Interactive Story Maps for Historical Musical Instruments: A 3D and Semantic Web Tool for Cultural Heritage Preservation

Valentina Bartalesi[1*], Emanuele Lenzi[1,2], Claudio De Martino[1], Gianpaolo Coro[1*]

[1]Institute of Information Science and Technologies (ISTI), National Research Council of Italy (CNR), Via G.Moruzzi, 1, Pisa, 56124, Italy.

[2]Department of Information Engineering (DII), University of Pisa, Via G.Caruso, 16, Pisa, 56122, Italy.

*Corresponding author(s). E-mail(s): valentina.bartalesi@isti.cnr.it; gianpaolo.coro@isti.cnr.it;
Contributing authors: emanuele.lenzi@isti.cnr.it; claudio.demartino@isti.cnr.it;

Abstract

Ancient stringed instruments from the 17th and 18th centuries are essential for global cultural heritage but pose preservation and study challenges due to their fragility and rarity. Digital reproductions and interactive story maps offer new ways to analyse these instruments, providing historical, mechanical, and acoustic insights while aiding modern luthiers in replication. This paper presents a novel methodology and software that integrates 3D models of historical instruments into interactive story maps using Semantic Web technologies. The system enables users to explore 3D models with embedded annotations from Sketchfab and supports the creation of formal narratives enriched with geospatial and multimedia content. Each narrative event is linked to a backend knowledge base (KB) structured with established ontologies. The tool is open-source and web-based, and allows interoperability between its story maps and external KBs such as Wikidata and Europeana. It supports multiple levels of descriptive details, from precise 3D model annotations to broad geographic and temporal references. An evaluation case study featured 25 participants exploring a story map about a 1737 Antonio Stradivari violin, connected to a KB of 1,716 triples. The results highlight the tool's usability and effectiveness in representing the violin's spatiotemporal and cultural context, showcasing the value of combining 3D models and geospatial data to enhance cultural heritage preservation.

Keywords: 3D models, Digital Twins, Narratives, Knowledge Representation, Ancient musical instruments

1 Introduction

The preservation of cultural heritage (CH) is essential for comprehending historical contexts, as it establishes a tangible link to the traditions, knowledge, and values of antecedent generations [1–3]. Concurrently, it contributes to shaping the future by fostering a sense of identity and continuity, which is fundamental for the development of resilient and inclusive societies. The protection of CH facilitates the bridging of cultural divides, enhances mutual understanding, and stimulates innovation through the integration of diverse human experiences. Specifically, the examination of ancient CH artefacts yields significant insights into ancestral lifestyles, production techniques, and socio-cultural practices while ensuring their conservation for posterity. Moreover, advancements in technology have enabled increasingly precise methods for preserving and representing CH artefacts, thereby improving their documentation, restoration, and replication, particularly concerning ancient musical instruments [4–7].

Among the most esteemed cultural heritage artefacts, ancient musical instruments are preserved in museum collections globally; however, they are predominantly no longer played. In contrast, stringed instruments such as violins, violas, and cellos continue to be actively performed in concerts worldwide and are rarely displayed in museums. Notably, these instruments often maintain their original acoustic quality and functionality centuries after their construction [8–10]. This phenomenon is particularly evident in violins crafted by renowned master luthiers from the seventeenth and eighteenth centuries. Contemporary luthiers frequently encounter challenges in replicating the superior craftsmanship of distinguished artisans such as Antonio Stradivari and Guarneri del Gesù. Furthermore, professional musicians seek high-quality instruments and aspire to perform on these historic masterpieces. Consequently, a comprehensive understanding of the techniques employed by these master luthiers would provide significant benefits to both instrument makers and performers. Additionally, ancient violins and the music they produce constitute an integral component of both tangible and intangible global cultural heritage, offering substantial cultural value to the broader public when contextualised within their historical and cultural frameworks [11, 12].

