

# Reconstructing Historical Events with CIDOC CRM

Yu Lee An<sup>1</sup> 

<sup>1</sup> School of Information and Library Science, University of North Carolina at Chapel Hill

## Abstract

Historical sources frequently contain uncertain, approximate, and ambiguous temporal information, posing a significant challenge for their representation in computational systems. This paper addresses this challenge by demonstrating how semantic web technologies and an event-centric ontology can systematically represent the pervasive temporal uncertainty inherent in primary sources. Using the CIDOC CRM ontology, this study outlines a structural approach to transforming linear historical narratives into explicit, machine-interpretable graphs. Drawing on the professional and commercial networks of the Georgian and Victorian British music trade as a case study, the paper identifies five distinct forms of temporal uncertainty: circa dates, bounded dates, time-spans, relative dates, and inherently fuzzy historical eras. Through the strategic interplay of CIDOC CRM's four temporal boundary properties (P82a, P81a, P81b, and P82b), the paper demonstrates how varying levels of temporal precision and modeling-imposed constraints can coexist within a single interval-based representation. Ultimately, this research argues that formalizing temporal uncertainty into explicit semantic statements does not flatten humanistic nuance. Instead, it provides a structural framework that makes complex historical data actionable for advanced workflows, such as automated reasoning and large-scale pattern analysis.

**Keywords:** temporal uncertainty, CIDOC CRM, semantic web, RDF, event modeling, historical data, British music trade, ontology

## 1 Introduction

As computational methods become increasingly embedded in humanities research, scholars have been compelled to reconsider how historical knowledge is structured and represented. This paper contends that data modeling does more than organize historical information consistently and transparently; it actively enables new forms of research by making historical evidence amenable to computational analysis and interpretation. Such a shift challenges traditional humanistic approaches to evidence and narrative, inviting scholars to rethink the relationship between interpretation and formal representation.

This growing reliance on computational methods raises a fundamental epistemological tension between interpretive practice and computational formalization. Elena Pierazzo highlights this tension, asserting that “to compute, we need to model, and we need to do it formally, explicitly, and accurately, as sloppiness or fuzziness strongly limits the potentials of the kind of process operated by machines” [17, p. 120]. From this perspective, computational processing requires historians to translate uncertainty into structured, explicit representations, approaching what Pierazzo calls a “deterministic certainty” [17, p. 120].

Johanna Drucker, on the other hand, emphasizes that humanistic data is not observer-independent, unlike data in the natural and social sciences. Nevertheless, she argues that this does not exclude such data from computation. Instead, its complexity “requires the same mathematical and computational models as other complex systems” [7, para. 18]. For Drucker, the goal is not

---

Yu Lee An. “Reconstructing Historical Events with CIDOC CRM.” In: *From Manuscript to Megabyte: Building Sustainable and Open Futures in the Digital Humanities*, ed. by James B. Harr III, Emma Stanley, Jenna Rinalducci, and Donna Kain. Vol. 5. Anthology of Computers and the Humanities. 2025, 38–47. <https://doi.org/10.63744/AfxZkbFK4dXG>.

© 2025 by the authors. Licensed under Creative Commons Attribution 4.0 International (CC BY 4.0).

to eliminate interpretation, but rather to encode it within a “humanistic interpretative frame that computational systems can meaningfully process” [7, para. 49].

Both Pierazzo and Drucker, in different ways, highlight the challenges and possibilities of representing ambiguity and uncertainty in historical data computationally. A central challenge concerns temporal uncertainty, which makes computational approaches to organizing and interpreting sources particularly valuable. Given these uncertainties, subjectivities, and gaps in the historical record, historians have often privileged narrative over formal modeling, as narrative allows them to convey the layered meanings embedded in historical sources. In response, this study draws on Semantic Web technologies, which enable the explicit encoding of meaning in machine-readable form and support the integration and reuse of data across heterogeneous sources.

Using the Georgian and Victorian British music trade as a case study, this paper demonstrates how semantic modeling and computational methods can facilitate the analysis of historical events characterized by temporal uncertainty. The goal is to show how such methods transform both the representation and the interpretation of historical evidence.

