) and place Regensburg `:place-Regensburg`. Additionally, Mühlbacher's and Schieffer's opinions, depicted in purple and orange, define the document as authentic (`sebi:Authentic`) and forged (`sebi:Forgery`). All entities not included in opinions, and therefore not surrounded by coloured dashed lines in Figure 5.3, are stored in the `:factual_data` named graph. That information includes descriptive metadata of the document (e.g. title and description, indicated as literals by the data properties `dct:title` and `dct:description`) or specifications of opinion contents (e.g. the declaration of beginning `time:hasBeginning` and end `time:hasEnd` of the datings indicated by the yellow and green opinions).

5.2.2 Representation of opinions contextual information

Contextual information about the opinion concern those aspects which have been highlighted at the beginning of this section, such as the evidence collected by the scholar to reach a certain conclusion (using HiCo[5] and SEBI), as well as the author of the opinion and relevant bibliographic entries (using PROV-o)[6] and Dublin Core.

As shown in Figure 5.4, each graph storing a scholarly opinion is categorised as a `rdf:type hico:InterpretationAct`. The bibliographic source (instance of `fabio:WorkCollection` or `fabio:Expression`) from which the opinion is extracted is represented via the property `prov:wasQuotedFrom`. Similalry, cited bibliographic resources are modelled via the property `prov:wasDerivedFrom` and the opinion responsible entity is recorded via the property and class `prov:wasAttributedTo prov:Agent`. Each bibliographical resource is represented with a set of object and data properties from Dubin Core vocabulary representing main features of the work, such as the title `dct:title`, a brief description `dct:description`, publishing date (`dct:date`), the language `dct:language`, main subjects `dct:subject` and involved agents (`dct:creator`, `dct:publisher`, `dct:contributor`). When dealing with superseded opinions, the source of the opinion is ideally stated in the document itself, for this reason, the source of such opinions is represented as `prov:wasQuotedFrom fabio:ArchivalDocument`.

The evidence supporting the opinion is recorded as (`sebi:support sebi:Evidence`). Document features and their evaluation are considered as critical components of the ontology. Document features (`sebi:Feature`) are either extrinsic features (`sebi:ExtrinsicFeature`), intrinsic ones

[5] https://marilenadaquino.github.io/hico/
[6] https://www.w3.org/TR/prov-o/

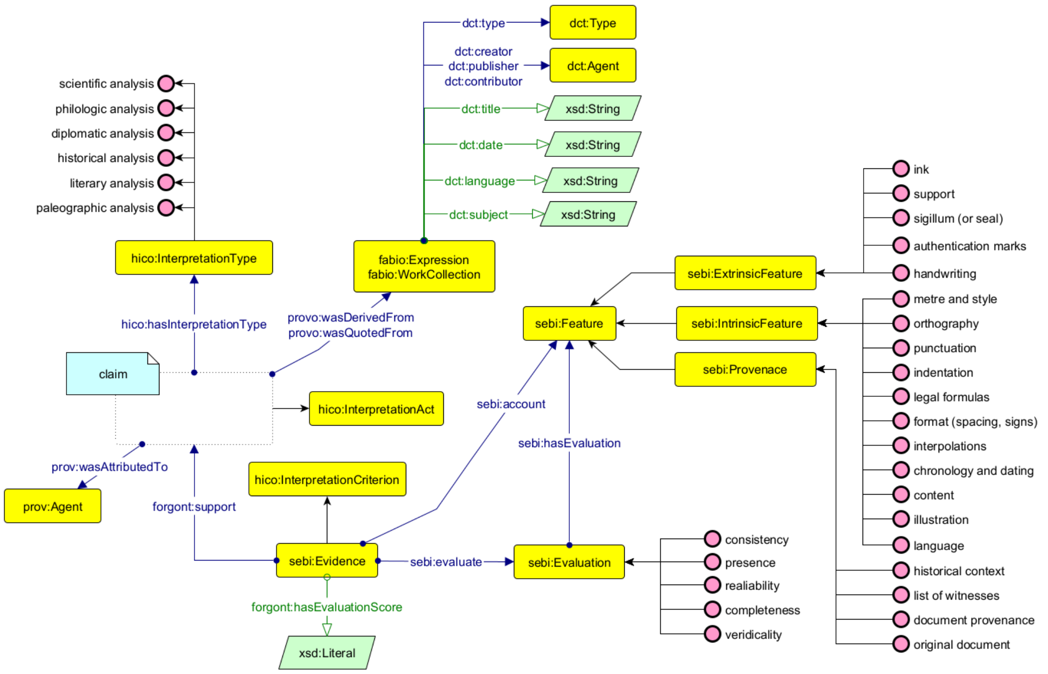

Figure 5.4: Ontological classes and properties for representing scholarly arguments and contextual information about document authenticity

(`sebi:IntrinsicFeature`), or provenance information (`sebi:Provenance`), capturing aspects such as ink, support, handwriting, and orthography. Each feature is evaluated on a set of established criteria (`sebi:Evidence`) such as consistency, presence, completeness, veridicality, and reliability. A score is associated to each evidence as `xsd:Literal` using the property `sebi:hasEvaluationScore`. The evaluation score indicates a measure on each collected evidence, allowing the integration of negatives (e.g. the absence of the signature in a document is represented as an evidence based on the feature "authentication marks" with evaluation "presence", with score `false` or 0).

Therefore, the contextual metadata concerning the four opinions represented in Figure 5.3 are represented as shown in Figure 5.5. In particular, Haider's opinion (in yellow) is enriched with its source (`prov:wasQuotedFrom :ref-Urkundenverzeichnis`), its author (`prov:wasAttributedTo :person-SiegfriedHaider`), the bibliographic resources mentioned

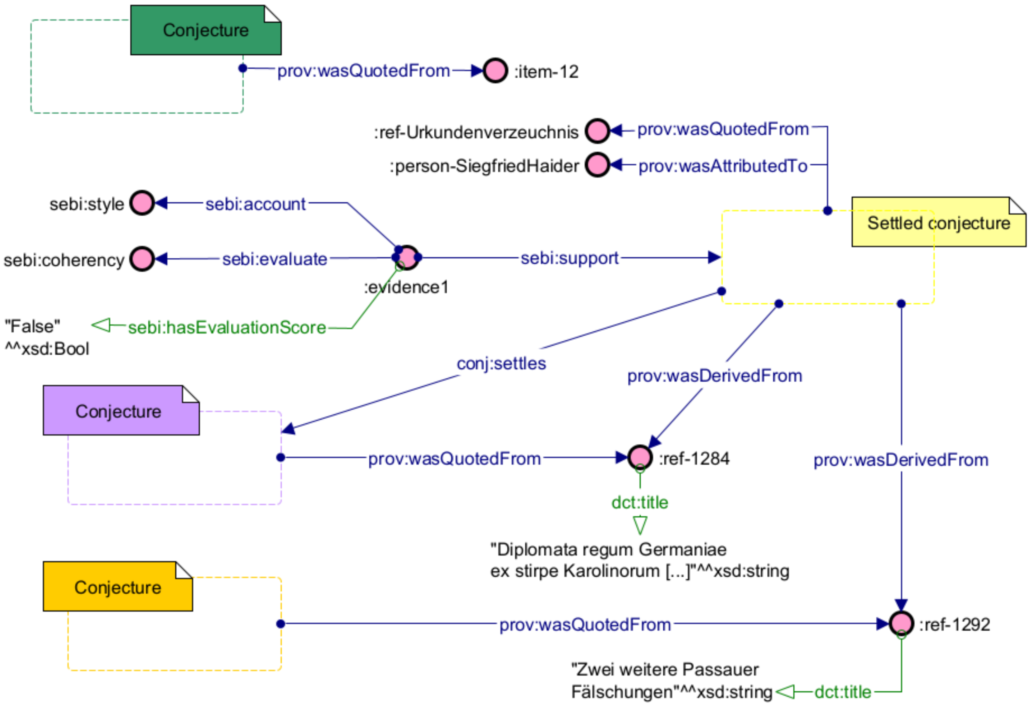

Figure 5.5: Selection of classes and properties to represent the contextual information about four opinions tackling authenticity inquiry of document 12 in the collection

(`prov:wasDerivedFrom pub:1284, pub:1292`) and the evidence collected to support its opinion (`:evidence1 sebi:support`). Haider's collected evidence is based on the account of the stylistic features of the document (`sebi:account sebi:style`, which is an instance of the class `sebi:IntrinsicFeatures`), produces an evaluation (`sebi:evaluate sebi:coherency`, which is an instance of the class `sebi:Evaluation`). The evaluation score of this evidence is "False" (`sebi:hasEvaluationScore "False"8sd:boolean`), therefore stating the evidence as "incoherent style".

Haider's opinion correlates with Mühlbacher's and Schieffer's (respectively, marked in purple and orange) by the bibliographic resources `pub:1284, pub:1292`. These resources are the source from which the two opinions are derived (`prov:wasQuotedFrom`), while the source of the superseded opinion, marked in green, is reported to be the document itself (`prov:wasQuotedFrom doc:12`). Each bibliographic resource is accompanied by a set of descriptive metadata previously described. These are exemplified in Figure 5.5 by the titles of the resources (`dct:title`), while other metadata is omitted for clarity. Notably, Haider's inquiries may believe or disbelieve such scholarly opinions (e.g., bibliographic sources claiming the document as a forgery). As described in the definition of a settled conjecture (provided in Section 4.2.3), a settled conjecture can explicitly settle other conjectures (in particular, see the agreeing opinions towards the fraudulence of the "Donation of Constantine" by Lorenzo Valla and Nicholas of Cusa represented in Listing 4.13). The agreement has been established between Haider's opinions (settled conjecture) and other scholars' opinions (conjectural graphs), which share the same triples. For the sake of the example, the relation of agreement between Haider's (in yellow) and Mühlbacher's (in orange) is represented by the object property `conj:settles` to mark their agreement explicitly.

5.3 Data extraction and Knowledge Graph population

This stage of the data management process outlined by the mythLOD methodology [81] focuses on data cleaning, entity linking, and dataset production activities. This section details converting the catalogue's unstructured contents from a DOCX file into a structured KG and describes the methods and processes used to convert raw, unstructured data into structured data.

The DOCX file was parsed as a ZIP archive containing XML elements, where targeted paragraphs (`<p>`) were identified within a specified range corresponding exclusively to the catalogue section, excluding introductory content, the table of contents, and bibliographic references. Catalogue items were mapped using their numerical identifiers, such as item 12, located in lines 140–146 of the DOCX structure.

Each catalogue entry was classified as either *complete* or *incomplete* based on the presence of all paragraphs containing all key components highlighted in

Figure 5.1 (e.g., alleged information, proposed information on the document, cited references). Manual input was applied to address missing information for incomplete entries (e.g., the Literature paragraph is missing), ensuring their inclusion in the following process stages.

After classification, the entries were reformatted into a standardised structure (JSON format). Key elements, corresponding to the record categories detailed in the source analysis source Table 5.1, were extracted and saved in an intermediate JSON file for further processing. Listing 5.2 illustrates the initial structuring of relevant information for catalogue item 12.

```
{
    "index": "12",
    "disputed_information": [
        "901 Januar 19, Regensburg"
    ],
    "settled_information": [
        "2. Hälfte 10. Jh."
    ],
    "regest_abstract": "König Ludwig das Kind schenkt dem Kloster St. Florian auf
    Bitten des Bischofs Richar von Passau die Ennsburg und den Besitz eines Hörigen
    nördlich der Donau, die beide der bischöflichen Herrschaft unterstehen sollen.",
    "comment": "stilistisch beeinflusst von WC = Bischof Pilgrim von Passau?
    Abschrift",
    "critical_edition": "Druck: UBLOE 2, 46f. Nr. 34 (zu [? 901] 900 Janner 19); D.
    LK. 108--110 Nr. 9",
    "regest_list": "Regest: Bohmer -- Mühlbacher, Regesta Imperii 1, 800 Nr. 1994",
    "literature": "Lit.: Fichtenau, Urkundenwesen 130: "So bleibt es nicht
    ausgeschlossen, daß Pilgrim hier seine Hand im Spiele hatte, wohl nur im Sinne
    einer stilistischen Verbesserung der Diktion"; Schieffer, Vorbemerkung zu MGH, D
    . LK. 108 Nr. 9 (echt); Mühlbacher, Passauer Fälschungen 424ff. (gefälscht)"
}
```

Listing 5.2: JSON intermediate file storing information related to document 12

The Bibliography section was processed using the same approach, resulting in its extraction and indexing in a separate intermediate JSON file. Listing 5.3 illustrates the JSON entry for Mühlbacher's work, *Zwei weitere Passauer Fälschungen*, indexed as bibliographic resource number 1292.

```
    "1292": {
        "ref_text": "Engelbert Mühlbacher, Zwei weitere Passauer Fälschungen. In: MI
    ÖG 24 (1903), 424-432."
    },
```

Listing 5.3: JSON intermediate file storing information related to the bibliographic resource

Some bibliographic resources were grouped in the catalogue (e.g. reference 1331: *Urkundenbuch des Landes ob der Enns, Bd. 1, Wien 1852; Bd. 2, Wien 1856; Bd. 3, Wien 1862; Bd. 4, Wien 1867; Bd. 6, Wien 1872; Bd. 7, Wien 1876; Bd. 8, Wien 1883*). These were semi-authomatically parsed, extracted and indexed separately due to their different publication dates (e.g. index number 1331-1 for *Urkundenbuch des Landes ob der Enns, Bd. 1, Wien 1852* for volume 1, index number 1331-2 for *Bd. 2, Wien 1856* and so forth for each volume).

Following the data conversion guidelines outlined in Table 5.1, additional data cleaning operations were conducted across different record categories. Specifically:

- For the **superseded information about the document**, both dates and locations were extracted, standardised, and indexed. These cleaned data were then stored in an intermediate JSON file. Locations were further reconciled against Wikidata. For instance, the date "901 Januar 19" has standardised as `"year": "901"` and the location "Regensburg" was encoded as `"location": ["Q2978", "Regensburg"]` to correspond to its Wikidata entity.

- For the **settled information about the document**, date expressions were extracted, standardised, and indexed into an intermediate JSON file. As these datings followed less regular patterns, this process was conducted semi-automatically. For instance, the dating "2. Hälfte 10. Jh.", meaning second half of 10th century, was standardised as `"start_year": "950", "end_year": "1000"`.

- The **regest** section, providing a brief summary of the document's content, were extracted and then translated in English.

- The **comment** section included several relevant information such as the document's classification, the collected evidence and the suspected author.

  - **Document's classification** terms as *Abschrift* (Copy), *Or.* (Original), *angebliches Or.* (alleged original), *verdächtiges Or.* (suspicious original) *insert* (insert) were extracted and standardised into English. Extracted data have been added to the JSON file storing the catalogue data.
  - **Collected evidence** related terms were manually reviewed and recorded in a CSV file. This process facilitated the refinement of SEBI-related taxonomies, focusing on characteristics of evidence collection, such as documentary features and evaluation-related terms. Terms too vague or not clearly stated have not been taken into consideration, e.g. "innovation".
  - **Suspected author** terms were selected manually from the text field. Since only Bishop Pilgrim of Passau is marked as the suspected author, this entity has been reconciled against Wikidata. Extracted data have been added to the JSON file storing the catalogue data.

- For each category of the **cited bibliography** (printed editions, regests, literature), records have been extracted, indexed and mapped to their corresponding entries in the bibliography section. The disambiguation process was conducted semi-automatically due to inconsistencies in citation formats. For instance, from the edition section, the string "MÜHLBACHER, Passauer Fälschungen 424ff. (gefälscht)" was isolated from other citations. The author's surname (MÜHLBACHER) and the title (Passauer Fälschungen) were cross-checked against the bibliography to match the complete reference, indexed as 1292 in the corresponding JSON file. When volume numbers were provided, they were further verified against the bibliography. Additionally, in instances like "UBLOE 2, 704-708 Anhang Nr. 4," the string was split differently from previous case, matching the title "UBLOE" and the volume number "2" and reconciling the entity to its complete reference, indexed as 1331-2 in the corresponding JSON file. Unmatched items were flagged as "unknown" and later manually reconciled.