Historical stringed instruments currently played by musicians in concert settings are inherently fragile, rendering direct access and examination challenging. Nonetheless, the investigation of these instruments can be facilitated through the creation of a virtual, functional counterpart known as *digital twin*. A digital twin is a three-dimensional (3D) model generated via 3D scanning techniques [13, 14], such as photogrammetry or micro-computed tomography [15–18], of the physical instrument, enabling comprehensive analysis of its mechanical and acoustic characteristics alongside its historical and cultural significance. A digital twin extends beyond a mere geometric representation by integrating historical, socio-cultural, and acoustical data [19, 20]. Typically, it encompasses a physical model detailing mechanical and acoustic sound production mechanisms [21–23]. The contextual information embedded within includes: (i) a *general contextualisation* outlining the temporal and geographical origins of the instrument; (ii) an account of relevant *socio-political circumstances*; and (iii) a *detailed contextualisation* concerning individual components as well as construction and maintenance interventions performed by luthiers over time. These layers of contextualisation can effectively be represented by embedding the three-dimensional model of the instrument within a digital narrative framework, supported by a geographical visualisation such as a map to chronologically and spatially track narrative events. Such map-based digital narratives, commonly named *story maps*, are typically implemented as interactive online maps enriched with textual content, images, videos, data, and multimedia elements [24]. The three-dimensional digital model of a musical instrument may be augmented through a story map comprising a sequence of events annotated with spatial coordinates and temporal references, thereby situating the instrument and its components both geographically (on the map) and historically (over time). The events incorporated into the story map can convey detailed information regarding the materials and distinctive features of the instrument, as well as the locations and conditions that influenced its artisan's craftsmanship, thus facilitating contextualisation from general to specific levels. Drawing upon Korzybski's General Semantics theory [25] and corroborating studies [26–30], story maps enable a perception of emotive dimensions beyond those afforded by a three-dimensional model alone, rendering them particularly suitable for describing musical instruments. Furthermore, story maps can exploit contemporary technological advancements in knowledge representation to construct open and interconnected knowledge graphs that link an individual instrument to related objects and events [31–35].

This paper introduces a methodology and accompanying software for the integration of three-dimensional models of historical musical instruments into story maps. We enhanced the Story Map Building and Visualising Tool

(SMBVT) — a platform that uses Semantic Web technologies to represent spatiotemporal narrative data [30, 36] — to enable the management of story map events linked to 3D models and their constituent components, thereby facilitating the depiction of their geographic and cultural contexts. Consequently, we developed and released a novel, open-source, and freely accessible tool designed to construct story maps based on 3D models of musical instruments as well as cultural heritage artefacts more broadly.

As a case study, we present a story map developed for a 1737 violin crafted by Antonio Stradivari. The story map elucidates the instrument's historical dimensions, geographic trajectories, temporal progression, and socio-cultural contexts, while facilitating the exploration of high-resolution three-dimensional models. Through a user-test evaluation, we examine the effectiveness of our tool in representing the instrument's spatiotemporal and historical-cultural attributes and enhancing the overall experience of 3D model visualisation, inspection, and navigation. The findings indicate that our tool holds significant potential for advancing the comprehension of cultural heritage artefacts. Although our story maps do not constitute digital twins — owing to the absence of physical or acoustic modeling - they serve as essential contextual frameworks supporting digital twin development.

This paper is organised as follows: Section 2 reports about state-of-the-art 3D model integration with story maps and musical instrument description through Semantic Web technologies. Section 3 describes our methodology and software to accomplish this task. It also describes the case study and the evaluation methodology. Section 4 reports the evaluation results. Finally, Section 5 comments on the results and concludes.

1.1 Principal research objectives

In the following, we list the principal objectives of this research:

- Preservation and Representation of Cultural Heritage (CH): We aim to enhance the study and preservation of CH artefacts, especially ancient musical instruments. We provide an interactive way to explore CH artefacts by integrating 3D models with historical, geographical, and socio-cultural contexts.
- Development of a 3D Story Mapping Tool: To reach the previous objective, we propose an extension of the Story Map Building and Visualising Tool to integrate 3D models and spatiotemporal events with digital narratives.
- Knowledge Representation and Interoperability: We facilitate open, interconnected access to information about CH artefacts, leveraging Semantic Web technologies and a knowledge representation approach. This allows the structuring and interlinking of historical and contextual information within a knowledge graph.
- User Evaluation and Enhancement of Digital Interaction: We assess the effectiveness of 3D story maps in conveying historical and cultural contexts through a case study describing a 1737 Stradivari violin. We evaluate the quality of the user interaction with the story map creation tool and the story's effectiveness in improving the understanding of the CH artefact.