## 2 About the Data

Members of the British music trade included music publishers, engravers, music sellers, musical instrument makers, and accessory makers—in short, anyone participating in the trade. The data describes connections among music traders, capturing successor–predecessor and family relationships, business associations, and the events that shaped their professional lives.

The dataset covers events tied to individuals in the trade during the Georgian and Victorian eras. Reconstruction of their activities draws on diverse primary and secondary sources, including newspaper advertisements, *London Gazette* legal notices, personal correspondence, music publishers’ catalogs, and Stationers’ Hall registers, which recorded publications entered for copyright protection.

These events fall into four main categories. Life events include milestones such as birth, marriage, death, career changes, and relocations, providing context beyond business activities. Critical business events mark significant changes, including the formation of new partnerships, the purchase of the stock-in-trade of other music sellers, and business closures, all of which signal strategic shifts. Everyday business events encompass routine commercial and professional activities such as music printing, publishing, and engraving; musical instrument making, selling, and hiring out, placing newspaper advertisements for new publications, and maintaining circulating music libraries. Although individually minor, these activities collectively reveal long-term patterns in the music trade. Business failure events include bankruptcy, insolvency, and imprisonment for debt, financial loss and bankruptcy discharges. Together, these four categories capture the range of events that shaped music traders’ lives and professional activities in Georgian and Victorian Britain.

## 3 How Semantic Web Technology Can Help

Semantic Web technologies can help historians by providing ways to encode data with precise contexts—embedding the meaning of every event, entity and relationship directly into the data itself. As Albers, Große, and Wagner observe, semantic modeling is the key to connecting the data content and the context [1, p. 239]. This is achieved by translating implicit, ambiguous human-language knowledge into explicit, machine-interpretable representations, shifting part of the information-processing burden from humans to machines.

Central to this approach is the Resource Description Framework (RDF), which represents information as an interrelated graph structure. As Martin, Szekely, and Allemang note, “A graph is the lowest common denominator for data representation,” and “is not limited to a particular

shape” [14, p. 16]. As such, it provides a universal form capable of describing a wide variety of entity types, attributes, and relationships with full fidelity. Graphs provide expressive power for modeling complex scenarios and indirect relationships that are difficult to capture with relational database models, and they scale naturally as datasets grow. Creating a complex semantic model is now feasible through both commercial and non-proprietary graph databases and graph processing software, making semantic web technologies increasingly accessible to non-technical domain experts [13, p. vii].

Open standards such as RDF Schema (RDFS), the Web Ontology Language (OWL), and the SPARQL Protocol and RDF Query Language (SPARQL) further enable interoperability among datasets and support the linking of related information across heterogeneous sources. Additionally, RDF reification allows scholars to represent multiple interpretations and record the provenance of historical information through formal vocabularies and ontologies [2, pp. 54, 57][12, pp. 81–82], thereby supporting the integration of data-driven approaches with traditional narrative history.

## 4 Modeling


Events in the British music trade can be modeled as RDF triples, forming a graph using the event-oriented CIDOC CRM ontology developed by the International Council of Museums (CIDOC). CIDOC CRM is widely used in the cultural heritage sector and is particularly well-suited for modeling historical data. This research employs the Erlangen CIDOC CRM, an RDF Schema expression of the CRM [6]. CIDOC CRM treats events as primary entities.

Modeling historical events involves defining and relating entities through their temporal relationships—a process complicated by pervasive uncertainty. The National History Standards state, “Chronology provides the mental scaffolding for organizing historical thought” [15, p. 7]. But historians’ scaffolding is often unstable. Historian David Fischer warns against the “fallacy of misplaced precision” [8, p. 61], emphasizing that overstating empirical certainty is problematic. Effective modeling, therefore, formalizes uncertainty rather than eliminating it, enabling machines to process uncertainty through logic-based languages. By encoding these uncertainties, historians can better manage inherent ambiguities.