- **Citations reporting scholarly opinions** were identified using textual markers such as parentheses, double quotes, and crosses. These elements were extracted and recorded in a CSV file, which was manually reviewed and annotated. This approach was necessary due to the lack of a typical pattern across comments and proved helpful in further refining the SEBI taxonomy. For example, the citation *MÜHLBACHER, Passauer Fälschungen 424ff. (gefälscht)* was annotated in the CSV as forgery, while the citation *SCHIEFFER, Vorbemerkung zu MGH, D. LK. 108 Nr. 9 (echt)* was annotated as authentic. The same strategy was adopted to store opinions including concurrent datings and authorship.

- Uncertainty markers have been extracted, standardised and translated in English. The following indicators have been preserved: parentheses to represent interpreted contents, fuzzy markers about dates (e.g., before, after, circa), uncertain markers (e.g., ?, possibly, allegedly).

- All bibliographic resources listed in the Bibliography section have been extracted and indexed into a JSON file, as previously described. These entries were cross-referenced with the OPAC Regesta Imperii database[7], and relevant metadata were extracted and incorporated into the JSON file. Contributors, publishers, and creators of the bibliographic resources have been standardised and aligned with Wikidata. For instance, *En-*

[7] https://opac.regesta-imperii.de/

gelbert Mühlbacher has been standardised as *Mühlbacher, Engelbert* and reconciled with the Wikidata ID (`Q87151`[8]).

The process involved multiple iterations to reach the final output, with manual checks and refinements applied to all generated JSON files. Upon completion, the data were converted into the RDF using RDFLib[9]. The resulting KG follows the data model presented in the previous section (Section 5.2). Opinions included in the debate were stored in distinct named graphs, while uncontested information was incorporated into the factual data graph.
Currently, no dedicated parser for Conjectures exists. Therefore, named graphs were converted into conjectures and settled conjectures using regular expressions.

Listing 5.4 shows a snippet of the resulting KG according to the data model concerning the representation of scholarly opinions on item 12. In particular, the document 12 claiming its own authenticity (`CONJ :ass-12`) is shown in lines 1-6, Haider's settled opinion is shown (`SETT :ass-12-2`) with its contextual information and evidence collected (lines 9-17), Schieffer's opinion (`CONJ :ass-lit-12-ref2`) with its contextual information (lines 40-43) and Mühlbacher's opinion (`CONJ :ass-lit-12-ref3`, lines 45-49). Agreement between Haider's settled opinion and Mühlbacher is established via the predicate `conj:settles` (line 16).
Undisputed data are stored in `GRAPH :factual_data`, which includes timespans definition (`:time-950-1000` and `:time-901` in lines 147-157), superseded location definition (`:Regensburg`, lines 139-163) and cited bibliographic resources definition (`pub:1284` and `pub:1292`, lines 165-189).

```
1 # Document 12 claiming its own authenticity
2 CONJ :ass-12 {
3     doc:12 dct:date :time-901 ;
4         dct:coverage :Regensburg .
```

[8] Engelbert Mühlbacher record is available at: `http://www.wikidata.org/entity/Q87151`

[9] `https://rdflib.readthedocs.io/en/stable/`

```
}
kb:ass-12 prov:wasQuotedFrom doc:12 .

# Acceped opinion by Haider
SETT :ass-12-2 {
    doc:12 dct:date :time-950-1000 ;
        a sebi:Forgery .
        dct:creator people:PilgrimOfPassau;
        dct:type :Copy .

    people:PilgrimOfPassau ov:confidence "?"^^xsd:string .
    :ass-12-2 conj:settles :ass-lit-12-ref3 .
}

:ass-12-2 :ass-12-2 a hico:InterpretationAct ;
    prov:wasAttributedTo people:SiegfriedHaider ;
    prov:wasDerivedFrom :ass-lit-12-ref1,
        kb:ass-lit-12-ref2,
        kb:ass-lit-12-ref3,
        kb:ref-1206,
        kb:ref-1229,
        pub:1284,
        pub:1292,
        kb:ref-1331-2 ;
    prov:wasQuotedFrom pub:Urkundenverzeichnis .

# Evidence collected by Haider
:ev1-ass-12-2 a sebi:Evidence ;
    rdfs:label "Stylistically influenced"^^xsd:string ;
    sebi:support kb:ass-12-2 ;
    sebi:account sebi:style ;
    sebi:evaluate sebi:coherency ;
    sebi:hasEvaluationScore "False"^^xsd:boolean .

# Reported opinion by Schieffer
CONJ :ass-lit-12-ref2 {
    doc:12 a sebi:Authentic .
}
:ass-lit-12-ref2 prov:wasQuotedFrom pub:1284 .

```

```
# Reported opinion by Mülbacher
CONJ :ass-lit-12-ref3 {
    doc:12 a sebi:Forgery .
}
:ass-lit-12-ref3 prov:wasQuotedFrom pub:1292 .

GRAPH :factual_data {

    # Accepted timespan definition
    :time-950-1000 a time:Interval ;
        rdfs:label "950-1000"^^xsd:string ;
        time:hasBeginning "0950"^^xsd:gYear ;
        time:hasEnd "1000"^^xsd:gYear .

    # Currenlty disputed timespan definition
    :time-901 a time:Interval ;
        rdfs:label "901"^^xsd:string ;
        time:hasBeginning "0901"^^xsd:gYear ;
        time:hasEnd "0901"^^xsd:gYear .

    # Location definition
    :Regensburg a dct:Location ;
        rdfs:label "Regensburg"^^xsd:string ;
        forgont:hasLatitude "POINT (12.083333333 49.016666666)"^^geo:wktLiteral ;
        owl:sameAs wd:Q2978 .

    # Schieffer's work
    pub:1284 a fabio:ArchivalRecordSet ;
        dct:date "1960"^^xsd:string ;
        dct:language "ger"^^xsd:string ;
        dct:publisher kb:person-SchiefferTheodor ;
        dct:source "https://www.dmgh.de/mgh_dd_zwent__dd_ldk"^^xsd:string ;
        dct:subject "Deutschland; 9. Jahrhundert; 10. Jahrhundert; Urkunden;
    Karolinger; Zwentibold <Lothringen, König> (870-900); Schieffer, Theodor
    (1910-1992)"^^xsd:string ;
        dct:title "Diplomata regum Germaniae ex stirpe Karolinorum (Die Urkunden der
     deutschen Karolinger). Bd. 4: Die Urkunden Zwentibolds und Ludwigs des Kindes (
    Zwentiboldi et Ludowici Infantis Diplomata)"^^xsd:string ;
        dct:description "MGH, Die Urkunden der deutschen Karolinger, Bd. 4: Die
    Urkunden Zwentibolds und Ludwigs des Kindes, bearb. von Theodor Schieffer,
```

```
    Berlin 1960." ;
        owl:sameAs riopac:55065 .

    # Muhlbacher's work
    pub:1292 a fabio:Article ;
        dct:creator kb:person-MuhlbacherEngelbert ;
        dct:date "1903"^^xsd:string ;
        dct:language "ger"^^xsd:string ;
        dct:relation "Mitteilungen des Instituts für Österreichische
    Geschichtsforschung 24 () S. 424-432"^^xsd:string ;
        dct:source "http://archive.org/stream/mitteilungendesi24univuoft#page/424/
    mode/2up"^^xsd:string ;
        dct:subject "Passau; Mittelalter; Urkunden; Diplomatik; Mühlbacher,
    Engelbert (1843-1903)"^^xsd:string ;
        dct:title "Zwei weitere Passauer Fälschungen"^^xsd:string ;
        dct:description "Engelbert Mühlbacher, Zwei weitere Passauer Fälschungen. In
    : MIÖG 24 (1903), 424-432." ;
        owl:sameAs riopac:114100 .
}
```

Listing 5.4: RDF representation of competing information about document 12 in the catalogue

5.4 Testing and Knowledge Exploration

As defined by the mythLOD methodology, a visual interface named Broast (*Broad Representation Of Authenticity aSsessmenT*) was developed to validate the outcomes of the data management activities. In this section, the development and functionalities of a web application are presented[10]. In particular, the application was designed to test and explore modelling choices (defined in Section 5.2) and the consistency and correctness of the produced RDF dataset (described in Section 5.3). The web-application is not intended for general public users. Still, it is a prototypical device that serves as both a demonstration tool and a practical platform for testing the produced knowledge. It acts as a proof-of-concept about including non-asserted content in online catalogues, demonstrating the potential for integrating the modelled

[10]The web-app is available at `https://projects.dharc.unibo.it/broast/`

information into online catalogues, which are central to this Dissertation.

The web-app is organised into four main components, as depicted in Figure 5.6 and is structured as follows.

- **Back-end Configuration:** The back-end is powered by Python Flask 3.0[11]. It is initiated and configured by the `run.py` file. Routes are managed through the `__init__.py` script within the `routes` folder. Each route is handled by individual Python scripts (`main.py`, `document.py`, `documents.py`, `scholar.py`, `scholars.py`, `publication.py`, and `publications.py`), which address specific functionalities of the website which will be described later in this section[12].
- **Front-end and Templates:** The front-end is developed using HTML 5.1[13], Bootstrap 5.1[14], and Jinja 3.1 templating[15]. HTML templates located in the `templates` folder (e.g., `index.html`, `document.html`, `documents.html`, `scholar.html`, `scholars.html`, `publication.html`, `publications.html`) interact with the back-end Python scripts to provide dynamic content updates[16].
- **Static Files:** The `assets` folder contains static files such as CSS, JavaScript[17], and images, which handle front-end styling and functionality.
- **GraphDB Integration:** The `graphdb` folder holds the triplestore managed by a stand-alone version of GraphDB[18]. This component integrates the Conjectures parser for handling RDF data.

[11] `https://flask.palletsprojects.com/en/3.0.x/`

[12] Other routes are present in the back-end to handle the templating system but they are omitted for clarity

[13] `https://dev.w3.org/html5/spec-LC/`

[14] `https://getbootstrap.com/docs/5.0/getting-started/introduction/`

[15] `https://jinja.palletsprojects.com/en/3.1.x/`

[16] Other templates are present in the app to handle the templating system but they are omitted for clarity

[17] `https://www.w3.org/wiki/The_web_standards_model_-_HTML_CSS_and_JavaScript`

[18] `https://graphdb.ontotext.com/documentation/9.8/enterprise/run-stand-alone-server.html`

```
broast
├── app
│   ├── __init__.py
│   ├── assets
│   ├── graphdb
│   ├── routes
│   │   ├── main.py
│   │   ├── document.py
│   │   ├── documents.py
│   │   ├── scholar.py
│   │   ├── scholars.py
│   │   └── reference.py
│   └── templates
│       ├── index.html
│       ├── document.html
│       ├── documents.html
│       ├── scholar.html
│       ├── scholars.html
│       └── reference.html
└── run.py
```

Figure 5.6: Diagram of the main components of the web-app structure: Python web application with organised components including initialisation, assets, database interaction, routing, and templating

This section primarily shows the contents provided in the web-app, mainly focusing on the interactions between the triplestore and the back-end through SPARQL queries, which serve as the principal method for evaluating this work as well as this theoretical framework.
After a brief introduction indicating the platform themes and main objectives, the homepage showcases two access points to the data collection corresponding to the data model's structure:

- **Exploration of the contents of the dispute: Document's catalogue**, which enables the exploration of documents and their related opinions with their categorisation based on their logical status (as detailed in Section 5.2.1)[19]

- **Exploration of the context of the dispute: Index of scholars and List of publications** which enables the exploration of scholars involved in the debate[20] and their publications (as outlined in Section 5.2.2)[21]

[19] `https://projects.dharc.unibo.it/broast/documents/`
[20] `https://projects.dharc.unibo.it/broast/scholars/`
[21] `https://projects.dharc.unibo.it/broast/publications/`

- **Collection Insights**: A set of visualisations and SPARQL queries on the critical inquiries performed on the collection[22]

5.4.1 Exploration of the contents of the dispute: Document's catalogue

The catalogue enables the exploration of the collection documents through the lenses of the inquiries performed on them and their historical evolution. It relies on two main components: the list of archival documents with their title and abstracts (regest) and a set of filters acting on the catalogue. Each filter refines the list of documents concerning the inquiries performed by the scholars to determine the documents' authenticity. In particular, each filter relies on two main components:

- The **metadata fields** describing the documents in the collection. They are identified by object properties from the data model (is part of `dct:isPartOf`, type `rdf:Type`, creator `dct:creator`, date `dct:date`, coverage `dct:coverage`). Such metadata fields have been chosen manually and declared in the back-end as a Python list.

- The **logical statuses** associated to the inquiries. Available logical statuses are the same as declared in Section 4.3 and, in particular, in Table 4.3, namely valid, disputed, superseded, settled, and undisputed. With Conjectures, these logical statuses can be queried by addressing (`GRAPH` keyword for valid opinions, `CONJ` for disputed opinions, `SETT` for settled opinions) or can be retrieved by the interaction between different graph types (Conjectures which are not settled (`CONJ-SETT`) for superseded opinions and graphs which are not settled (`GRAPH-SETT`) for undisputed claims). These represent the possible categorisations of inquiries concerning their evolution through history. They are declared in the back-end via a template (Python list of dictionaries).

[22] `https://projects.dharc.unibo.it/broast/insights`

Two main SPARQL queries build the catalogue in the back-end. The first checks all available logical statuses for each metadata field. For instance, the ASK query, shown in Listing 5.5, checks if settled inquiries concerning the categorisation of documents exist.

```
ASK {
    ?s a fabio:ArchivalDocument.
    SETT ?g {
        ?s  <http://www.w3.org/1999/02/22-rdf-syntax-ns#type> ?o .
    }
}
```

Listing 5.5: ASK query to check if valid opinions concerning the typisation of documents exist

The result of this dynamic query is a matrix, shown in Table 5.2. The interface shows the filter if a `True` value is retrieved, enabling 13 available perspectives on 4 metadata fields. For instance, dates occur in valid, disputed, settled and superseded inquiries. Similarly, Documents' categorisations (Type) occur in all available inquiries' logical statuses since they are both part of the dispute (e.g., the document is a forgery or authentic) and undisputed (e.g., the document is an archival document).

	Valid GRAPH	Disputed CONJ	Settled SETT	Undisputed GRAPH-SETT	Superseded CONJ-SETT
Type	True	True	True	True	True
Creator	True	False	True	False	False
Date	True	True	True	False	True
Coverage	False	True	False	False	True

Table 5.2: Intersection between available logical statuses (columns) for declared metadata fields (rows) in Haider's collection

The second query, shown in Listing 5.6, selects all archival documents (`?s`, instances of the class `fabio:ArchivalDocument`), their title (`?title`) and abstract (`?abstract`) in English.

```
SELECT DISTINCT ?s ?o ?abstract (COUNT(DISTINCT ?g) as ?n)
```

```
WHERE {
?s dct:title ?o.
FILTER (lang(?title) = 'en')
OPTIONAL {
    ?s dct:description ?abstract.
    FILTER (lang(?abstract) = 'en')
}.
}
GROUP BY ?s ?o ?abstract ORDER BY DESC (?n)
```

Listing 5.6: SELECT query to retrieve all documents in the collection with their title and abstract (regest)

Upon page load, two query results are sent to the front-end, displaying the catalogue in Figure 5.7. In particular, the available filters are displayed on the left (the result of the SPARQL query shown in Listing 5.5) with all filters deactivated. Additionally, the list of documents is displayed on the right (the result of the SPARQL query shown in Listing 5.6). The number of retrieved documents is reported on the top-right (153). Each available filter can be activated via a switch button. By default, the switch toggles all valid opinions about the available metadata fields (e.g., all valid dates) as usual in online catalogues.