2 Related works

This section provides an overview of software that integrates 3D models into story maps, potentially utilising knowledge representation technologies. These approaches have features and capabilities that can be fairly compared with our solution to highlight the complementary aspects of our approach.

In [37], the author proposed enhancing user engagement with story maps by converting traditional two-dimensional map components into dynamic three-dimensional animations. The study employed machine learning to differentiate pictorial maps from other types, subsequently focusing on detecting objects within historical maps and transforming them into interactive 3D models designed to guide users through quizzes and virtual reality experiences.

An investigation into story maps and 3D data visualisation was conducted and documented in [38]. This research examined the ability of 3D visualisation to effectively represent complex spatiotemporal and semantic constructs. By developing novel 3D prototypes based on Space-Time Cubes [39], the study demonstrated the efficient aggregation and presentation of spatiotemporal datasets, thereby enhancing the understanding of historical phenomena. Story maps served as tools for geospatial storytelling, integrating interactive and multimedia components to create engaging geographic narratives.

In [40], the authors developed a GIS-based web application to convey the history, biodiversity, geodiversity, and cultural features of Salamis Island (Greece) through an interactive story map. Using photogrammetry, they reconstructed 3D models of culturally significant sites to support the narrative. Similarly, [41] introduced a story map that investigated the geomorphology and history of the Methana Peninsula (Greece), incorporating 3D representations that enabled interactive engagement with the peninsula's distinct geological characteristics.

ArcGIS StoryMaps is a web-based application allowing users to create interactive narratives using maps and geographic content [42, 43]. Users can integrate 3D model files into applications like ArcGIS CityEngine [44, 45] and embed 3D content via URL or HTML from external repositories. A limited free version is available, but it restricts map interaction and theme customisation, as well as accessibility and external content embedding, which, however, are crucial for effective story map creation.

Cesium Stories is another commercial platform for creating 3D story maps [46]. It streamlines the creation of 3D narratives from high-resolution data and supports web-based streaming. The platform provides a free account with 5 GB of storage for 3D models and 15 GB for data streaming. Cesium also integrates CesiumJS, an efficient open-source library for visualising 3D story maps [47]. A key limitation of this system is its reliance on geographic representation, which prevents disabling background map visualisation for non-geospatial objects.

The Google Maps Platform offers a digital storytelling tool enabling users to create story maps using photorealistic 3D tiles [48]. This tool caters to (i) journalists for location-based article contextualisation, (ii) travellers for interactive logs, and (iii) educators for immersive experiences. Google also provides a visualisation service for story maps, allowing associations with textual information and YouTube videos, though it restricts other multimedia types and multiple images per event.

Matterport is commercial software creating immersive 3D reproductions of physical spaces using specialised cameras [49, 50]. It allows adding information on a 3D model, enhancing the digital story through model augmentation. Primarily aimed at real estate visualisation and construction documentation, this software offers high-quality representation and augmentation of physical spaces but lacks certain interactive capabilities.

Concerning knowledge representation, the Omeka S web publishing platform offers a visual representation of semantically enriched multimedia data, designed for cultural heritage collections [51]. In the context of musical instruments, semantic descriptions have been explored in studies characterizing violin timbers through sound qualities and instrument components [52], and proposing taxonomies of musical instruments [53]. Research has also described diverse contextual dimensions of instruments, including provenance, usage in compositions, performing artists, and accompaniments [54, 55]. For instance, the Melody platform enables the embedding of semantic descriptions for musical instruments, representing them as narratives [56]. While these approaches have proven effective, they are mainly tailored to musical instruments and specific aspects of cultural heritage preservation, lacking integration of spatial and three-dimensional representations.

Overall, these tools effectively fulfil their designated functions; however, they show limitations in interoperability between narratives generated within the same software and those from different platforms. We have drawn insights from these systems to develop a solution with complementary features: (i) a freely accessible software and web service for embedding externally sourced 3D models; (ii) support for spatiotemporal, geographic, and model-based story maps; (iii) the ability to associate extensive multimedia content with events; (iv) an expanded ontology for semantic description of cultural heritage artifacts and contexts, beyond musical instruments; and (v) a standard-based knowledge graph integrated into the backend that connects all stories created with the tool, enabling interoperability and embeddability across knowledge bases.