Within the CIDOC CRM framework, the `E5_Event` class serves as the structural nexus connecting diverse historical entities. To illustrate: the 1781 keyboard duel between Mozart and Clementi—arranged by Emperor Joseph II of Austria at the Viennese court [9, p. 357]—is modeled as an `E5_Event` (see Figure 1). Often referred to as the “Clementi–Mozart duel,” the event was a keyboard competition in the presence of Emperor Joseph II, the Grand Duke Paul of Prussia, Grand Duchess Maria Fyodorovna of Prussia, and members of the imperial court. During the evening, the two musicians performed solo works, improvised at the keyboard, and participated in a two-piano performance. Although the contest was officially declared a draw, it became one of the celebrated musical events of the late eighteenth century.

The relevant actors—Mozart, Clementi, the Emperor, and the Duke and Duchess of Prussia—are represented as instances of `E39_Actor`; the event’s date, December 24, 1781, as an `E52_Time-Span`; and the Imperial Court in Vienna as an `E53_Place`. The `E5_Event` acts as the central hub, using semantic properties such as `P12_occurred_in_the_presence_of`, `P7_took_place_at`, and `P4_has_time-span` to integrate distinct entities into a structured, machine-interpretable graph.

While humans may effortlessly synthesize the Mozart and Clementi piano duel into a single mental image, computational systems require this event to be decomposed into explicit data connections. Computational systems cannot reliably infer implicit relationships unless these are encoded as structured data elements. By representing events as distinct entity nodes and applying properties such as `P12_occurred_in_the_presence_of` (linking `E39_Actor`) and `P7_took_place_at` (linking `E53_Place`), a linear historical narrative is transformed into a



**Figure 1:** The 1781 musical encounter between Muzio Clementi and Wolfgang Amadeus Mozart, expressed in CIDOC CRM.

highly contextualized graph. This approach yields a model representing entities, events, and a network of relationships—including those involving music publishers, printers, instrument makers, and other participants.

The diagram in Figure 1 represents the following elements:

- E5\_Event: the overarching event itself—the duel/meeting.
- E39\_Actor: the people involved—Mozart, Clementi, the Emperor, and the Duke and Duchess of Prussia.
- E53\_Place: where it happened (the Austrian Court/Palace).
- E52\_Time-Span: when it happened (December 24, 1781).
- P12\_occurred\_in\_the\_presence\_of: the semantic property linking the event to the actors who were there.
- P7\_took\_place\_at: the semantic property linking the event to its physical location.

Each triple connects actors (people or organizations), events, and places, and represents their relationships. Actor data identifies individuals or organizations and records the various roles they played, such as music publisher, engraver, seller, or musical instrument maker. Event data describes what happened and when, including time-spans. Relationship properties capture partnerships, family ties, successor-predecessor links, and broader business associations. Spatial data records business locations and changes over time. Temporal uncertainties are captured by CIDOC CRM properties.

Time itself requires careful consideration; it can be represented as either a quantitative or qualitative entity [4, p. 3]. When considering time quantitatively, we usually refer to either a point in time or an interval. In this study, all events are modeled as having beginnings and endings and cannot possess zero duration. This research aligns with the point of view expressed by Freksa [10], Allen [3], and Hamblin [11], asserting that representing time as truly instantaneous points with zero

duration is illogical. Papadakis, Doerr and Plexousakis, too, argue that “there are no instantaneous observable phenomena” [16, p. 4], meaning all events must have a spatio-temporal extent. This position forms the foundation for modeling time in this study.

## 5 The Challenge of Representing Time in Historical Data

The British music trade dataset contains numerous temporal references and modeling them presents a significant challenge because historical dates are often imprecise, incomplete, and uncertain. Although chronology is fundamental to understanding historical processes, exact dates and precise temporal boundaries are often unavailable in historical sources. Temporal information is commonly expressed through qualitative, natural-language phrases, such as “in the early days of,” “up to about,” and “was for a time.” Historians routinely interpret such expressions by drawing on contextual knowledge and linguistic cues. Computers, however, cannot reliably interpret these meanings without formal representations.

This research identifies five principal forms of temporal uncertainty.