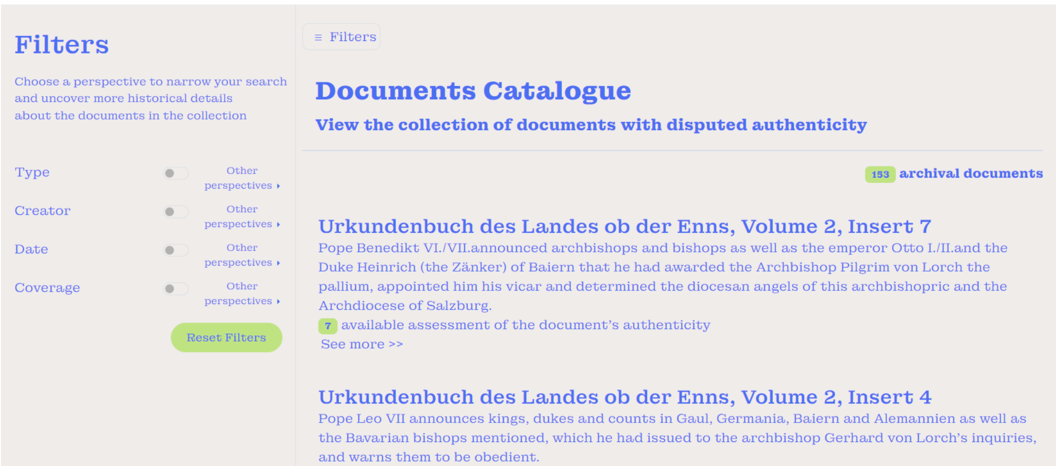

Figure 5.7: Screenshot of the documents' catalogue filtering interface

Other logical statuses can be enabled using the "Other Perspectives" button (settled, disputed, superseded)[23]. The dropdown menu shows only available logical statues for the selected information corresponding to the "True" values from the matrix in Table 5.2. For instance, Figure 5.8 shows the available logical statuses for the metadata field "date" (superseded, disputed, settled). Whenever a switch button is toggled, a `filter` function is called, and the selected filter (e.g., disputed dates) is sent to the back-end via two parameters: selected logical status (e.g., `disputed`) and selected metadata field (e.g., `dct:date`).

Figure 5.8: Screenshot of the documents' catalogue filtering interface with toggled "Other perspectives" button. It shows available logical statuses (superseded, disputed, settled, of inquiries concerning the dating of documents (date)

In the back-end, the first parameter is mapped with a logical statuses template. This associates the parameter ID (e.g., `disputed`) with the corresponding Conjectures keyword (e.g., `CONJ`) and eventually a filter (e.g., `currently_disputed` is mapped with the keyword `CONJ` and the filter `?g2 conj:settles ?g`). Parameters are added to a SPARQL query template, which is run against the endpoint. For instance, the SPARQL query in Listing 5.7, shows one of the possible queries generated by this templating system if the parameters `disputed` and `dct:date` are received.

[23]The undisputed category of claims is missing due the lack of occurrences for undisputed ratings in the dataset

```
SELECT ?s ?title ?abstract ?obj ?obj_label ?start
WHERE {
?s a fabio:ArchivalDocument ;
    dct:title ?title ;
    dct:description ?abstract.
    FILTER (lang(?title) = 'en')
    FILTER (lang(?abstract) = 'en')

    CONJ ?g { ?s dct:date ?obj }
    ?obj rdfs:label|dct:title ?obj_label.
    OPTIONAL {?obj time:hasBeginning ?start}
}
GROUP BY ?s ?title ?abstract ?obj ?obj_label ?start ORDER BY ASC (?start) ASC (?
    obj_label)
```

Listing 5.7: SELECT query to retrieve all dating attempts (date) expressed on the documents which are part of the debate (CONJ)

The SPARQL query (5.7) selects all disputed dates (`CONJ` graph including a `dct:date` predicate) retrieving the documents (subject of the opinion, `?s`), with their title (`?title`) and abstract (`?description`) in English (`FILTER (lang(?title) = 'en')` and `FILTER (lang(?abstract) = 'en')`), as well as the values declared by the opinion (`?obj`) with their labels (`?obj_label`). Optionally, the query retrieves also the date beginning (`OPTIONAL {?obj time:hasBeginning ?start}`). Finally, the results are sorted chronologically (if the value is a date) or alphabetically `ORDER BY ASC (?start) ASC (?obj_label)`.

Results of the query in Listing 5.7 are sent back to the front-end and parsed to fit the interface structure. Each document (e.g., *Urkundenbuch des Landes ob der Enns, Volume 2, Insert 1*) is displayed with the list of all concurring selected values (e.g., dates 777 and 970-971) in green badges as shown on the right of Figure 5.9. Values are displayed in a list in the "Filters" section (on the left) they are sorted chronologically (dates) or alphabetically (all other filter categories).

Figure 5.9: Screenshot of the document catalogue filtering interface with the accepted creator filter activated

Values can then be used to further filter the documents. This interaction is handled in front-end hiding elements which do not include the selected value. For instance, by selecting the date "777" the documents involved in the opinion are shown and the selected value is highlighted with respect to their concurring opinions, as shown in Figure 5.10.

Figure 5.10: Screenshot of the document catalogue filtering interface with the accepted creator filter activated and selected value (777)

The description of the archival documents is stored in a separate page (handled by `document.py` route and `document.html` template) which can be reached by the "See more" button below each document record in the catalogue.

Each document description contains three main elements: descriptive metadata of the document (undisputed), a list of related publications and a section with opinions on the evolution of the document's authenticity through history. The page shows existing undisputed opinions about the documents retrieved in the back-end via a SPARQL query matching the Conjectures pattern `GRAPH-SETT`. For instance, as shown at the top of Figure 5.11, the document (`doc:12`) has title *Urkundenbuch des Landes ob der Enns, Volume 2, number 34*, is an archival document and a copy, is part of "Urkundenbuch des Landes ob der Enns. Bd. 2: 777-1230" and is available in monasterium.net[24].

Additionally, the page showcases the list of related publications addressing the document (Section *Related Publications* at the left of Figure 5.11).

The section *Evolution of document's authenticity* reports the inquiries addressing the documents' authenticity through history. This information is retrieved in the back-end by querying all inquiries part of the dispute on the document (`CONJ ?g {doc:12 ?p ?o}`), as shown in Listing 5.8. Additionally inquiries' contextual information are retrieved, in particular the publication from which the inquiry is extracted (`prov:wasQuotedFrom`) and the date of publication (`dct:date`) and its author (`prov:wasAttributedTo`) and sorted in descending chronological order. The agreement between the settled opinion and other inquiries is retrieved via `OPTIONAL {?g2 conj:settles ?g}`. Inquiries not fitting the condition (therefore, not being settled by other opinions) are considered disbelieved by Haider's settled opinion.

```
SELECT *
WHERE {
    CONJ ?g { doc:12 ?p ?o . }
    OPTIONAL {?g2 conj:settles ?g}.
```

[24] https://www.monasterium.net//mom/AT-StiASF/StFlorianCanReg/0900-0901_I_19/charter?q=O%C3%96UB%202%20(Wien%201856),%2046

```
    OPTIONAL {
        ?g prov:wasQuotedFrom ?ref ;
            prov:wasAttributedTo ?author .
        ?ref dct:date ?date.
        ?date time:hasBeginning ?beginning
        }
    }
    ORDER BY DESC (?beginning)
```

Listing 5.8: SPARQL query retrieving disputed opinions addressing the date of the documents in the collection

Additionally, the SPARQL query shown in Listing 5.9 retrieves the settled inquiry (SETT).

```
SELECT DISTINCT ?g
    WHERE {
        SETT ?g { doc:12 ?p ?o . }
    }
LIMIT 1
```

Listing 5.9: SPARQL query retrieving settled inquiry for document number 12

In the interface shown in Figure 5.11, the document *Urkundenbuch des Landes ob der Enns, Volume 2, number 34* (doc:12) the inquiries (retrieved from query 5.8) are shown with their contents and their provenance information. They are sorted in descendant chronological order (from the latest to the oldest depending on the date of publication of their source). The settled inquiry is represented with a star on the top-left, is quoted from *Verzeichnis der den oberösterreichischen Raum betreffenden gefälschten, manipulierten oder verdächtigten mittelalterlichen Urkunden* published in 2022 by Siegfried Haider and states the document as a forgery, made between 950 and 1000, by Bishop Pilgrim of Passau. Additionally, agreeing inquiries (therefore stating the document is a forgery) are marked with a green border, while disagreeing inquiries are marked with red border (therefore believing the document as authentic).

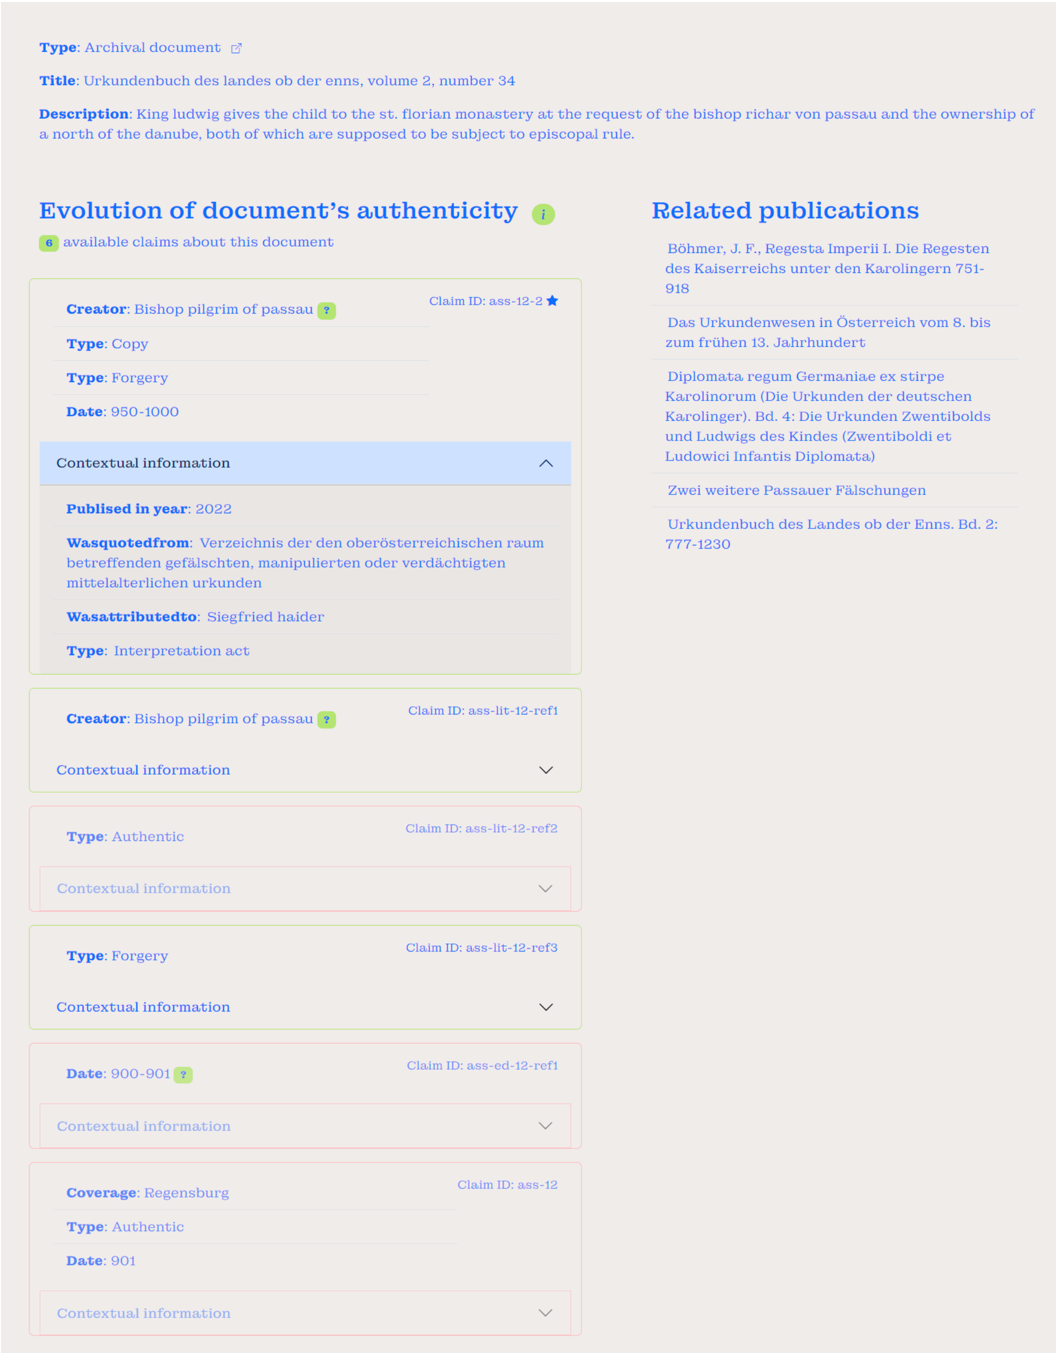

Figure 5.11: Screenshot of *Urkundenbuch des Landes ob der Enns, Volume 2, number 34* (item 12) document page in roast web-application

For instance, `doc:12` includes six opinions, with three classifying the document as a forgery and three affirming its authenticity. The evolution of the scholarly opinion towards the document authenticity is reflects by the chronological sorting of the related opinions, revealing how perspectives on the document's authenticity have developed through centuries.

5.4.2 Exploration of the context of the dispute: Index of scholars and List of publications

Broast consists of two key components: an index of scholars representing the actors of the debate and a list of publications that tackle the discussion. Together, these elements offer a comprehensive view of the academic landscape related to the documents.

The index of scholars collects and displays the scholars credited in the bibliographic references documenting the collection in alphabetical order. In particular, the index is retrieved by querying all scholars (`?scholar`) related to the publications (`dct:creator, dct:contributor, dct:publisher`) stored in the collection. Additionally, the query retrieves the number of publications (`?refCount`) and the number of archival documents (`?docCount`) addressed by the publications, as shown in Listing 5.10.

```
SELECT DISTINCT ?scholar ?label (COUNT(distinct ?ref) as ?refCount) (COUNT(distinct
    ?document) as ?docsCount)
WHERE {
    VALUES (?property) {(dct:creator) (dct:contributor) (dct:publisher)}
    ?ref ?property ?scholar ;
         cito:documents ?document .
    ?scholar rdfs:label ?label .
}
GROUP BY ?scholar ?label ORDER BY ASC (?label)
```

Listing 5.10: SPARQL query retrieving the index of scholars

For instance, as shown in Figure 5.12, Engelbert Mühlbacher is responsible for 1 publication addressing 10 archival documents of the collection.