3 Methods

This section describes (i) the starting point of our tool for 3D story map building (Section 3.1), (ii) the tool itself (Sections 3.2-3.5), (iii) the case study (Section 3.6), and (iv) the evaluation methodology (Sections 3.7 and 3.8).

3.1 Baseline software

The Story Map Building and Visualising Tool (SMBVT) is a semi-automated system designed for the construction and visualisation of story maps via an intuitive web interface, underpinned by Semantic Web technologies [30]. This tool systematically organises the knowledge that forms the narrative through a formal representation grounded in the Narrative Ontology (NOnt) [57–59], a vocabulary based on CIDOC CRM specifically developed to represent narratives. SMBVT is open-source software providing a freely accessible web service for creating story maps to describe territories beyond conventional cartographic representations [30, 60]. Within this framework, a story is conceptualised as a sequence of events, each annotated with spatial and temporal metadata and visualised on a geographic map. Furthermore, each event is augmented with textual descriptions, hyperlinks, images, and videos displayed upon selection. The narrative can be represented either temporally as a timeline to elucidate the plot or spatially as a map to illustrate its geographical progression. User interaction with events is facilitated through sequential scrolling along the temporal axis or by directly selecting corresponding reference markers on the map.

Each event is linked to a set of entities that constitute its essential components, such as locations, individuals, and organisations. These entities are derived from the Wikidata knowledge base [61] to guarantee the use of standardised vocabulary and concepts. The SMBVT backend populates the NOnt by generating a knowledge base (KB) comprising entities extracted from the events. Given that events may share entities with other narratives created by SMBVT, events and stories are inherently interconnected within the KB. Furthermore, individual stories can be linked to external KB resources, including those available in the extensive Europeana cultural heritage digital library [62–64].

SMBVT was deemed an appropriate foundational software to develop a 3D story map-building tool due to its provision of essential intra- and inter-connection capabilities. Moreover, it had openness features that aligned with our requirements. These attributes rendered it complementary to the alternative systems discussed in Section 2. Further elaboration on SMBVT is provided in [30] and Section 3.2.

To embed 3D models and additional geographic features in the SMBVT narratives, we extended it by introducing new functionalities, summarised as follows:

1. We added the possibility to integrate 3D model navigation, annotations, and user interaction from an external online model visualisation platform [65];
2. We enabled the SMBVT story creation interface to manage 3D scenes, geospatialised geometries, and non-geospatial background images;
3. We introduced new types of events (3D models, map-based events, and general descriptive events), while allowing the creation of sequences of mixed event types within the same story.
4. We enabled the import of overlaying points, polygons, and geometries to define event locations on a map.

By supporting the creation of diverse event types, our system facilitates the construction of narratives in which each event is associated with a detailed and contextually appropriate spatiotemporal description. This methodological framework allows content creators to enhance 3D models by integrating general contextual information and model-specific details through event-related information and targeted annotations. For example, an event may describe the restoration of a particular component of a violin, indicated precisely on the 3D model via a localised annotation, and subsequently linked to a geographical reference on a global map identifying the location where the repair took place.

3.2 The 3D Story Map Building and Visualising Tool

The SMBVT provides a graphical interface that facilitates the creation of narratives by users. The story creation process initiates when the user inputs a title that encapsulates the central theme of the story. Subsequently, the system automatically associates this title with a relevant Wikipedia entity, referred to as the *subject entity*, and presents it to the user for confirmation. Following this step, the tool employs Wikipedia's official APIs and SPARQL Query Service [66, 67] to retrieve all pages that are directly linked to the subject entity, alongside pertinent information regarding associated entities, including their names, descriptions, and images. Utilising a knowledge representation approach, the tool also acquires the Internationalized Resource Identifiers (IRIs) for the extracted entities. Furthermore, users are empowered to define new entities by providing their names, descriptions, IRIs, and NOnt classes.

We extended SMBVT to integrate and contextualise 3D models and geospatial geometries within semantic narratives. In particular, we introduced the capability to create three distinct types of narrative events within a single story (Fig. 1-a): 3D-type events, Map-type events, and Slide-type events.