### 5.1 Circa Dates

The first category consists of circa dates, which may be expressed numerically or in natural language. In these cases, a date is provided in the source, but some degree of uncertainty surrounds its accuracy. Examples include “about 1778,” “circa 1800,” or “near the year 1800.” Such dates establish a likely temporal reference point while acknowledging potential deviation from the stated date. When derived from reliable sources, these dates are often considered highly probable, although some uncertainty remains.

### 5.2 Dates Expressed through Temporal Boundaries

The second category includes dates qualified by expressions such as “from,” “before,” or “after” a specific date. Here, a temporal boundary is known, but the duration or extent beyond that boundary remains uncertain. For example, the phrase “from 1650 onward” establishes a lower temporal bound but does not indicate how long the associated event continued. Similar examples include “before the year xxxx,” “after xxxx,” “shortly after this date,” “prior to this time,” “sometime about or prior to the year xxxx,” and “up to about xxxx.”

### 5.3 Time-Spans

A third form of uncertainty involves time-spans, of which two distinct cases can be identified. The first concerns events known to have occurred at a particular point in time, although the exact date remains unknown: historians may know the earliest possible date (*terminus post quem*) and the latest possible date (*terminus ante quem*). The second case concerns events that span a period—such as wars, monarchical reigns, and large-scale social, economic, technological, and cultural transformations—in which uncertainty may exist regarding their beginning, their end, or both. One of the temporal boundaries may therefore be represented by approximate or fuzzy dates.

### 5.4 Relative Dates

A fourth category consists of relative dates, in which temporal information is defined through relationships to other events rather than through explicit calendar dates. Although the exact timing may be unknown, historians may know that event A bears temporal relation to event B, such as occurring before, after, during, overlapping with, starting, or finishing it. Such relationships are

commonly analyzed using Allen’s Interval Algebra [3] and provide valuable temporal information even when absolute dates are unavailable.

### 5.5 Inherently Fuzzy Dates

The final category encompasses inherently fuzzy dates. Some temporal expressions, such as “at a later date” and “many years before,” lack sufficient specificity to support precise dating. More importantly, certain historical phenomena are inherently uncertain, and historians generally do not attempt to assign precise start and end dates. The Industrial Revolution provides a notable example. Historians no longer accept the view that it began suddenly in the 1750s and concluded in the 1830s; rather, they regard it as a long-term process that began earlier and extended into the early twentieth century [5, pp. 24–25][18, p. 26][19, p. 196]. In such cases, assigning precise start and end dates is neither realistic nor historically meaningful.

Throughout the modeling process, a balance is maintained between two approaches: a conservative approach that encodes only information explicitly stated in the sources, and an interpretive approach that fills gaps using informed domain knowledge, provided such inference is supported by evidence.

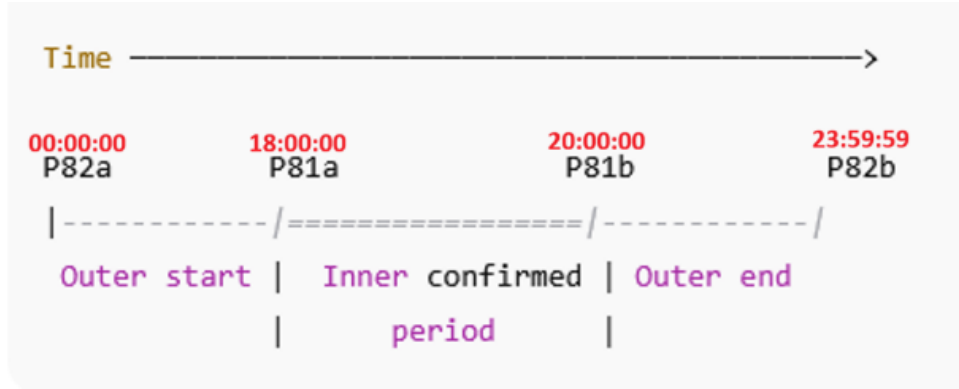
## 6 Modeling Temporal Uncertainty

The preceding discussion demonstrates that historical dates are rarely precise and often appear as bounded, relative, or inherently fuzzy temporal expressions. A semantic model must therefore accommodate varying degrees of certainty rather than assume exact dates and times. CIDOC CRM is particularly well-suited to this task because it allows events and their temporal extents to be described explicitly while preserving inherent uncertainty. As an event-centric ontology, it enables the documentation not only of what occurred, but also of participants, locations, and broader interrelationships associated with an event.