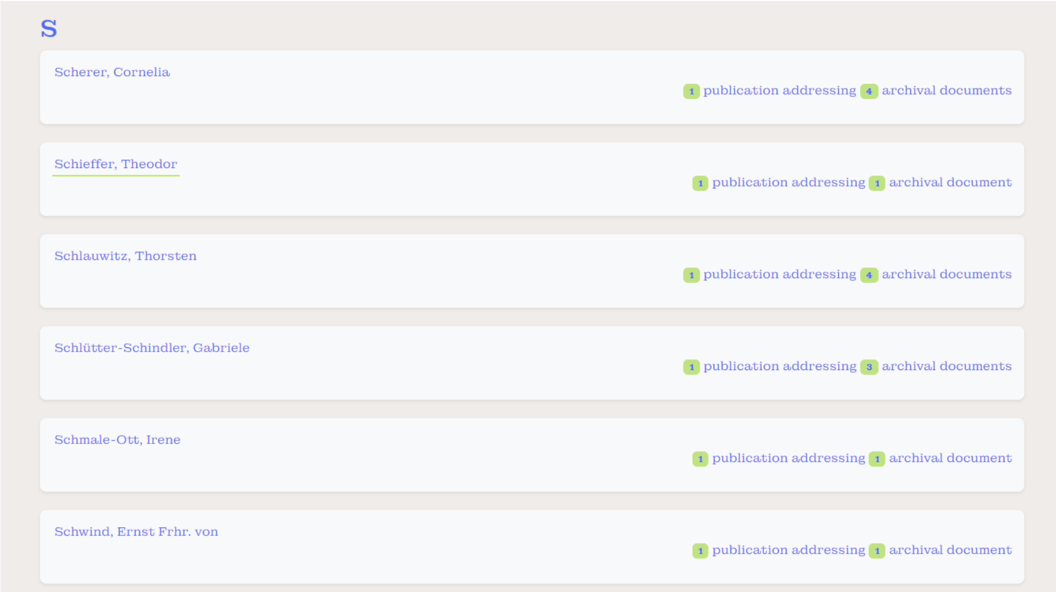

Figure 5.12: Screenshot of Engelbert Mühlbacher scholar page in Broast web-application

Each scholar is characterised by a dedicated view showing the scholar's biographical information, available publications, related scholars and archival documents in the collection. For instance, Engelbert Mühlbacher (1843 Gresten, 1903 Vienna) is a medievalist, historian and university teacher, as shown in Figure 5.13[25]. Three of his publications address the collection. In particular, he is the creator of the article *Zwei weitere Passauer Fälschungen*, contributor and publisher of two archival records sets *Regesta Imperii 1* and *Die Urkunden der Karolinger, Bd. 1*. Additionally, Mühlbacher is related to other scholars of the collection as Heinrich Fichtenau, Johann Friedrich Böhmer, Johann Lechner, Albert Brackmann, and Alfons Brackmann[26]. Mühlbacher is therefore related to 10 archival documents of the collection[27].

[25]Biographical information is extracted on the fly from Wikidata via the Wikidata API (`https://www.wikidata.org/w/api.php`)

[26]Relations between scholars are calculated with a SPARQL query retrieving other scholars' publications citing the same archival documents. Results are ranked by the number of common documents

[27]Relations between the scholar and archival documents in the collection are calculated with a SPARQL query retrieving archival documents addressed by the scholar's publications

Figure 5.13: Screenshot of the index of scholars in Broast web-application

The list of publications collects, as shown in Figure 5.14 and displays the publications about the documents. These publications are organised by category, which includes anthologies, archival record sets, articles, books, and book chapters. Similar to the scholars' index, this list is generated by querying all publications that belong to the specified categories, along with their titles.

Figure 5.14: Screenshot of the List of publications in Broast web-application

Finally, a dedicated page is available for each publication about the collection.

For instance, Mühlbacher's article *Zwei weitere Passauer Fälschungen* is shown in Figure 5.11. Each publication is displayed with the metadata extracted from the OPAC Regesta Imperii. The article is shown with all the relevant descriptive metadata, a reference to the OPAC Regesta Imperii record, and a link to the full text and the documents of the addressed in the publication. [28]. Mühlbacher is therefore related to 10 archival documents of the collection[29].

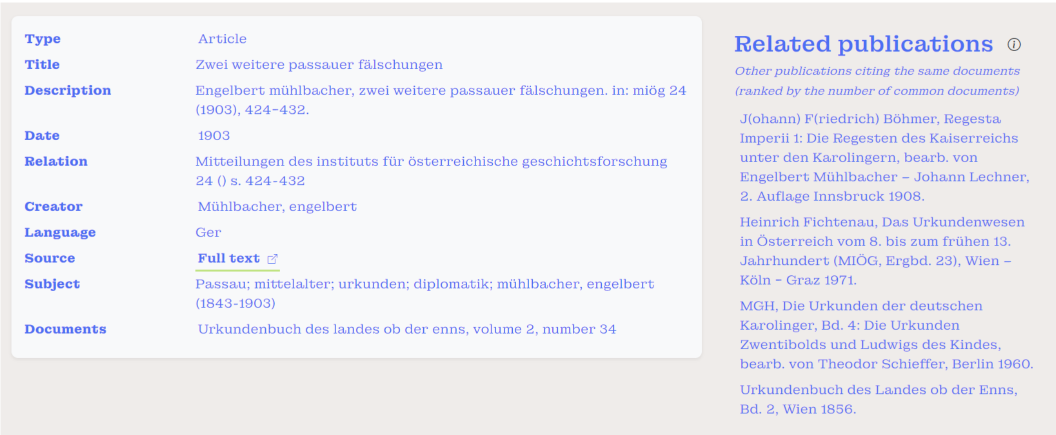

Figure 5.15: Screenshot of the publication "Zwei weitere Passauer Fälschungen" in Broast web-application

5.4.3 Collection insights

Three main visualisations constitute the insights on the collection, mainly addressing spatial (location of creation) and temporal information (year of creation) about the documents and temporal information about the scholarly opinion addressing the documents (year of publication).

A map, shown in Figure 5.16, illustrates the alleged locations associated with the creation of historical documents whose authenticity has been disputed, revealing the geographical context of these documents and the debates surrounding their provenance.

[28]Relations between scholars are calculated with a SPARQL query retrieving other scholars' publications citing the same archival documents. Results are ranked by the number of common documents

[29]Relations between the scholar and archival documents in the collection are calculated with a SPARQL query retrieving archival documents addressed by the scholar's publications

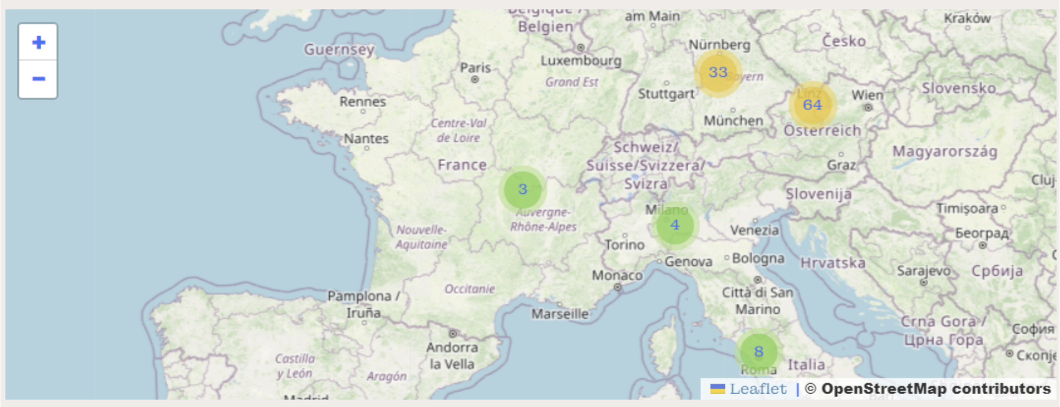

Figure 5.16: Map of the alleged places of creation of the documents in the collection

Temporal information (year of creation) of the documents are retrieved in the back-end by the means of two SPARQL queries, respectively addressing the alleged perspective on documents (documents believed authentic, now deemed superseded) as shown in Listing 5.12 and the settled perspective on documents (documents believed forged) as shown in Listing 5.11.

According to the scholarly community, the first (Listing 5.12) retrieves the number of documents for each century according to the earliest dates provided by the settled sources. Since this study assumes Haider's perspective as the most authoritative and up-to-date source, this will inherently represent the scholar's perspective. Listing 5.11, retrieves all settled dates (`SETT ?g {?s dct:date ?date}`) stating the document as a forgery (`sebi:Forgery` and the related century (`?centuryStart`). The latter is calculated over the start year (`?startYear`) divided by 100, using the FLOOR function to round down to the nearest whole number and then multiply by 100.

Similarly, the second query (Listing 5.12) retrieves the number of documents for each century according to the information now deemed questioned by the community (usually provided by the documents themselves). Listing 5.11 retrieves all superseded dates (`CONJ ?g {?s dct:date ?date}`) claiming the document as authentic (`sebi:Authentic`) and the filter (`FILTER NOT EXISTS ?g2 conj:settles ?g`).

```
SELECT ?label (COUNT(?date) AS ?n)
 WHERE {
 SETT ?g {
    ?s dct:date ?date ;
        a sebi:Forgery .
    }
 ?date a time:Interval ;
 time:hasBeginning ?start ;
 time:hasEnd ?end .

 BIND(YEAR(?start) AS ?startYear)
 BIND(FLOOR(?startYear / 100) * 100 AS ?
    centuryStart)
 BIND(?centuryStart AS ?label)
 }
 GROUP BY ?label
 ORDER BY ?label
```

Listing 5.11: SPARQL query to retrieve all superseded datings of documents deemed to be authentic

```
SELECT ?label (COUNT(?date) AS ?n)
WHERE {
CONJ ?g {
    ?s dct:date ?date ;
    ?s a sebi:Authentic .
}
?date a time:Interval ;
time:hasBeginning ?start ;
time:hasEnd ?end .

BIND(YEAR(?start) AS ?startYear)
BIND(FLOOR(?startYear / 100) * 100 AS ?
    centuryStart)
BIND(?centuryStart AS ?label)
FILTER NOT EXISTS {?g2 conj:settles ?g}}
GROUP BY ?label
ORDER BY ?label
```

Listing 5.12: SPARQL query to retrieve all settled datings of documents deemed to be a forgery

Figure 5.17 shows respectively the results of Listings 5.12 (in purple) and 5.11 (in green). The x-axis shows the century of creation of the documents, and the y-axis shows the number of documents normalised under the total scholarly opinions for each century. Results show two different progressions of documents' creation dates depending on the two perspectives. As a result, the superseded perspective on the data (acknowledging the documents as authentic) shows a notable peak in documents' production around 1100. Differently, acknowledging the settled perspective on the documents as forgeries, the peak is registered one century later (1200). In simpler terms, most of the forgeries in the collection were produced in the 13th century. Still, they were either backdated or entirely made to appear as if they were created in the 12th century, showing a 100-year shift between superseded and settled perspectives on documents. Unfortunately, the catalogue lacks information on why these documents were forged. Understanding these motivations would provide additional context for this analysis.

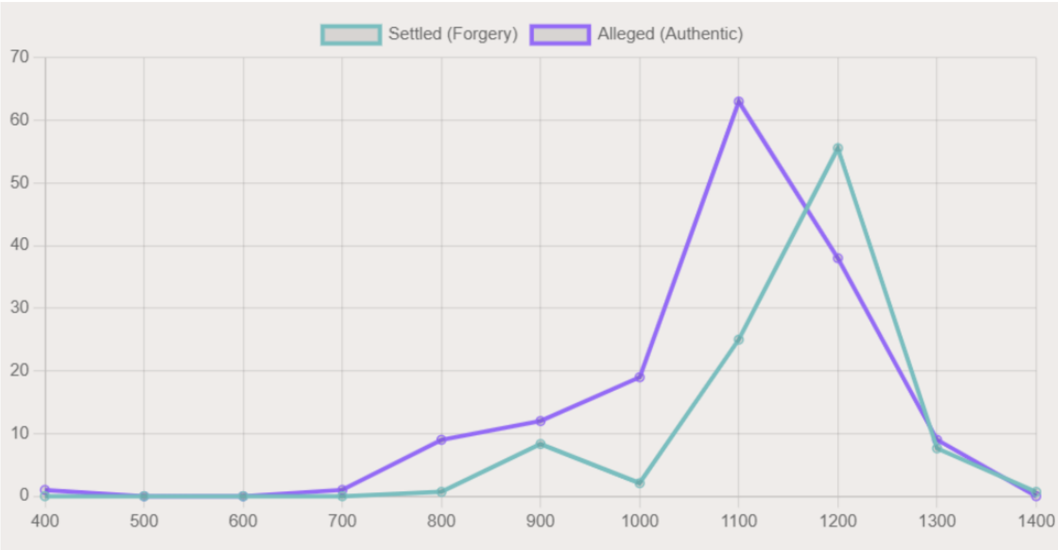

Figure 5.17: Line-chart showing settled and superseded perspectives on the most productive centuries (earliest date) for the documents of the collection

Scholarly opinions have trended throughout history. The dotted chart, shown in Figure 5.18, illustrates the evolution of scholarly opinions regarding the classification of the documents in the collection. The x-axis shows the date of publication in which the scholarly opinion is stated. The y-axis shows the Document ID in the collection. Overall, the trend shows that established views classifying the document as forged (indicated in green) began to emerge around 1880 and have steadily increased in prevalence up to the present day[30]. Despite this, several scholarly opinions addressing the documents' authenticity (indicated in purple) continue to exist, highlighting the ongoing debate over their categorisation.

[30]Note that Haider's scholarly opinions have been removed for clarity

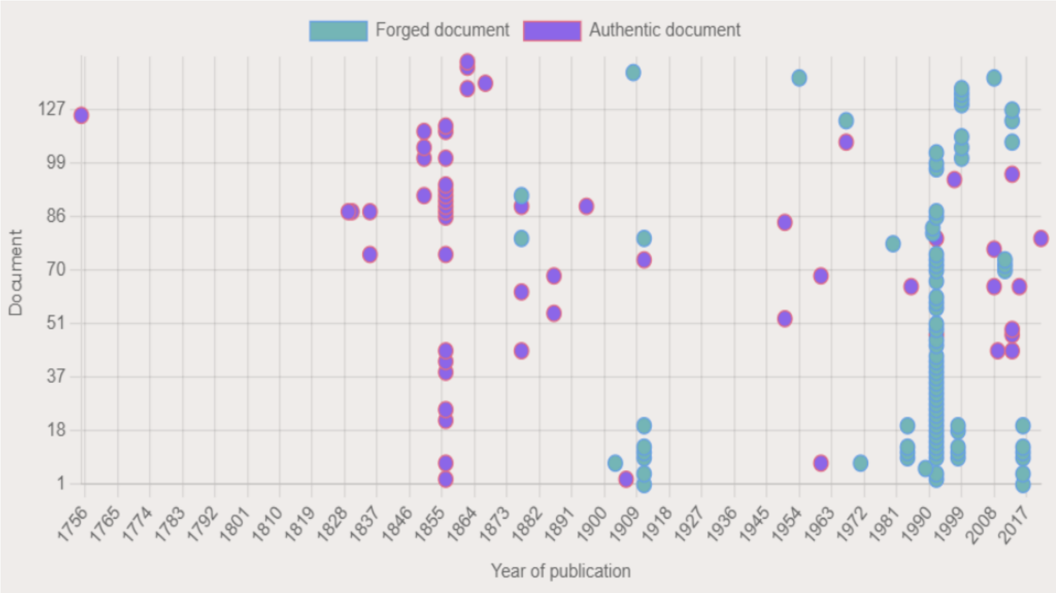

Figure 5.18: Chart showing the evolution of scholarly opinions towards the classification of the documents in the collection

Additionally, the web page provides a range of analytical insights derived from SPARQL queries run in the back-end, comprehending various trends related to opinion contents and their logical status (e.g., settled and superseded opinions about the most productive centuries) document debate (e.g., most debated documents), scholarly productivity and citation (e.g., most productive scholars, most cited scholars, number of scholarly opinions per scholar), evidence usage (e.g., most accounted features, most accounted features types, most performed evaluations, most used evidence), and temporal aspects of referenced publications (e.g., newest publications on the documents, oldest publications on the documents). Each insight has a title, a brief description, the results and their counting, and the related SPARQL query syntax.