3D-type events involve three-dimensional models, enabling users to navigate and interact with the models hosted on Sketchfab [65], a widely employed online platform for displaying and sharing 3D models within virtual environments. This platform supports models equipped with materials, lighting, and scalable elements, accessible across mobile, desktop, and virtual/augmented reality devices. Sketchfab facilitates the incorporation of point-specific annotations on the model's surface, each linked to a unique identifier and accompanied by explanatory text. For our research, we have integrated the Sketchfab scene visualisation within the SMBVT interface. This integration prevents users from exiting the story visualisation when a 3D scene is displayed. When users create a 3D-type event, they are required to complete a series of fields common to all event types (Fig. 1-b), including: (i) the event title, (ii) a textual description, (iii) the start and end dates, which denote a reference time frame with granularity ranging from seconds to years, (iv) associated Wikidata entities drawn from a continuously updated list generated by the SMBVT at story creation, (v) the primary sources for the textual description, and (vi) related digital objects (such as web pages, audio files, and entities from external KBs) and media objects (including multiple images and videos to be displayed within the narrative). Furthermore, 3D-type events necessitate additional configuration fields for model scene visualisation, specifically the Sketchfab unique model identifier and the Sketchfab annotation identifier. These fields inform the SMBVT of the specific model to be embedded and the annotation on which the visualisation should concentrate. This methodology also enables model navigation through the annotations within a story (Section 3.3). Moreover, it enhances the information associated with the individual model annotations through supplementary fields and semantic information.

Map-type events are events tied to specific geographic locations indicated on a map. The creation of a map-type event necessitates the completion of the *common* fields required for other event types, from the event title to the associated digital objects. Additionally, it introduces a dedicated *location* field, wherein users are required to input geometry in the Well-Known Text (WKT) format, as specified by the Open Geospatial Consortium [68]. The geometry must conform to the WGS84/EPSG:4326 reference system. To facilitate this process, a user-friendly widget is available, which can be activated via a side button. This tool enables users to either select a point on a global map or delineate a polygon. The widget subsequently converts the drawn geometry into a WKT string. The chosen geometry is then visually represented and zoomed on, providing a reference location or area associated with the event, effectively overlaid on a background world map (Section 3.3).

Slide-type events are inherited from the previous version of the SMBVT. These events are not associated with 3D models or maps; rather, they serve to enhance the overall narrative context and content. Slide-type events ask the creator-users to fill in the *common* fields of all events to define the time and semantic frame of the event. Additionally, they require users to specify a background image as a visual reference for the event. This image should be distinct from those indicated in the multimedia-objects field, which will appear in the side description text (Section 3.3). The background image may be uploaded directly via an import form or supplied through an external URL.

The software architecture of the newly developed 3D Story Map Building and Visualising Tool is depicted in Figure 2. The tool is implemented as an open-source Java application running on a

Tomcat service [69]. Its web interface interacts with back-end services to store narrative information, generate semantic representations, and create story map visualisations. Narratives and their associated events are represented internally using a JSON [70] schema aligned with the Narrative Ontology. The entities involved in a narrative are extracted from Wikipedia and Wikidata during the story creation process. These entities, along with the narrative data, are stored in a PostgreSQL database [71]. All event-related data (e.g., textual data, links to audio, videos, and 3D models) entered via the tool interface are also saved in the database. The tool retrieves the JSON representation of the story from the database. Then, using a Java-implemented triplifier [72], the tool converts the JSON representation of the story into a Web Ontology Language (OWL) graph [73] compliant with the NOnt, and stores it in an Apache Jena Fuseki SPARQL server [74]. The conversion pipeline is inherited from the SMBVT and is described in the related paper [30]. Finally, the tool generates a public URL where the resulting story map can be visualised. If links to 3D models are included in the story, the tool uses the Sketchfab Viewer API [75] to embed the Sketchfab model visualisation in the story map.