CIDOC CRM represents varying degrees of temporal uncertainty through the combined application of four temporal boundary properties: `P82a_begin_of_the_begin`, `P81a_end_of_the_begin`, `P81b_begin_of_the_end`, and `P82b_end_of_the_end`. The P82 pair establishes the outer limits within which an event may have occurred, while the P81 pair defines the interval during which the event is known to have been ongoing based on available evidence (see Figure 2). Taken together, these properties distinguish what is possible from what is certain, providing a formal mechanism for representing temporal boundaries with differing levels of precision.

To illustrate how CIDOC CRM represents uncertain time-spans, consider the aforementioned encounter between Wolfgang Amadeus Mozart and Muzio Clementi on Christmas Eve 1781. Although this event occurred in Vienna and predates Clementi’s publishing career, it is included in the dataset because of his later prominence within the British music trade. Following his move to London, Clementi became one of the most influential music publishers, piano manufacturers, and musical entrepreneurs in Georgian London. The example illustrates how an event-centric model can capture historically significant events associated with important actors, even when those events occur outside the dataset’s immediate geographical or chronological focus.

The duel also exemplifies a common type of temporal uncertainty found in historical datasets. Contemporary sources—specifically Mozart’s correspondence with his father—confirm that the event took place on the evening of 24 December 1781 [9, p. 357], yet they omit the precise start and end times. The date of the event, therefore, can be established with considerable confidence (the outer boundary), while its internal temporal extent remains uncertain. Because the encounter is known to have occurred in the evening, it is reasonable to assume that the duel was ongoing



**Figure 2:** Expressing time-span (Clementi–Mozart duel).

between 6:00 PM (18:00:00) and 8:00 PM (20:00:00). The four temporal boundary properties mentioned above can formalize this hypothesis.

Together, these four properties define a temporal window within which the event may have occurred. They indicate that the duel was plausibly in progress between 18:00:00 and 20:00:00 while acknowledging that the exact duration cannot be determined from the available evidence. Importantly, the same four properties apply regardless of whether the temporal boundaries are uncertain or known with precision. If the event were documented to have started precisely at 18:00:00 and concluded at 20:00:00, the same four properties would be deployed. In that scenario, the values of P82a and P81a would coincide, as would the values of P81b and P82b, thereby expressing temporal certainty rather than uncertainty.

- P82a\_begin\_of\_the\_begin = the earliest possible time (*terminus post quem*).
- P81a\_end\_of\_the\_begin = the earliest confirmed time.
- P81b\_begin\_of\_the\_end = the latest confirmed time.
- P82b\_end\_of\_the\_end = the latest possible time (*terminus ante quem*).

A different temporal structure is exhibited by the bankruptcy notice of Andrew Loder—a music seller, dealer, and chapman of Bath—published in the *London Gazette*. The notice required Loder to surrender himself to the Court of Commissioners of Bankrupt at three scheduled sittings, each starting at 11:00:00 AM at the White Lion Inn in Bath (Figure 3).

In CIDOC CRM terms, each sitting is represented as an interval whose beginning is precisely specified in the source. This is encoded as a fully constrained lower boundary, where P82a\_begin\_of\_the\_begin and P81a\_end\_of\_the\_begin converge due to complete certainty about the interval’s onset. However, the notice does not specify the end time of these sittings. To maintain a consistent interval structure under partial specification, this study introduces a modeling convention of a minimum duration of an hour. This is a formal constraint on the representation of temporal extents, partially derived from historical evidence of bankruptcy proceedings. Under this convention, each interval is extended by a minimum modeled duration of one hour, so P81b\_begin\_of\_the\_end is set to 12:00 PM for an 11:00 AM start.