For instance, in addressing the question, "Which are the most accounted features in addressing fraudolence?" the SPARQL query detailed in Listing 5.13 retrieves the most reported evidence is interpolations, followed by chronology and dating as shown in Table 5.3. Despite this, the poor number of argumentations does not allow us to make any additional consideration of the data due to the poor number of occurrences.

```
SELECT ?label (COUNT(?feature) AS ?n)
WHERE {
    CONJ ?g { ?s a forgont:Forgery }

    ?g a hico:InterpretationAct.
    ?evidence sebi:account ?feature;
      sebi:support ?g.
    ?feature rdfs:label ?label
```

Listing 5.13: SPARQL query retrieving the most accounted documentary features used to assess a document's authenticity

label	n
interpolations	17
chronology and dating	7
authentication marks	6
list of witnesses	5
original document	5
addition	3
metre and style	2
content	1
document provenance	1

Table 5.3: Results obtained by the SPARQL query retrieving the most accounted documentary features used to assess a document's authenticity

5.5 Discussion

This chapter analysed several aspects related to the second research question of this Dissertation, RQ2 - *How can critical inquiry be effectively and efficiently represented with minimal complexity?*.

In particular, concerning RQ2e - *Which are the typical patterns that scholars adopt to challenge a document's authenticity?* - the analysis of the source materials demonstrated that several relevant patterns scholars adopt to challenge a document's authenticity are critically related to past and superseded perspectives on the studies subject.

The catalogue, *Index of Medieval Documents Concerning the Upper Austrian Region that is Damaged, Tampered with, or Altered*[31] [52], compiled by Siegfried Haider, presents 153 inquiries challenging the authenticity of historical documents. It documents relevant aspects of the scholarly practice of gathering and evaluating information from multiple perspectives to produce well-reasoned analysis and understanding. In the catalogue, several relevant aspects for the topics of this Dissertation are considered: alleged information about the document (e.g., declared or alleged date and place of creation and underlying juridical categorisation as authentic), proposed information about the document contrasting alleged knowledge (e.g., dating attempt about the date of creation of the document), study and report of existing secondary sources documenting the subject (e.g., archival record sets reporting the document under analysis), notes on the inquiry (e.g., collected evidence, documents' classification concerning its transmission), previous inquiries on the documents (e.g., existing inquiries tackling the discussing the document's authenticity). In summary, the catalogue demonstrated to be a valuable case study for examining various prevalent issues, particularly the coexistence of contrasting opinions, their logical statuses as determined by the community consensus, and the sources and arguments supporting critical inquiry.

A set of patterns related to critical inquiry practices have been identified from the surveyed materials:

- Critical inquiry on documents' authenticity consists in the challenge of alleged information and in the proposal of new information potentially replacing them (e.g., allegedly produced in 901, now thought to be produced in the 2nd half of the 10th century)

[31] Original title in German *Verzeichnis der den oberösterreichischen Raum betreffenden gefälschten, manipulierten oder verdächtigen mittelalterlichen Urkunden: ein Arbeitsbehelf*

- Questioning a document's authenticity implies negating the alleged information about the document itself. Such information includes its juridical categorisation (authentic) and possibly other aspects of the document, such as its date and place of creation and creator.

- Scholars support their opinions by collecting argumentations and evidence (e.g., the presence of interpolations in the document)

- Each opinion has a current status in the critical debate over the same topic (e.g. the document is nowadays considered a forgery, or the document has been considered authentic, but this information is currently considered superseded by the community)

- The perception of the document's authenticity through history changes through time, revealing disagreements between scholars (e.g., the document allegedly produced in 901 has been declared a forgery by Mühlbacher in his article in 1903, therefore addressed as authentic by Schieffer in 1963).

Therefore, such patterns have been used to define a data model, construct a KG, and test it through a series of dynamic queries implemented in a web application that demonstrates various interactions with the contents. RQ2f - *How can EWA be introduced to represent such patterns?* - further tackles these aspects.

In addressing RQ2f, the findings demonstrate that combining document-centric ontologies with the EWA approach provides a structured and flexible means of modelling competing scholarly inquiries challenging documents' authenticity.

Domain ontologies (Dublin Core and Fabio) have been reused to describe opinions related to the documents' information despite their being superseded, declared by the document, and proposed by a scholarly inquiry.

Additionally, the data model presented in this chapter (Section 5.2) considers the inherent fuzziness arising from historical primary sources (for example, dating attempts modelled using Time ontology). Besides this, a set of custom classes `sebi:Forgery`, `sebi:Authentic` and `sebi:SuspectedForgery`) and properties have been introduced through the SEBI ontology to address the juridical status of the documents. Each coherent set of information has been grouped via named graphs, thus avoiding the introduction of fictitious entities to represent such information (e.g., `crm:E13_Attribute_Assignment`). This approach also avoids the need for SPARQL queries to retrieve all single opinions (e.g., date, creator) tackling the authenticity of the same document while sharing the same provenance (e.g., same source).

This work represents a set of relevant provenance information (e.g., argumentations, cited sources) to contextualise the critical inquiry performed by scholars. Such representation includes inquiry attribution (e.g., `prov:wasAttributedTo prov:Agent`), the source of the inquiry (e.g., `prov:wasQuotedFrom` to link the inquiry to related sources, FaBio ontology to define them and Dublin Core to represent their descriptive metadata), type of analysis (e.g., `hico:InterpretationType`, not tackled by this dataset, but addressed by the data model), and argumentations supporting the conclusions of the inquiry (`sebi:Evidence` as a specification of `hico:InterpretationCriterion`). SEBI ontology allows the representation (for instance, see the SPARQL query in Listing 5.13) of evidence (`sebi:Evidence`) by associating an evaluation (`sebi:Evaluation`, e.g., coherency) to a certain feature concerning the document or its history (`sebi:Feature`, e.g., style). The property `sebi:hasEvaluationScore` allows for the potential negation of elements forming the evidence, as the process of challenging involves either accepting or negating an evaluation of a particular feature (e.g., `coherency False` indicates stylistic incoherence).

The scholar index and the publication list provided in the Broast web-app demonstrate how contextual information can enhance navigation within the KG, offering a novel perspective on the data. This aspect is non-trivial as it considers scholarly work in the first place. Details on scholars' engagement in the critical debate, their roles, and the references they have worked on are crucial for humanities models, catalogues, and platforms, as their role in the critical discourse is as significant as the inquiries they develop. In a broader context, interpretations in the humanities should always be represented alongside this related information.

EWA—achieved with Conjectures—introduce concurring opinions thereby preventing the emergence of conflicting and incompatible stances in the KG. In particular, settled and conjectural graphs introduce the evolution of opinions through history (allegedly authentic, currently believed to be a forgery) and define the current status of the critical debate concerning the documents (superseded as authentic, settled as a forgery).
EWA extends the domain ontologies' expressivity by representing critical inquiry and its evolution (e.g., settled date and superseded date). Hiders' perspective has been chosen as the settled point of view over the dispute, therefore representing the current status of the debate on each document. Comparisons between superseded and settled perspectives on documents' juridical classifications (allegedly authentic versus settled fraudulent document) can be therefore performed as demonstrated by the results of the queries shown in Listings 5.11 and 5.12. Results showcase a 100-years shift between the alleged (now, superseded) and its settled creation date.

The categorisation of undisputed, superseded, and settled opinions (represented in this work respectively via plain named graphs, conjectural graphs and settled conjectures) allow for the intersection of such knowledge categories (also called logical statuses) over the aspects tackled by the inquiries (metadata describing the archival documents, introduced by domain ontologies).

The Broast web-app catalogue, described in Section 5.4.1 explores such representation by intersecting the opinions' logical statuses — valid, settled, disputed, undisputed and superseded — and the aspects tackled by the inquiries — creator, date, coverage, type. The catalogue in Broast provides proof of concept on how the inclusion of EWA can be leveraged to model scholarly challenges and evolving knowledge in LOD catalogues. Additionally, the dynamic query execution implemented in the web-app catalogue tests the desired predicates and their logical status, allowing systematic reviewing of the data model and the KG correctness. The Broast web-app aims to supersede the flat representation of historical records usually adopted in online catalogues.

In the scope of this case study, it is given that the dispute concentrates on two main perspectives on the documents' juridical classification (authentic vs. forgery or suspected forgery), and given that such documents are nowadays classified as fraudulent, as described in Section 5.3. As exemplified in Figure 5.5. Opinions have been categorised depending on their agreement with Haider's concurrent opinion. Such categorisation does not rely only on the mere statement of the document as authentic or forged but can also be retrieved by other aspects tackled by the inquiries (e.g., creation dates, creators) and their proximity to the information stated by the one or other perspective. This categorisation reveals the perception of the document's authenticity through history according to the settled point of view on the dispute, revealing disagreements between scholars.
The section *Evolution of document's authenticity* in the document dedicated page of the Broast web-app showcases this aspect (Section 5.4, Figure 5.11). As, for instance, the perception of the 12th document of the collection has been shifted many times through history as allegedly produced in 901, then declared as a forgery by Mühlbacher in his article in 1903, therefore addressed as authentic by Schieffer in 1963 and so on. This formalisation indicates the

degree of consensus or alignment among sources on a given settled agreement on a certain document. These relations have been represented via the means of the object property `conj:settles`, as described in Section 5.3.

The recording and structuring of critical discourse and scholarly observation reveal the evolution of knowledge towards a specific topic and, therefore, reconstruct and preserve the scholarly narrative developed through history.

Despite some promising results in knowledge representation, this work presents some limitations. In the case study, Haider's inquiries have been defined as settled by considering the temporal aspect of publications. Therefore, the agreement or disagreement with the settled inquiries has been modelled accordingly. More sophisticated systems may be required to define a settled perspective by considering several contextual aspects, such as ranking their motivations and criteria, relations with cited sources, and scholars' backgrounds. Similarly, the representation of forged documents implies two contrasting points of view: the document is considered authentic or is recognised as a forgery (or at least a suspicious forgery). Therefore, only two perspectives can be created on data, so inquiries can be divided into agreeing and disagreeing sources. In other domains, this distinction can be more nuanced and more sophisticated systems may be required to define agreements and disagreements.
While an effective proof-of-concept for testing the data model and using EWA, the current catalogue is inherently limited by scope and size. As a result, it cannot fully support comprehensive analyses of forgery patterns or the broader historiographical shifts related to these phenomena. A larger and more diverse dataset is required to generate meaningful insights into critical debates, the evolution of scholarly inquiries, and methodological approaches that scholars have employed over time. While this work demonstrates the potential of including EWA methods and, in particular, Conjectures, its constraints necessitate the inclusion of a more expansive dataset to enable the detailed, systematic

study of forgery trends and the critical discourse surrounding these works over time.

Chapter 6

Conclusions and Future Work

This Dissertation explored how humanities critical inquiries are formally represented in RDF, emphasising how scholarly narratives evolve over time, grounded in and influenced by pre-existing knowledge in the field. In particular, two hypotheses have been proposed in the research statement of this work (Section 1.4), which are discussed as follows.

(H1) - *Annotators in CH may approach data observations, incompleteness and uncertainty differently, favouring qualitative, context-rich representations of competing hypotheses.* The analysis conducted on Wikidata (Chapter 3) demonstrates that CH annotators use the surveyed WLS methods more heterogeneously than those working with astronomical data, which appears to rely on rankings (and thus asserted and non-asserted claims) to signify new information as knowledge in the field evolves. This reflects the epistemological differences highlighted in the theoretical framework (Section 1.1) and demonstrates that historiographical uncertainty should not be minimised but organised to accurately represent CH critical inquiry and observations as situated constructive knowledge.

(H2) - *A comprehensive representation of scholarly opinions can reflect the evolution of critical inquiry on a given topic.* Given that historiographical uncertainty should neither be diminished nor omitted in CH discourse, critical inquiry can be organised around three main concepts: (1) the representation of scholarly opinions (observations and conclusions of the scholar); (2) the context of the opinion (methods, collected evidence, surveyed primary and secondary sources); and (3) the categorisation of the opinion concerning the scholarly context in which it occurs (e.g., whether the community has dismissed the scholarly opinion, implies a WLS as a hypothesis, or is established within the community).

The research questions tackled by this Dissertation have been addressed in the discussions within each chapter. This section summarises the key findings already explored throughout the thesis.

RQ1 - *What are the current limitations and potential improvements in the representation of humanities critical inquiry in current Knowledge Graphs?*

CH KGs have made strides by incorporating reification methods like n-ary relations and named graphs to represent critical inquiry and argumentations, uncertainty and interpretative contents. In particular, this work recognised nine main features in addressing the representation of critical inquiry in CH KGs (Section 2.3), namely: factual knowledge (data), scholarly opinion as a questionable entity (capta), scholarly method, studied sources, confidence, reification method, complex scholarly opinions, competing scholarly opinions, opinions logical status. CH data models often fail to provide a solution for representing competing scholarly opinions and their logical status (settled and disputed). Ontologies such as CIDOC CRM provide a solid foundation but need more specific guidelines for handling contradictory interpretations, limiting their utility in representing the dynamic nature of historical research and hampering the representation of ambiguous or disputed information. However, such situations are the norm rather than the exception in humanities discourse (Chapter 1).

Wikidata offers a set of representational patterns to cope with uncertain and potentially contradicting knowledge (referred to as WLS claims). The analysis of WLS claims in Wikidata cultural heritage records (Chapter 3, and in particular the discussion Section 3.3.1) demonstrates that uncertainty and evolving knowledge cover a wide range of topics (e.g., dating, attributions, object categorisations) covering all CH object categories (Visual, Audiovisual, Textual). Such inconsistent application of WLS claims across CH datasets introduces ambiguity, making searching, retrieving, and interpreting WLS data difficult. In contrast, the same WLS patterns are overlooked in astronomy-related objects, where information updates are the most frequently recorded aspect via ranked statements. This inconsistency is especially problematic in CH datasets, where historiographical uncertainty and scholarly debate affect many descriptions, such as attributions, dating, and objects' categorisations (H1). Additionally, such WLS patterns should be utilised more in CH datasets, leading to an overemphasis on plain facts. Potential improvements to the Wikidata data model include simplifying the handling of WLS claims, disambiguating their application, and implementing systematic updates to datasets better to reflect the evolving nature of knowledge in CH (Chapter 3, Section 3.4).

In simpler terms, CH KGs use reification methods, such as n-ary relations and named graphs, to represent critical inquiries that do not support EWA. Consequently, the logical status of these scholarly opinions is not explicitly included in the knowledge graph. Still, it can be retrieved by considering contextual information, such as time validity and the certainty of the assertion via SPARQL query retrieval or further processing with some impractical drawbacks (See in particular the discussion of EWA effectiveness in Section 4.2.6). Human intervention is always required to (1) understand definitions of uncertainty, which may differ from one field to another, and (2) to disambiguate statements whose truth value is not explicit. In contrast, Wikidata employs EWA by using its reification method to indicate opinions that have been discarded or superseded. Thus, EWA is proposed as a more streamlined approach for marking WLS claims. This aspect is explored in greater depth in the discussion on the second research question of this work (RQ2).

RQ2 - *How can critical inquiry be effectively and efficiently represented with minimal complexity?*

The main contribution of this work addresses the representation of the evolution of critical inquiry in the humanities using RDF by addressing its constructive, partial perspective. In particular, EWA is recognised as a powerful tool of some ontology-independent reification methods for comprehensively representing CH inquiries and related data. It considers the underlying historiographical uncertainty and ongoing debates within scholarly discourse as a nuanced expression without asserting absolute claims. In particular, EWA streamline the querying and retrieval processes, enabling a clearer distinction between current and outdated interpretations and avoiding the need to search for contextual information (e.g., most updated claim).
Chapter 4 (in particular Section 4.2.6) surveyed the effectiveness of existing reification methods in achieving EWA and highlighted their shortcomings, particularly in representing evolving knowledge (exemplified by the "Girl Reading a Letter at an Open Window") and scholarly challenges (exemplified by the "Donation of Constantine"). EWA offers direct access to claims related to critical inquiry and their logical status. Conjectures, in weak and strong form, is proposed as a framework relying on named graphs and defining them into three distinct types—plain, non-asserted (conjectures), and asserted (settled conjectures) — to represent various types of claims related to critical inquiry and their evolution over time. Conjectures proves to be particularly effective in representing disputes.