Our software architecture inherits the separation between the narrative authoring environment and the semantic KB from the SMBVT. During narrative creation, curators only interact - behind the scenes - with structured tabular and JSON-based representations of entities and events, which support rapid editing, incremental refinement, and usability for non-expert users. The semantic KB is instead conceived as a publication-level artefact, generated once the narrative reaches a coherent state and intended for interoperability, long-term preservation, and semantic querying. This delayed semantic binding allows the system to align narrative content with the domain ontology and external knowledge bases at publication time, while avoiding early and potentially rigid semantic commitments during authoring. Previous deployments of SMBVT [24, 29, 76] have demonstrated that this design supports the medium-term extension of narratives through iterative publication cycles, semantic merging of newly generated graphs, and ontology evolution, while balancing curator usability, software maintainability, and the robustness and extensibility required for complex narrative pathways.

The next section explains how our new version of the SMBVT manages the visualisation of a sequence of multi-type events within one story.

3.3 Visualising mixed event types in a story map

We enhanced the SMBVT story-visualisation layout to accommodate the management of sequences of multiple and diverse event types. The revised layout employs a vertical format for story maps, facilitating navigation through narrative events, in contrast to the prior horizontal configuration, which was predominantly oriented to timeline visualisation. We determined that this vertical layout was more user-friendly and intuitive for managing heterogeneous event sequences, particularly in contexts involving 3D scenes. Such a design better aligns with conceptual event sequences rather than strictly temporal ones, thereby enabling a more fluid perception of the visualisation as it transitions between different 3D annotations.

During the navigation of stories, the visualisation dynamically adapts according to the selected event type. In the case of *3D-type* events, a 3D model is prominently positioned at the centre of the screen. Accompanying descriptions, multimedia elements, and Wikidata entities are systematically arranged on the left-hand side (Fig. 3-a). The 3D model is fully interactive due to the seamless integration of the Sketchfab scene, which empowers users to rotate, resize, and zoom in or out using mouse controls for a comprehensive exploration from various perspectives. When a specific annotation is identified within the event form, the viewer's focus is automatically directed towards the pertinent section of the 3D model. Furthermore, the annotation and associated textual information are displayed within the 3D scene, thereby enriching the contextual information provided in the left-hand panel. This integration augments the Sketchfab 3D scene and annotations with additional contextual text, images, references, and semantic annotations (Wikidata/user-defined entities). Entity correspondence within the KB facilitates the interconnection of stories and the linking to other KBs that include similar objects.

In *Map-type* events, a geographic map occupies the right side of the screen, while textual descriptions and related digital or multimedia objects are presented on the left side (Fig. 3-b). The map highlights the event location by overlaying either a pin or a polygon, as specified in the event form. These locations are enriched by semantic annotations that facilitate connections to other narratives, KBs, and other event types. For instance, this framework enables a direct linkage between a narrative concerning a Stradivari violin produced in Cremona (Italy) and another narrative (possibly in Europeana) regarding Northern Italy that also references Cremona.

In *Slide-type* events, the background image is intentionally darkened to enhance the readability of the overlaid text (Fig. 3-c). Textual descriptions, alongside additional images and multimedia content, may be displayed either centrally or on the left side of the screen, depending on the creator-user's selection within the event form. This event type corresponds to a reference image that is augmented with textual, multimedia, and semantic information, thereby establishing connections to other narratives and KBs.

We implemented the interfaces for three events utilising the StoryMapJS open-source library [77], which facilitated the representation of narrative data in a slide-based format. The original version of this library was limited to the display of slides featuring geographic maps exclusively. Consequently, we extended the functionality of the library to accommodate all event types, enabling the system to (i) manage temporal indicators, (ii) showcase multimedia elements associated with each event, (iii) seamlessly integrate Sketchfab scenes, and (iv) handle events lacking geographic information. We substituted the previous SMBVT visualisation module with this new visualisation module. Furthermore, we expanded the SMBVT story-building interface to incorporate the new event types and execute the new visualisation module. Subsequently, we included these extensions within an updated SMBVT software released as a new open-source solution (Supplementary Material). Additionally, we provide a dedicated, free-to-use service instance designed for constructing and publishing story maps through this tool (Supplementary Material).