Because the outer boundary is bounded at the day level, P82b\_end\_of\_the\_end is mapped to the final moment of the attested calendar day (i.e., 1827-03-19T23:59:59 for the first sitting). The result is an asymmetrical temporal structure in which the beginning of the event interval is precisely pinpointed to the hour, while the end remains bounded at the broader granularity of the calendar day. This configuration explicitly separates source-attested temporal information from

**W**HEREAS a Commission of Bankrupt is awarded and issued forth against Andrew Loder of Orange Grove in the City of Bath Music-Seller, Dealer and Chapman and he being declared Bankrupt is hereby required to surrender himself to the Commissioners in the said Commission named, or the major part of them, on the 19th and 31st of March instant, and on the 27th days of April next, at Eleven in the Forenoon on each day, at the White Lion Inn, in the said City of Bath, and make a full discovery and disclosure of his estate and effects, when and where the Creditors are to come prepared to prove their debts, and at the second sitting to choose Assignees, and at the last sitting the said Bankrupt is required to finish his examination and the Creditors are to assent to or dissent from the allowance of his certificate. All persons indebted to the said Bankrupt, or that have any of his effects, are not to pay or deliver the same but to whom the Commissioners shall appoint, but give notice to Mr. Thomas Turner Solicitor, 6, Harington-Place, Queen-Square, Bath, or Mr. William Price, Solicitor 1, New-Square Lincoln's-Inn, London.

**Figure 3:** Andrew Loder bankruptcy notice. *London Gazette*, March 27, 1827, <https://www.thegazette.co.uk/London/issue/18344/page/646>. Different data elements have been color-coded: events in yellow, dates and times in blue, places in red, and actors in green.


modeling-imposed constraints. It demonstrates how CIDOC CRM accommodates hybrid temporal configurations in which varying levels of precision coexist within a single interval-based representation and are systematically expressed through the relationships between the P81 and P82 property pairs.

## 7 Conclusion

Historians encounter unique challenges as information modelers, given that documentary evidence is often scarce, fragmented, or inconclusive. Temporal uncertainty is not merely a hurdle to overcome, but an enduring aspect of reconstructing the past. This study has shown how diverse forms of temporal uncertainty—such as circa dates, dates expressed through temporal boundaries, time-spans, relative dates, and inherently vague historical periods—can be systematically formalized for computational purposes.

By utilizing the event-centric framework of CIDOC CRM—specifically, the interplay between the P81\_ongoing\_throughout and P82\_at\_sometime\_within temporal properties and their associated boundary subproperties (P81a, P81b, P82a, and P82b)—ambiguous natural language and incomplete chronological markers can be transformed into rigorous, machine-interpretable data structures without imposing unwarranted precision. The two case studies presented here, the Clementi–Mozart keyboard duel and the Andrew Loder bankruptcy notice, illustrate how this framework handles both partial and asymmetrical temporal knowledge within a consistent formal structure.

Ultimately, this modeling process is a crucial first step in making complex historical texts actionable for machine intelligence. By formalizing temporal ambiguity within a standardized ontology, historians enable advanced computational workflows—such as executing sophisticated SPARQL queries across diverse datasets, using automated reasoners to infer missing chronological connections, and analyzing long-term patterns within the music trade. Transforming temporal uncertainties in historical sources into explicit semantic RDF triples does not diminish historical



**Figure 4:** Graph representation of the Andrew Loder bankruptcy notice.

nance; rather, it supplies the computational scaffolding needed to extend humanistic interpretation into the digital era.