Regarding EWA efficiency (Section 4.3), the analysis reveals that while RDF-star and Wikidata are among the most efficient methods, the strong form of Conjectures demonstrates notable performance improvements concerning its weak form. The strong form of Conjectures excels in expressing debates, showing quick retrieval for various queries regarding disputed and settled claims, albeit with a minor performance drop for asserted claims. Additionally, the strong form remains competitive regarding the number of triples and dataset weight, although it exhibits some performance loss in loading times for large datasets.

Chapter 5 presents a case study using EWA to model critical inquiry in the humanities, focusing scholarly challenges on the authenticity of historical documents. Conjectures are used to introduce EWA by addressing undisputed, disputed and settled contents as they have evolved through history, successfully representing concurring and evolving perspectives (as an instance of partial and constructive knowledge). Domain-specific ontologies (such as Dublin Core and FaBiO) are reused to formalise the domain. The model represents evolving debates, avoiding the creation of fictitious entities, allowing SPARQL queries to reflect changing scholarly positions. Ultimately, the EWA approach is leveraged to represent critical inquiry's dynamic, evolving nature in the humanities, particularly related to contested perspectives. The Broast web application demonstrates the model's practicality by enabling users to navigate evolving knowledge of contested documents. It reinforces EWA's role in modelling critical inquiry as it develops. Broast is a prototypical tool for navigating scholarly opinions, highlighting the contributors behind these viewpoints and the sources they engage with and produce.

This case study demonstrates the proof-of-concept for formalising the definitions C[1:6] introduced in Section 1.3, primarily through the implementation of EWA, Conjectures, and domain ontologies.

- (C1, C2) Critical inquiry is often marked by uncertainty, as it stems from the inherent ambiguity in historiographical interpretation. This leads to the concept of conjecture as a formalisation of WLS claims (see the definition of a conjecture in Section 4.2). This phenomenon is represented in the case study by representing all conflicting opinions as conjectural graphs, as shown in Figure 5.3 (Section 5.2.2) representing the opinions of Haider, Mühlbacher, Schieffer and the document 12 itself towards the authenticity of the same document.

- (C3) Each scholarly opinion is supported by evidence, sources, and methodologies that justify the conclusions drawn. The SEBI, PROV-o, and FaBio ontologies formalise this aspect by representing consulted bibliography, collected evidence and provenance information about scholarly opinions as described in Section 5.2.2.

- (C4) The relation of scholarly inquiries to previous opinions, either believing or disbelieving past perspectives, is formalised in Conjectures through explicit settled agreements (as defined in Listing 4.14, Section 4.2.3). This is illustrated in Figure 5.5 (Section 5.2.2), which shows the agreement between Haider and Mühlbacher.

- (C5) Potentially divergent opinions and their respective positions in the ongoing debate are formalised within the Conjectures framework. These views are categorised as either disputed, undisputed, or settled, as detailed in the definition of Conjectures (Section 4.2). As exemplified in Figure 5.3 (Section 5.2.1), Haider's opinion is categorised as a settled conjecture, among other conflicting conjectural graphs. This categorisation aids in retrieving superseded, valid, or unsettled statements, highlighting the evolving nature of knowledge. This categorisation is assessed through queries designed to test Conjectures efficiency (summarised in Table 4.3, Section 4.3), and through the matrix used to develop the Broast catalogue (Table 5.2, Section 5.4.1).

- (C6) A set of fairly common situations can lead to constrasting opinions. The case study explores the representation of differing opinions within scholarly debates often stem from challenging documents' authenticity (C6).

Future developments will move in three main directions, mostly addressing Conjectures and automatic extraction of scholarly disputes and forgeries. Despite Conjectures showing promising results in representing the evolution of humanities critical inquiries, particularly in addressing disputes, the work with the framework has just begun. First, the loading time of Conjectures in the strong form needs further analysis and refinement due to non-competitive results, as discussed in Section 4.3.3. Currently, GraphDB is the only parser accepting Conjectures in their strong form. Future works on Conjectures will extend other Triplestores' grammars (e.g., Blazegraph[1], Apache Jena Fuseki[2]) to accept the framework. Similarly, the Conjectures parser[3] will be the basis for extending existing RDF frameworks, e.g. rdflib. Additionally, the KG produced in Chapter 5 will be converted in RDF-star (in particular Trig-star and Turtle-star) and Wikidata statements, and their effectiveness will be compared with the Conjectures framework.

Section 1.3 documents several approaches that annotators may adopt when no appropriate representation is provided to annotate complex statements (e.g., tackling disputes or interpretative contents), namely reticence, flattening, coercion and dumping. A common practice is to relegate such discussions to free-text notes (e.g., in the description of a CH record). These aspects have not been further explored in this work as they are not in the scope of the research questions. Still, their shortcomings have been discussed throughout this work (e.g., giving a low representation of the phenomenon in structured or semistructured sources). Wikipedia stores several articles describing forged documents, medieval charters, and the disputes that involved them throughout history. These articles include various historical documents, including me-

[1] `https://blazegraph.com/`
[2] `https://jena.apache.org/documentation/fuseki2/`
[3] `https://rdflib.readthedocs.io/en/stable/`

dieval charters, maps, and contemporary books, spanning different centuries (from late antiquity to the contemporary period) and various scholarly debates. Examples include the "Donation of Constantine"[4], the Ireland, Shakespeare forgeries[5], and the Protocols of the Elders of Zion[6]. Recently, the rise of large language models (LLMs) has expanded possibilities for extracting structured data from free text (e.g., text-to-KG alignments). I am currently working on a baseline that leverages LLMs to extract and structure debates over document authenticity assessment.

Recent advancements have expanded the scope of Aspect-Based Sentiment Analysis (ABSA) by applying it to various datasets, including news articles and tweets [38, 47, 53]. The task is also addressed in these works as Target-Dependent Sentiment Classification (TDSC). The introduction of context-aware language models like BERT has significantly improved the performance in this task, setting new benchmarks in the field [47]. Considering the complex language and the data sparseness of the typical humanities discourse, LLMs can be used to widen the possibility of retrieval from unstructured natural language descriptions. Many tests have proved that LLMs perform promisingly for KG generation from unstructured text [70, 69].

The baseline proposes a subclass of TDSC named Open Claim and Stance Classification (OCSC), which is aimed at extracting scholarly inquiries by identifying the claiming entity (e.g., Lorenzo Valla), classifying stances (e.g., positive, negative), identifying the stance related aspects (e.g., hypothesised collected evidence) and classifying the extracted term to a provided terminology (e.g., incoherent language). The baseline will use the data model presented in the last chapter (Chapter 5) as its classes, properties, and named individuals are used to instruct the prompts and validate the extracted data. The extracted dataset will define parameters and metrics to automatically categorise the claims as disputed and settled. This work

[4]https://en.wikipedia.org/wiki/Donation_of_Constantine
[5]https://en.wikipedia.org/wiki/Ireland_Shakespeare_forgeries
[6]https://en.wikipedia.org/wiki/The_Protocols_of_the_Elders_of_Zion

will refine the model and expand its expressivity to represent frauds and hoaxes, hopefully integrating such content in the Broast web application and Wikidata descriptions.

Finally, this work tackled the domain of forgeries only in the scope of the case study represented in Chapter 5. However, the topic includes several intriguing aspects, such as art forgeries, hoaxes, fake news, conspiracy theories, and counterfeits. In the future, I will study these different types of forgeries, revisit the data model to include their representation and refine the SEBI ontology-related terms, broadening the range of about collected evidence (hypernyms and hyponyms).

My main goal is to develop an open-collaborative online portal that leverages RDF and EWA to allow scholars to record and share their inquiries regarding forged artworks. The portal aims to build a comprehensive catalogue of known or suspected forged documents and scholarly inquiries addressing them by fostering collaboration among researchers.

Acknowledgements

This dissertation has already covered considerable ground, so I will keep these acknowledgements concise. Still, many individuals' contributions deserve mention, and I extend my sincere gratitude to all of them.

To the researchers at the Digital Humanities Research Center (/DH.arc) at the Department of Classical Philology and Italian Studies (FICLIT) at the University of Bologna. I am deeply grateful to have spent my PhD years with you, benefiting from your knowledge, humanity, and kindness.

Above all, to my supervisor, Professor Francesca Tomasi, for her constant support throughout this journey. My gratitude extends far beyond academic supervision or achievements. Thank you for teaching me a profession and a passion to share with others.

To my co-supervisor, Professor Fabio Vitali, for his *unquestionable* passion for Conjectures, his creative mind, and his dedication to constantly rethinking and refining the project.

To the Institut für Digitale Geisteswissenschaften (ZIM) members in Graz for hosting me during my research period. In particular, I thank Professors Georg Vogeler and Daniel Luger for their thoughtful support in developing the case study for this dissertation.

To Professor Siegfried Haider and the Oberösterreichisches Landesarchiv for providing the source materials of this dissertation's case study.

To the thesis students, Alessio Di Pasquale, Gerald Manzano, and Eduart Uzeir, for contributing to the Conjectures project and sharing a bit of this

journey with me.
To the IT staff of the FICLIT Department, particularly Tommaso Vitale, for their always timely and kind assistance with the application deployment and server maintenance throughout this project.
To the reviewers of this dissertation and the related publications, for their rigorous feedback and precise observations that have substantially improved this research.
To all of my co-workers for standing by through the ups and downs of this three-year roller coaster. Your support has made all the difference.
To my family, especially my parents, for accompanying me every step of the way, even through concepts and choices that may have seemed incomprehensible. I am particularly grateful that they understood and came to appreciate digital humanities — I guess this is the definition of unconditional love.
For every hiking trail and mountain summit I shared with my long-lasting friends — you provided all the fresh air I needed, even without realising it.

I am grateful for how this journey has enriched me, from exploring new places to meeting researchers worldwide and learning from their varied perspectives and knowledge.
If I have inadvertently left anyone unmentioned, please attribute this to weariness rather than a lack of gratitude. My deepest thankfulness goes to everyone who supported me — or those dear to me — throughout this journey.

Bibliography

[1] Abián, D., Guerra, F., and Trillo-Lado, Raquel. "Wikidata and DBpedia: A Comparative Study". In: *Semantic Keyword-Based Search on Structured Data Sources*. Ed. by Szymański, Julian and Velegrakis, Yannis. Springer International Publishing, 2018, pp. 142–154. ISBN: 978-331974496-4. DOI: 10.1007/978-3-319-74497-1_14.

[2] Aljalbout, Sahar, Falquet, Gilles, and Buchs, Didier. "Handling Wikidata Qualifiers in Reasoning". In: *arXiv Preprint* 2304.03375 (2023). URL: https://arxiv.org/abs/2304.03375.

[3] Andrews, Tara. "Modelling Historical Data in the RELEVEN Project (IT)". In: *Programming and Data Infrastructure in Digital Humanities* (2023), p. 16.

[4] Andrews, Tara. "The Structured Assertion Record (STAR) Model for Event-based Representation of Historical Information". In: *GrapHNR 2023*. Mainz, Germany, 2023. URL: https://graphentechnologien.hypotheses.org/files/2023/05/GrapHNR-2023-32-Andrews-STAR.pdf.

[5] Arndt, Dörthe and Van Woensel, William. "Towards supporting multiple semantics of named graphs using n3 rules". In: *13th RuleML+ RR 2019 Doctoral Consortium and Rule Challenge*. Vol. 2438. CEUR. 2019. URL: https://ceur-ws.org/Vol-2438/paper6.pdf.

[6] Auer, Sören, Bizer, Christian, Kobilarov, Georgi, Lehmann, Jens, Cyganiak, Richard, and Ives, Zachary. "Dbpedia: A nucleus for a web

of open data". In: *The Semantic Web: 6th International Semantic Web Conference, 2nd Asian Semantic Web Conference, ISWC 2007+ ASWC 2007, Busan, Korea, November 11-15, 2007. Proceedings*. Springer. 2007, pp. 722–735. DOI: `10.1007/978-3-540-76298-0\_52`.

[7] Balaraman, Vevake, Razniewski, Simon, and Nutt, Werner. "Recoin: Relative Completeness in Wikidata". In: *Companion Proceedings of the The Web Conference 2018*. WWW '18. Lyon, France: International World Wide Web Conferences Steering Committee, 2018, 1787–1792. ISBN: 9781450356404. DOI: `10.1145/3184558.3191641`.

[8] Barabucci, Gioele, Tomasi, Francesca, and Vitali, Fabio. "Supporting Complexity and Conjectures in Cultural Heritage Descriptions". In: *Proceedings of the International Conference Collect and Connect: Archives and Collections in a Digital Age*. Leiden, the Netherlands: CEUR Workshop, Nov. 2021, pp. 104–115. URL: `http://ceur-ws.org/Vol-2810/paper9.pdf`.

[9] Barone, Nicola. "Intorno alla falsificazione dei documenti ed alla critica di essi. Memoria letta all'Accademia Pontaniana nella tornata del 21 gennaio 1912". In: *Atti Dell'Accademia Pontaniana* 42 (1912). URL: `http://www.rmoa.unina.it/4359/`.

[10] Benslimane, Djamal, Sheng, Quan Z., Barhamgi, Mahmoud, and Prade, Henri. "The Uncertain Web: Concepts, Challenges, and Current Solutions". In: *ACM Trans. Internet Technol.* 16.1 (Dec. 2016). ISSN: 1533-5399. DOI: `10.1145/2847252`.

[11] Berners-Lee, Tim, Connolly, Dan, Kagal, Lalana, Scharf, Yosi, and Hendler, Jim. "N3Logic: A logical framework for the World Wide Web". In: *Theory Pract. Log. Program.* 8.3 (2008), pp. 249–269. DOI: `10.1017/S1471068407003213`.

[12] Blau, Adrian. "Uncertainty and the history of ideas". In: *History and Theory* 50.3 (2011), pp. 358–372. ISSN: 00182656, 14682303. DOI: `10.`

1111/j.1468-2303.2011.00590.x. URL: http://www.jstor.org/stable/41300100.

[13] Borgman, Christine L. and Wofford, Morgan F. "From Data Processes to Data Products: Knowledge Infrastructures in Astronomy". In: *Harvard Data Science Review* 3.3 (July 2021). DOI: 10.1162/99608f92.4e792052.

[14] Bruseker, George, Guillem, Anaïs, and Carboni, Nicola. "Semantically Documenting Virtual Reconstruction: Building a Path to Knowledge Provenance". In: *ISPRS Annals of the Photogrammetry, Remote Sensing and Spatial Information Sciences* II-5/W3 (2015), pp. 33–40. DOI: 10.5194/isprsannals-II-5-W3-33-2015.

[15] Carroll, Jeremy J., Bizer, Christian, Hayes, Pat, and Stickler, Patrick. "Named graphs". In: *Journal of Web Semantics* 3.4 (2005). World Wide Web Conference 2005——Semantic Web Track, pp. 247–267. ISSN: 1570-8268. DOI: 10.1016/j.websem.2005.09.001.

[16] Chapman, David. "Information is provisional". In: *Kybernetes* 43.6 (2014). Ed. by Ali, Mustafa, pp. 895–910. ISSN: 0368-492X. DOI: 10.1108/K-08-2013-0178.