3.4 Ontology extension

Reusing the SMBVT provided us with the advantage of inheriting its knowledge management technology. The baseline ontology (NOnt) already supported the description of spatiotemporal events interconnected by semantic relations. We extended NOnt to incorporate concepts corresponding to 3D models and their annotations, which were not natively supported by the previous version of the ontology. Specifically, we integrated appropriate classes and properties from reference vocabularies such as the CIDOC CRM [78] and RDF [79]. In particular, we reused the CIDOC CRM E36 "Visual Item" class to represent pairs of 3D models and model annotations. This class encapsulates the intellectual and conceptual dimensions of signs, images, or other visual works, independent of their physical medium.

Furthermore, we extended the SMBVT to generate a unique IRI for each 3D model. We also enabled the generation of an IRI for each 3D model-annotation pair by concatenating the annotation number to the base 3D model IRI. We reused the RDF [80] "label" property to associate each instance of the "Visual Item" with a human-readable label. To denote the relationship between a 3D model and its annotations, we reused the CIDOC CRM "is composed of" property. Moreover, we reused the CIDOC CRM P67 "refers to" property to connect the 3D models and their annotations to a narrative event. Events were already represented by the CIDOC CRM E5 "Event" class in NOnt. We defined a new property, "hasDescription", within NOnt to attach an overall textual description to a 3D-type event. We described multimedia data (images, videos, audio files, and other digital objects) associated with a 3D-type event by reusing the CIDOC CRM E73 "Information Object" class and linked them to the event through the P67 "refers to" property.

The SMBVT already pre-categorises the entities involved in the event — corresponding to those available in Wikidata or defined by the user — as actors, places, objects, and immaterial items. It represents these entities through the CIDOC CRM classes E39 "Actor", E53 "Place", E19 "Physical Object", and E89 "Propositional Object", respectively. When these correspond to

Wikidata entities, the KB uses the "hasWikidataEntity" NOnt property to connect the corresponding Wikidata IRI to the entity. One new version of SMBV inherited these features.

Upon the publication of a narrative via a dedicated *publication* button, the SMBVT automatically generates an OWL graph of the knowledge encapsulated in the created narrative. This graph is compliant with the NOnt and guarantees adherence to the Linked Open Data paradigm. The OWL graph is automatically stored in a Fuseki triplestore [81], which serves as the system's backend KB and provides a SPARQL endpoint [82, 83] for executing semantic queries. We modified the publication process to incorporate the newly added event types, classes, and properties, thereby populating the KB with the updated version of the NOnt that included the classes and properties explained above.

3.5 Theoretical and practical implications of our research

The added value of our approach, emerging from the previous sections, can be summarised as follows:

1. We allow users to build a story map as a sequence of heterogeneous events involving 3D-model-associated events, geospatial events, and generic multimedia events. The stories support different granularity levels for each event: spatial (from one point to a large area), temporal (from a specific instant to an entire year), 3D (from one point annotated on the model mesh to an entire scene), abstract (from small images accompanying an event to an overall conceptual image).
2. A KB that extends the CIDOC CRM-based NOnt ontology is used as the backend service to augment each event with semantic information. This allows each event to be automatically connected with other reference KBs such as Wikidata and Europeana, independently of the event type. The connection with Wikidata further augments the properties and the descriptions associated with the event through additional data from international communities and Wikipedia. The connection with Europeana, one of the largest CH digital libraries in Europe, allows a user to connect a story (down to a single event) to a larger cultural and touristic context. This aspect is particularly crucial to advertise small territories internationally.
3. The semantic annotations establish a robust connection between diverse event types. For instance, the 3D reproduction of a musical instrument is automatically associated with the event indicating its place of manufacture, as both share the same location entity. Furthermore, the KB ensures that these events are interconnected with all other stories and events created with our tool that reference the same entities. This KB simplifies the retrieval of these interconnections through a single SPARQL query, as demonstrated in [24]. By building upon the SMBVT, we enhance interoperability *between* heterogeneous events and stories, as well as *across* different narratives. This characteristic is distinctive within the realm of story map creation.
4. We released our tool as open-source, free-to-use software, usable without limitations through an open web interface (Supplementary Material). Given the limitations of the tools reported in Section 2, this is another distinguishable feature of our approach.

Overall, the points above indicate that our proposal is novel in the current scenario of story map-building.