## References

- [1] Albers, Laura, Große, Peggy, and Wagner, Sarah. “Semantic Data-Modeling and Long-Term Interpretability of Cultural Heritage Data – Three Case Studies”. In: *Digital Cultural Heritage*, ed. by Horst Kremers. Cham, Switzerland: Springer, 2020, pp. 239–253.
- [2] Allemang, Dean, Hendler, James, and Gandon, Fabien. *Semantic Web for the Working Ontologist: Effective Modeling for Linked Data, RDFS, and OWL*. 3rd. New York, NY: ACM, 2020.
- [3] Allen, James F. “Maintaining Knowledge About Temporal Intervals”. In: *Communications of the ACM* 26, no. 11 (1983), pp. 832–843.
- [4] Batsakis, Sotiris, Petrakis, Euripides G. M., Tachmazidis, Ilias, and Antoniou, Grigoris. “Temporal Representation and Reasoning in OWL 2”. In: *Semantic Web – Interoperability, Usability, Applicability* (2009), pp. 1–19. DOI: 10.3233/SW-160248.
- [5] Berg, Maxine and Hudson, Pat. “Rehabilitating the Industrial Revolution”. In: *The Economic History Review* 45, no. 1 (1992), pp. 24–50. DOI: 10.2307/2598327.
- [6] Doerr, Martin, Light, Richard, and Hiebel, Gerald. “Implementing the CIDOC CRM Conceptual Reference Model in RDF Version 1.1”. Tech. rep. CIDOC CRM Special Interest Group, 2020. URL: <https://cidoc-crm.org/implementing-the-cidoc-conceptual-reference-model-in-rdf>.
- [7] Drucker, Johanna. “Humanities Approaches to Graphical Display”. In: *Digital Humanities Quarterly* 5, no. 1 (2011).
- [8] Fischer, David Hackett. *Historians’ Fallacies: Towards a Logic of Historical Thought*. New York, NY: Harper Perennial, 1970.

- [9] Foix, Georges De Saint and Martens, Frederick H. “Muzio Clementi (1752-1832)”. In: *The Musical Quarterly* 9, no. 3 (1923), pp. 350–382. URL: <http://www.jstor.org/stable/738494>.
- [10] Freksa, Christian. “Temporal Reasoning Based on Semi-Intervals”. In: *Artificial Intelligence* 54, no. 1 (1992), pp. 199–227. DOI: 10.1016/0004-3702(92)90090-K.
- [11] Hamblin, C. L. “Instants and Intervals”. In: *The Study of Time: Proceedings of the First Conference of the International Society for the Study of Time, Oberwolfach (Black Forest) – West Germany*, ed. by J. T. Fraser, F. C. Haber, and G. H. Müller. Berlin, Heidelberg: Springer, 1972, pp. 324–331.
- [12] Hogan, Aidan. *The Web of Data*. Cham, Switzerland: Springer, 2020.
- [13] Kejriwal, Mayank. *Domain-Specific Knowledge Graph Construction*. SpringerBriefs in Computer Science. Cham, Switzerland: Springer, 2019.
- [14] Martin, Sean, Szekely, Ben, and Allemang, Dean. *The Rise of the Knowledge Graph: Toward Modern Data Integration and the Data Fabric Architecture*. 1st. Sebastopol, CA: O’Reilly, 2021.
- [15] “National History Standards, Part I: Standards in Historical Thinking”. In: *OAH Magazine of History* 9, no. 3 (1995), pp. 7–12. URL: <http://www.jstor.org/stable/25163028>.
- [16] Papadakis, Manos, Doerr, Martin, and Plexousakis, Dimitris. “Fuzzy Times on Space-Time Volumes”. In: *eChallenges e-2014 Conference, 29-30 October 2014, Belfast, Ireland*, ed. by Paul Cunningham and Mariam Cunningham. Danvers, MA: IIMC International Information Management Corporation, 2014, pp. 1–11.
- [17] Pierazzo, Elena. “How Subjective Is Your Model?” In: *The Shape of Data in the Digital Humanities*, ed. by Julia Flanders and Fotis Jannidis. Digital Research in the Arts and Humanities. Abingdon, Oxon: Routledge, 2019, pp. 117–132.
- [18] Stobart, Jon. “An Eighteenth-Century Revolution? Investigating Urban Growth in North-West England, 1664-1801”. In: *Urban History* 23, no. 1 (1996), pp. 26–47. DOI: 10.1017/S0963926800011664.
- [19] Withers, C. W. J. “Mapping the Industrial Revolution”. In: *Journal of Historical Geography* 11, no. 2 (1985), p. 196.