[17] Checkland, Peter and Holwell, Sue. "Data, capta, information and knowledge". English. In: *Introducing Information Management: the business approach*. Elsevier, 2006, pp. 47–55. ISBN: 0-7506-6668-4. DOI: 10.4324/9780080458397-10.

[18] Chiesa, Angela Ottino della. *The Complete Paintings of Leonardo Da Vinci*. Classics of World Art Series. Penguin, 1985. ISBN: 9780140086492. URL: https://books.google.it/books?id=nr3vAAAACAAJ.

[19] Consortium, World Wide Web et al. "RDF 1.1 concepts and abstract syntax". In: *W3C Proposed Recommendation* (2014). URL: https://www.w3.org/TR/rdf11-concepts/.

[20] Costa, Paulo C. G., Laskey, Kathryn B., Blasch, Erik, and Jousselme, Anne-Laure. "Towards unbiased evaluation of uncertainty reasoning: The URREF ontology". In: *15th International Conference on Information Fusion*. 2012, pp. 2301–2308. ISBN: 978-1-4673-0417-7. URL: https://ieeexplore.ieee.org/document/6290584.

[21] Daquino, Marilena. *Mining Authoritativeness in Art Historical Photo Archives: Semantic Web Applications for Connoisseurship*. Vol. 40. IOS Press, 2019. ISBN: 978-1-64368-010-1.

[22] Daquino, Marilena. "A computational analysis of art historical linked data for assessing authoritativeness of attributions". In: *Journal of the Association for Information Science and Technology* 71.7 (2020), pp. 757–769. DOI: 10.1002/asi.24301.

[23] Daquino, Marilena, Giovannetti, Francesca, and Tomasi, Francesca. "Linked Data per le edizioni scientifiche digitali. Il workflow di pubblicazione dell'edizione semantica del quaderno di appunti di Paolo Bufalini". In: *Umanistica Digitale* 7 (2019), pp. 49–75. DOI: 10.6092/issn.2532-8816/9091.

[24] Daquino, Marilena, Pasqual, Valentina, and Tomasi, Francesca. "Knowledge Representation of Digital Hermeneutics of Archival and Literary Sources". In: vol. 11. 3. 2020, pp. 59–76. URL: 10.4403/jlis.it-12642.

[25] Daquino, Marilena, Pasqual, Valentina, Tomasi, Francesca, and Vitali, Fabio. "Expressing Without Asserting in the Arts". In: *Proceedings of the Italian Research Conference on Digital Libraries. Padova, Italy*. CEUR-WS.org, 2022. URL: https://ceur-ws.org/Vol-3160/short14.pdf.

[26] Daquino, Marilena and Tomasi, Francesca. "Historical Context Ontology (HiCO): A Conceptual Model for Describing Context Information of Cultural Heritage Objects". In: *Metadata and Semantics Re-*

search. MTSR 2015. Ed. by Garoufallou, E., Hartley, R., and Gaitanou, P. Vol. 544. Communications in Computer and Information Science. Springer, 2015, pp. 359–368. URL: 10.1007/978-3-319-24129-6_37.

[27] Darch, Peter T. and Sands, Ashley E. "Uncertainty about the Long-Term: Digital Libraries, Astronomy Data, and Open Source Software". In: *2017 ACM/IEEE Joint Conference on Digital Libraries (JCDL)*. 2017, pp. 1–4. DOI: 10.1109/JCDL.2017.7991584.

[28] Delmas-Glass, Emmanuelle and Sanderson, Robert. "Fostering a community of PHAROS scholars through the adoption of open standards". In: *Art Libraries Journal* 45.1 (2020), 19–23. DOI: 10.1017/alj.2019.32.

[29] Demartini, Gianluca, Enchev, Iliya, Wylot, Marcin, Gapany, Joël, and Cudré-Mauroux, Philippe. "BowlognaBench—Benchmarking RDF Analytics". In: *Data-Driven Process Discovery and Analysis*. Ed. by Aberer, Karl, Damiani, Ernesto, and Dillon, Tharam. Berlin, Heidelberg: Springer Berlin Heidelberg, 2012, pp. 82–102. ISBN: 978-3-642-34044-4. DOI: 10.1007/978-3-642-34044-4_5.

[30] Di Pasquale, Alessio, Pasqual, Valentina, Tomasi, Francesca, and Vitali, Fabio. "On assessing weaker logical status claims in Wikidata cultural heritage records". In: *Semantic Web* 15 (2024). 6, pp. 2395–2417. ISSN: 2210-4968. DOI: 10.3233/SW-243686.

[31] Diefenbach, Dennis, Wilde, Max De, and Alipio, Samantha. "Wikibase as an Infrastructure for Knowledge Graphs: The EU Knowledge Graph". In: *The Semantic Web – ISWC 2021*. Ed. by Hotho, Andreas, Blomqvist, Eva, Dietze, Stefan, Fokoue, Achille, Ding, Ying, Barnaghi, Payam, Haller, Armin, Dragoni, Mauro, and Alani, Harith. Cham: Springer International Publishing, 2021, pp. 631–647. ISBN: 978-3-030-88361-4. DOI: 10.1007/978-3-030-88361-4_37.

[32] Dividino, Renata, Sizov, Sergej, Staab, Steffen, and Schueler, Bernhard. "Querying for provenance, trust, uncertainty and other meta knowledge in RDF". In: *Journal of Web Semantics* 7.3 (2009), pp. 204–219. ISSN: 1570-8268. DOI: 10.1016/j.websem.2009.07.004.

[33] Djebri, Ahmed El Amine, Tettamanzi, Andrea GB, and Gandon, Fabien. "Publishing uncertainty on the semantic web: blurring the LOD bubbles". In: *International Conference on Conceptual Structures*. Springer. 2019, pp. 42–56. DOI: 10.1007/978-3-030-23182-8_4.

[34] Doerr, Martin, Gradmann, Stefan, Hennicke, Steffen, Isaac, Antoine, Meghini, Carlo, and Sompel, Herbert Van de. "The europeana data model (EDM)". In: *World Library and Information Congress: 76th IFLA general conference and assembly*. Vol. 10. 2010, p. 15. URL: https://cdn.ifla.org/past-wlic/2010/149-doerr-en.pdf.

[35] Doerr, Martin, Kritsotaki, Athina, Rousakis, Yannis, Hiebel, Gerald, and Theodoridou, Maria. *CRMsci: The scientific observation model*. Tech. rep. FORTH, Tech. Rep., 2014.[Online], 2014. URL: https://cidoc-crm.org/crmsci/sites/default/files/CRMsci1.2.pdf.

[36] Doerr, Martin, Ore, Christian-Emil, and Fafalios, Pavlos. *Definition of the CRMinf*. Tech. rep. 2023. URL: https://www.cidoc-crm.org/crminf/sites/default/files/CRMinf%20v1.0%28site%29.pdf.

[37] Doerr, Martin, Ore, Christian-Emil, and Stead, Stephen. "The CIDOC conceptual reference model-A new standard for knowledge sharing". In: *26th international conference on conceptual modeling (ER 2007)*. Vol. 83. Auckland, New Zealand, Nov. 2007, pp. 51–56. DOI: 10.13140/2.1.1420.6400.

[38] Dong, Li, Wei, Furu, Tan, Chuanqi, Tang, Duyu, Zhou, Ming, and Xu, Ke. "Adaptive Recursive Neural Network for Target-dependent Twitter Sentiment Classification". In: *Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics (Volume 2: Short*

Papers). Ed. by Toutanova, Kristina and Wu, Hua. Baltimore, Maryland: Association for Computational Linguistics, 2014, pp. 49–54. URL: https://aclanthology.org/P14-2009.

[39] Drucker, Johanna. "Humanities Approaches to Graphical Display". In: *Digital Humanities Quarterly* 5.1 (Mar. 2011). ISSN: 1938-4122. URL: https://www.digitalhumanities.org/dhq/vol/5/1/000091/000091.html.

[40] Erxleben, Fredo, Günther, Michael, Krötzsch, Markus, Mendez, Julian, and Vrandečić, Denny. "Introducing Wikidata to the Linked Data Web". In: *The Semantic Web - ISWC 2014*. Ed. by Mika, Peter, Tudorache, Tania, Bernstein, Abraham, Welty, Chris, Knoblock, Craig, Vrandečić, Denny, Groth, Paul, Noy, Natasha, Janowicz, Krzysztof, and Goble, Carole. Cham: Springer International Publishing, 2014, pp. 50–65. ISBN: 978-3-319-11964-9. DOI: 10.1007/978-3-319-11964-9_4.

[41] Färber, Michael, Bartscherer, Frederic, Menne, Carsten, Rettinger, Achim, Zaveri, Amrapali, Kontokostas, Dimitris, Hellmann, Sebastian, and Umbrich, Jürgen. "Linked Data Quality of DBpedia, Freebase, OpenCyc, Wikidata, and YAGO". In: *Semantic Web* 9.1 (Jan. 2018), 77–129. ISSN: 1570-0844. DOI: 10.3233/SW-170275.

[42] Fink, Eleanor E. "American Art Collaborative (AAC) Linked Open Data (LOD) Initiative, Overview and Recommendations for Good Practices". In: (2018). URL: https://repository.si.edu/bitstream/handle/10088/106410/OverviewandRecommendationsAccessible.pdf.

[43] Franzini, Greta, Zampedri, Federica, Passarotti, Marco, Mambrini, Francesco, and Moretti, Giovanni. "Græcissare: Ancient Greek Loanwords in the LiLa Knowledge Base of Linguistic Resources for Latin." In: *CLiC-it*. Ed. by J., Monti, F., Dell'Orletta, and F., Tamburini.

Vol. 2769. CEUR-WS, 2020. URL: http://ceur-ws.org/Vol-2769/paper_06.pdf.

[44] Frey, Johannes, Müller, Kay, Hellmann, Sebastian, Rahm, Erhard, and Vidal, Maria-Esther. "Evaluation of Metadata Representations in RDF Stores". In: *Semant. Web* 10.2 (Jan. 2019), 205–229. ISSN: 1570-0844. DOI: 10.3233/SW-180307.

[45] Gadamer, Hans G. *Truth and Method*. London: A&C Black, 2013. URL: https://books.google.it/books/?id=SkdMAQAAQBAJ.

[46] Galárraga, Luis, Razniewski, Simon, Amarilli, Antoine, and Suchanek, Fabian M. "Predicting Completeness in Knowledge Bases". In: *Proceedings of the Tenth ACM International Conference on Web Search and Data Mining*. WSDM '17. Cambridge, United Kingdom: Association for Computing Machinery, 2017, 375–383. ISBN: 9781450346757. DOI: 10.1145/3018661.3018739.

[47] Gao, Z., Feng, A., Song, X., and Wu, X. "Target-Dependent Sentiment Classification With BERT". In: *IEEE Access* 7 (2019), pp. 154290–154299. URL: 10.1109/ACCESS.2019.2946594.

[48] Garson, James. "Modal Logic". In: *The Stanford Encyclopedia of Philosophy*. Ed. by Zalta, Edward N. and Nodelman, Uri. Spring 2023. Metaphysics Research Lab, Stanford University, 2023. URL: https://plato.stanford.edu/entries/logic-modal/.

[49] Glaser, Barney G and Strauss, Anselm L. *The discovery of grounded theory: Strategies for qualitative research*. Routledge, 2017. ISBN: 1351522167.

[50] Grego, Francesca. *Galeotto fu il restauro: riemerge dopo 300 anni il cupido di Vermeer*. Aug. 2021. URL: https://www.arte.it/notizie/mondo/galeotto-fu-il-restauro-riemerge-dopo-300-anni-il-cupido-di-vermeer-18643.

[51] Groth, Paul, Gibson, Andrew, and Velterop, Jan. “The anatomy of a nanopublication”. In: *Information Services & Use* 30.1-2 (2010), pp. 51–56. DOI: 10.3233/ISU-2010-0613.

[52] Haider, Sigfried. *Verzeichnis der den oberösterreichischen Raum betreffenden gefälschten, manipulierten oder verdächtigten mittelalterlichen Urkunden.* Oberösterreichisches Landesarchiv, 2022. ISBN: 978-3-902801-45-6.

[53] Hamborg, F., Donnay, K., and Gipp, B. “Towards Target-Dependent Sentiment Classification in News Articles”. In: *Diversity, Divergence, Dialogue. iConference 2021.* Ed. by Toeppe, K., Yan, H., and Chu, S.K.W. Vol. 12646. Lecture Notes in Computer Science. Springer, 2021, pp. 157–169. URL: 10.1007/978-3-030-71305-8_12.

[54] Hartig, Olaf. “Foundations of RDF* and SPARQL*:(An alternative approach to statement-level metadata in RDF)”. In: *Proceedings of the 11th Alberto Mendelzon International Workshop on Foundations of Data Management and the Web*. Ed. by Reutter, Juan L. and Srivastava, Divesh. Vol. 1912. CEUR Workshop Proceedings. Juan Reutter, Divesh Srivastava. Montevideo, Uruguay: CEUR-WS.org, June 2017. URL: http://ceur-ws.org/Vol-1912/paper12.pdf.

[55] Hawke, Sandro, Herman, Ivan, Parsia, Bijan, Polleres, Axel, and Seaborne, Andy. *SPARQL 1.1 Entailment Regimes.* W3C Recommendation. W3C, Mar. 2013. URL: https://www.w3.org/TR/sparql11-entailment/.

[56] Hayes, Patrick. *RDF semantics, W3C recommendation.* Tech. rep. World Wide Web Consortium, 2004. URL: http://www.w3.org/TR/rdf-mt/.

[57] *Help:Ranking - Wikidata.* URL: https://www.wikidata.org/wiki/Help:Ranking.

[58] Hernández, Daniel, Gutierrez, Claudio, and Hogan, Aidan. "Certain Answers for SPARQL with Blank Nodes". In: *The Semantic Web – ISWC 2018: 17th International Semantic Web Conference, Proceedings, Part I*. Monterey, CA, USA: Springer-Verlag, 2018, 337–353. ISBN: 978-3-030-00670-9. DOI: 10.1007/978-3-030-00671-6_20.

[59] Hernández, Daniel, Hogan, Aidan, and Krötzsch, Markus. "Reifying RDF: What works well with wikidata?" In: *Proceedings of the 11th International Workshop on Scalable Semantic Web Knowledge Base Systems*. Ed. by Liebig, Thorsten and Fokoue, Achille. Vol. 1457. Bethlehem, Pennsylvania, USA: CEUR Workshop, June 2015, pp. 32–47. URL: https://ceur-ws.org/Vol-1457/SSWS2015_paper3.pdf.

[60] Hogan, Aidan, Umbrich, Jürgen, Harth, A., Cyganiak, Richard, Polleres, Axel, and Decker, S. "An empirical survey of Linked Data conformance". In: *J. Web Semant.* 14 (2012), pp. 14–44. DOI: http://dx.doi.org/10.2139/ssrn.3198962.

[61] Isaac, Antoine et al. *Europeana data model primer*. Tech. rep. Europeana, 2013. URL: https://pro.europeana.eu/files/Europeana_Professional/Share_your_data/Technical_requirements/EDM_Documentation/EDM_Primer_130714.pdf.

[62] Laskey, Kenneth J and Laskey, Kathryn B. "Uncertainty Reasoning for the World Wide Web: Report on the URW3-XG Incubator Group." In: *URSW*. Citeseer. 2008. URL: https://www.w3.org/2005/Incubator/urw3/.

[63] Lebo, T. et al. *PROV-O: The PROV Ontology*. W3C Recommendation. World Wide Web Consortium, 2013. URL: http://www.w3.org/TR/2013/REC-prov-o-20130430/.

[64] Malik, Sonika and Jain, Sarika. "A Bayesian network approach to handle uncertainty in Web Ontology Language". In: *International Journal of*

Reasoning-based Intelligent Systems 13.4 (2021), pp. 243–250. DOI: 10.1504/IJRIS.2021.118661.

[65] Manzano, Gerald. "Estensione per i parser RDF e SPARQL per esprimere senza asserire con le Congetture". PhD thesis. URL: http://amslaurea.unibo.it/30616/.

[66] Mariani, Fabio. "Introducing VISU: Vagueness, Incompleteness, Subjectivity, and Uncertainty in Art Provenance Data". English. In: *Computational Methods in the Humanities 2022*. Ed. by Rochat, Yannick, Metrailler, Coline, and Piotrowski, Michael. Vol. 3602. CEUR Workshop Proceedings. Germany: Sun Site Central Europe (RWTH Aachen University), 2023, pp. 63–84. URL: https://wp.unil.ch/llist/en/event/comhum2022/.

[67] Martin, Craig. "Conjecture, probabilism, and provisional knowledge in Renaissance meteorology". In: *Early Science and Medicine* 14.1-3 (2009), pp. 265–289. DOI: 10.1163/157338209X425588.

[68] McDermott, Drew V. and Dou, Dejing. "Representing Disjunction and Quantifiers in RDF". In: *Proceedings of the First International Semantic Web Conference on The Semantic Web*. ISWC '02. Berlin, Heidelberg: Springer-Verlag, 2002, 250–263. ISBN: 3540437606. DOI: 10.1007/3-540-48005-6_20.

[69] Meyer, L.-P., Stadler, C., et al. "LLM-assisted Knowledge Graph Engineering: Experiments with ChatGPT". In: *First Working Conference on Artificial Intelligence Development for a Resilient and Sustainable Tomorrow (AIDRST 2023)*. Ed. by Zinke-Wehlmann, C. and Friedrich, J. Informatik aktuell. Springer Vieweg, 2024, pp. 157–169. URL: 10.1007/978-3-658-43705-3_8.

[70] Mihindukulasooriya, N., Tiwari, S., Fernández, E.C., and Lata, K. "Text2KGBench: A Benchmark for Ontology-Driven Knowledge Graph Generation from Text". In: *The Semantic Web – ISWC 2023*. Ed. by

Payne, T.R., Huynh, D.D., Kim, J., Haddad, H., Afzal, Z., Pan, J.Z., Chapman, M., Gandon, F.L., Krishna, R., Dumontier, M., and Zhao, J. Vol. 14266. Lecture Notes in Computer Science. Springer, 2023, pp. 239–254. URL: 10.1007/978-3-031-47243-5_14.

[71] Möller, Cedric, Lehmann, Jens, and Usbeck, Ricardo. "Survey on English Entity Linking on Wikidata: Datasets and Approaches". In: *Semantic Web* 13 (6 2022). DOI: 10.3233/SW-212865.

[72] Nguyen, Vinh, Bodenreider, Olivier, and Sheth, Amit. "Don't like RDF Reification? Making Statements about Statements Using Singleton Property". In: WWW '14. Seoul, Korea: Association for Computing Machinery, 2014, 759–770. ISBN: 9781450327442. DOI: 10.1145/2566486.2567973.

[73] Nguyen, Vinh, Bodenreider, Olivier, Thirunarayan, Krishnaprasad, Fu, Gang, Bolton, Evan, Queralt-Rosinach, Núria, Furlong, Laura I, Dumontier, Michel, and Sheth, Amit. "On Reasoning with RDF Statements about Statements using Singleton Property Triples". In: *ArXiv* abs/1509.04513 (Sept. 2015). URL: https://arxiv.org/abs/1509.04513.

[74] Niccolucci, Franco and Hermon, Sorin. "Expressing reliability with CIDOC CRM". In: *International Journal on Digital Libraries* 18.4 (2017), pp. 281–287. DOI: 10.1007/s00799-016-0195-1.

[75] Noy, Natasha, Rector, Alan, Hayes, Pat, and Welty, Chris. "Defining n-ary relations on the semantic web". In: *W3C working group note* 12.4 (2006). URL: https://www.w3.org/TR/2006/NOTE-swbp-n-aryRelations-20060412/.

[76] Nussbaumer, Philipp and Haslhofer, Bernhard. *Putting the CIDOC CRM into Practice - Experiences and Challenges*. Technical Report. Sept. 2007. URL: http://eprints.cs.univie.ac.at/404/.

[77] Oldman, Dominic and Labs, CRM. "The CIDOC conceptual reference model (CIDOC-CRM): primer". In: *CRM Labs* 5 (2014). URL: https://www.cidoc-crm.org/sites/default/files/CRMPrimer_v1.1_1.pdf.

[78] Orlandi, Fabrizio, Graux, Damien, and O'Sullivan, Declan. "Benchmarking RDF Metadata Representations: Reification, Singleton Property and RDF". In: *2021 IEEE 15th International Conference on Semantic Computing (ICSC)*. IEEE. 2021, pp. 233–240. DOI: 10.1109/ICSC50631.2021.00049.

[79] Pasqual, Valentina, Daquino, Marilena, Tomasi, Francesca, et al. "MIMA: a data model to represent multi-disciplinary analysis on manuscripts. Use case on Pellegrino Prisciani's Historiae Ferrariae". In: *AIUCD 2021-DH per la società: e-guaglianza, partecipazione, diritti e valori nell'era digitale. Raccolta degli abstract estesi della 10° conferenza nazionale*. 2021, pp. 560–566. DOI: 10.6092/unibo/amsacta/6712.

[80] Pasqual, Valentina, Manzano, Gerald, Uzeir, Eduart, Tomasi, Francesca, Vitali, Fabio, et al. "Evaluation of Expressing without asserting approaches in RDF. The case of Conjectures." In: *Proceedings of the Conference on Information and Research science Connecting to Digital and Library science. Brixen, Italy*. CEUR-WS.org, 2024, pp. 60–73. URL: https://ceur-ws.org/Vol-3643/paper7.pdf.

[81] Pasqual, Valentina and Tomasi, Francesca. "Linked open data per la valorizzazione di collezioni culturali: il dataset mythLOD". In: *AIB studi* 62.1 (May 2022), 149–168. DOI: 10.2426/aibstudi-13301.

[82] Pasquale, Alessio Di, Vitali, Fabio, and Pasqual, Valentina. *Wikidata selection of Cultural Heritage, Stars, Galaxies and Random entities and claims*. Feb. 2023. DOI: 10.5281/zenodo.7624783.

[83] Patel-Schneider, Peter F. "Contextualization via Qualifiers". In: *Joint Proceedings of the International Workshops on Contextualized Knowledge Graphs, and Semantic Statistics co-located with 17th Interna-*

tional Semantic Web Conference (ISWC 2018). Ed. by Capadisli, Sarven, Cotton, Franck, Giménez-García, José M., Haller, Armin, Kalampokis, Evangelos, Nguyen, Vinh, Sheth, Amit P., and Troncy, Raphaël. Vol. 2317. CEUR Workshop Proceedings. CEUR-WS.org, 2018. URL: `https://ceur-ws.org/Vol-2317/article-01.pdf`.

[84] Peroni, Silvio and Shotton, David. "FaBiO and CiTO: Ontologies for describing bibliographic resources and citations". In: *Journal of Web Semantics* 17 (2012), pp. 33–43. ISSN: 1570-8268. DOI: `10.1016/j.websem.2012.08.001`.

[85] Petras, Vivien, Hill, Timothy, Stiller, Juliane, and Gäde, Maria. "Europeana–a Search Engine for Digitised Cultural Heritage Material". In: *Datenbank-Spektrum* 17 (2017), pp. 41–46. DOI: `10.1007/s13222-016-0238-1`.

[86] Piotrowski, Michael. "Accepting and modeling uncertainty". In: *Zeitschrift für digitale Geisteswissenschaften* (2019). Ed. by Andreas Kuczera Thorsten Wübbena, Thomas Kollatz von. DOI: `10.17175/sb004_006a`.

[87] Piotrowski, Michael. "Uncertainty as Unavoidable Good". In: *Universität Bielefeld, Center for Uncertainty Studies (CeUS)* 5 (2023), p. 10. DOI: `10.4119/unibi/2983506`.

[88] Piotrowski, Michael and Neuwirth, Markus. "Prospects for computational hermeneutics". In: *Atti del IX Convegno Annuale AIUCD. La svolta inevitabile: sfide e prospettive per l'Informatica Umanistica*. Ed. by Marras, Cristina, Passarotti, Marco, Franzini, Greta, and Litta, Eleonora. Quaderni di Umanistica Digitale. Milano: Università Cattolica del Sacro Cuore, Jan. 2020. ISBN: 978-88-942535-4-2. DOI: `10.6092/unibo/amsacta/6316`.

[89] Piscopo, Alessandro and Simperl, Elena. "What We Talk about When We Talk about Wikidata Quality: A Literature Survey". In: *Proceedings*

of the 15th International Symposium on Open Collaboration. OpenSym '19. Skövde, Sweden: Association for Computing Machinery, 2019. ISBN: 9781450363198. DOI: 10.1145/3306446.3340822.

[90] Ponza, Marco, Ferragina, Paolo, and Chakrabarti, Soumen. "A Two-Stage Framework for Computing Entity Relatedness in Wikipedia". In: *Proceedings of the 2017 ACM on Conference on Information and Knowledge Management*. CIKM '17. Singapore, Singapore: Association for Computing Machinery, 2017, 1867–1876. ISBN: 9781450349185. DOI: 10.1145/3132847.3132890.

[91] Rebele, Thomas, Suchanek, Fabian, Hoffart, Johannes, Biega, Joanna, Kuzey, Erdal, and Weikum, Gerhard. "YAGO: A multilingual knowledge base from wikipedia, wordnet, and geonames". In: *The Semantic Web–ISWC 2016: 15th International Semantic Web Conference, Kobe, Japan, October 17–21, 2016, Proceedings, Part II 15*. Springer. 2016, pp. 177–185. DOI: 10.1007/978-3-319-46547-0_19.

[92] Rolfini, Alessio. "Semantics of Conjectures". In: *CoRR* abs/2110.08920 (2021). arXiv: 2110.08920.

[93] Rother, Lynn, Koss, Max, and Mariani, Fabio. "Taking Care of History: Toward a Politics of Provenance Linked Open Data in Museums". In: *Perspectives on Data*. Ed. by Fry, Emily Lew and Canning, Erin. Art Institute of Chicago, 2022, chapter 6. DOI: 10.53269/9780865593152/06.

[94] Sahoo, Satya S, Bodenreider, Olivier, Hitzler, Pascal, Sheth, Amit, and Thirunarayan, Krishnaprasad. "Provenance Context Entity (PaCE): Scalable provenance tracking for scientific RDF data". In: *International Conference on Scientific and Statistical Database Management*. Springer. 2010, pp. 461–470. DOI: 10.1007/978-3-642-13818-8_32.

[95] Sangeeta Sen Mariana Curado Malta, Biswanath Dutta and Dutta, Animesh. "State-of-the-Art Approaches for Meta-Knowledge Assertion in

the Web of Data". In: *IETE Technical Review* 38.6 (2021), pp. 672–709. DOI: 10.1080/02564602.2020.1819891.

[96] Sartini, Bruno, Baroncini, Sofia, Erp, Marieke van, Tomasi, Francesca, and Gangemi, Aldo. "ICON: An Ontology for Comprehensive Artistic Interpretations". In: *Journal of Computational Cultural Heritage* 16.3 (2023), pp. 59–76. DOI: 10.1145/3594724.

[97] Schneider, Jodi, Groza, Tudor, and Passant, Alexandre. "A review of argumentation for the social semantic web". In: *Semantic Web* 4.2 (2013), pp. 159–218. DOI: 10.3233/SW-2012-0073.

[98] Schueler, Bernhard, Sizov, Sergej, Staab, Steffen, and Tran, Duc Thanh. "Querying for meta knowledge". In: *Proceedings of the 17th International Conference on World Wide Web*. WWW '08. Beijing, China: Association for Computing Machinery, 2008, 625–634. ISBN: 9781605580852. DOI: 10.1145/1367497.1367582.

[99] Shenoy, Kartik, Ilievski, Filip, Garijo, Daniel, Schwabe, Daniel, and Szekely, Pedro. "A study of the quality of Wikidata". In: *Journal of Web Semantics* 72 (2022), p. 100679. ISSN: 1570-8268. DOI: 10.1016/j.websem.2021.100679.

[100] Sikos, Leslie F and Philp, Dean. "Provenance-aware knowledge representation: A survey of data models and contextualized knowledge graphs". In: *Data Science and Engineering* 5.3 (2020), pp. 293–316. ISSN: 2364-1541. DOI: 10.1007/s41019-020-00118-0.

[101] Solomon, Tessa. *Restoration of Vermeer Painting in Germany Reveals Hidden Image of Cupid*. Aug. 2021. URL: https://www.artnews.com/art-news/news/vermeer-girl-reading-a-letter-cupid-restoration-1234602227/.

[102] Stinson, Alexander, Fauconnier, Sandra, and Wyatt, Liam. "Stepping Beyond Libraries: The Changing Orientation in Global GLAM-Wiki". In: *JLIS.it* 9.3 (Sept. 2018), 16–34. DOI: 10.4403/jlis.it-12480.

[103] Teller, Paul. "Provisional Knowledge". In: *Constituting Objectivity: Transcendental Perspectives on Modern Physics*. Ed. by Bitbol, Michel, Kerszberg, Pierre, and Petitot, Jean. Dordrecht: Springer Netherlands, 2009, pp. 503–514. ISBN: 978-1-4020-9510-8. DOI: 10.1007/978-1-4020-9510-8_30.

[104] Theoharis, Yannis, Christophides, Vassilis, and Karvounarakis, Grigoris. "Benchmarking Database Representations of RDF/S Stores". In: *The Semantic Web – ISWC 2005*. Ed. by Gil, Yolanda, Motta, Enrico, Benjamins, V. Richard, and Musen, Mark A. Berlin, Heidelberg: Springer Berlin Heidelberg, 2005, pp. 685–701. ISBN: 978-3-540-32082-1. DOI: 10.1007/11574620_49.

[105] Udrea, Octavian, Recupero, Diego Reforgiato, and Subrahmanian, V. S. "Annotated RDF". In: *ACM Trans. Comput. Logic* 11.2 (Jan. 2010). ISSN: 1529-3785. DOI: 10.1145/1656242.1656245.

[106] Vitali, Fabio and Pasqual, Valentina. "The Representation of Historical Uncertainties as the Outcome of Competing and Incompatible Certainties". In: *Edizioni Ca' Foscari* ([ACCEPTED]).

[107] Vrandečić, Denny and Krötzsch, Markus. "Wikidata: a free collaborative knowledgebase". In: *Commun. ACM* 57.10 (Sept. 2014), 78–85. ISSN: 0001-0782. DOI: 10.1145/2629489.

[108] Watkins, E. Rowland and Nicole, Denis A. "Named Graphs as a Mechanism for Reasoning About Provenance". In: *Frontiers of WWW Research and Development - APWeb 2006*. Ed. by Zhou, Xiaofang, Li, Jianzhong, Shen, Heng Tao, Kitsuregawa, Masaru, and Zhang, Yanchun. Berlin, Heidelberg: Springer Berlin Heidelberg, 2006, pp. 943–948. ISBN: 978-3-540-32437-9. DOI: 10.1007/11610113_99.

[109] Zhitomirsky-Geffet, Maayan and Minster, Sara. "Cultural information bubbles: A new approach for automatic ethical evaluation of digital

artwork collections based on Wikidata". In: *Digital Scholarship in the Humanities* (Dec. 2022). ISSN: 2055-7671. DOI: 10.1093/llc/fqac076.

[110] Zimmermann, Antoine. *RDF 1.1: on semantics of RDF datasets*. Tech. rep. 2014. URL: http://www.w3.org/TR/rdf11-datasets/.