3.6 Case Study

To evaluate our tool, we used data provided by "Il Poggio Montecastelli" [84], a Cultural Association located in Montecastelli Pisano, a small medieval village in the heart of Tuscany, Italy. Il Poggio organises events focusing on science and music, hosts scientific experiments in psycho-acoustics [85], and incorporates a violin-making workshop. In this workshop, ancient musical instruments are scanned, reproduced as 3D digital models, and subsequently crafted into physical replicas for use in public concerts. This Association serves as an international hub, attracting professional artists worldwide and hosts musicians who specialise in diverse ancient instruments.

The Association formally agreed to share their precious and unique 3D scan collection as part of a non-profit research initiative titled "ViolTwin". The initiative's objectives were threefold: (i) to assess the feasibility of narrating a story centred around the 3D reproduction of a musical instrument, (ii) to develop a product beneficial for both the general public (providing insight into the history, locales, and attributes of the instrument) and luthiers (enhancing their understanding of craftsmanship), and (iii) to connect the collection, the Montecastelli village, and the Association to a broader European context.

In this paper, we present a comprehensive evaluation of a story map constructed around a violin crafted by Antonio Stradivari in 1737, which is part of the Association's collection (Fig. 3-a). This violin represents one of the last instruments produced by Stradivari, who was 93 years old at the time of its construction. It is probable that his sons, Francesco and Omobono, assisted him due to his advanced age. Nonetheless, Stradivari's original signature remains visible inside the instrument's soundbox: *Antonius Stradivarius Cremonensis / Faciebat Anno 1737*. The physical violin is currently housed at the Staatskapelle Dresden, one of the oldest orchestras globally, established in 1548, and has also been performed in Montecastelli Pisano. The Association's interest in reproducing this violin was motivated by the opportunity to understand, through reproduction, a violin with a distinctive sound that has endured over the centuries, while also establishing a reference that may benefit future luthiers. This violin has a rich history, marked by multiple incidents; it sustained significant damage in two separate events — once in 1800 and again in 1910 — when two individuals accidentally sat on it. However, it has consistently been fully repaired and restored to a playable condition.

The Association provided two high-resolution 3D models in OBJ format: one was created using a *high-resolution* structured-light scanner that focused on the violin's external and internal structure, while the other was produced with a *lower-resolution* photogrammetric scanner that reconstructed the model's mesh and texture. Additionally, the Association shared four reference images and a video related to this instrument. To supplement the information found in official documents and multimedia references, we interviewed master luthier Philipp Bonhoeffer from the Association, who provided insights into significant historical developments as well as the phases of instrument building, maintenance, and restoration. Based on his expertise, the master luthier also identified and suggested the specific structural components and areas of the instrument to be annotated and highlighted within the 3D models. The information provided by the master drew on multiple complementary sources: (i) his direct knowledge gained through discussions with the current and previous owners of the instrument, (ii) specialised technical literature consulted during the reconstruction of a faithful copy of the violin, and (iii) authoritative reference works on the history of the violin and on Antonio Stradivari's production. We created four audio files in MP3 format from this interview to enrich the narrative events of the story. The 3D models were imported into two Sketchfab scenes, where lighting, materials, scale, and both external and internal annotations were configured in accordance with the luthier's recommendations (Supplementary Material). Despite this careful curation process, a degree of uncertainty inevitably remains regarding certain historical aspects of the instrument, as no single direct and fully authoritative source exists for its complete history. This uncertainty was therefore an intrinsic characteristic of the narrative.

Based on these data, we constructed a story map (available in English and Italian; Supplementary Material) that detailed key historical milestones and the geographic locations associated with the violin's construction, performance, and repairs over the centuries, while also elucidating its connection to the Montecastelli village. The story map encapsulated (i) the history of the violin, (ii) the cultural, temporal, and spatial context in which it was created and remains in use, and (iii) an overview of its components, enriched with relevant textual and multimedia annotations supplemented by the audio excerpts from the interview.

The story map comprised a total of 11 events, of which eight featured 3D models. The first story event was a Map-type event introducing master Stradivari, with a reference polygon over the city of Cremona (Fig. 3-b). This event included several entities linked to Wikidata entries relating to the artists and locations mentioned in the accompanying text, such as Niccolò Paganini, Giovanni Battista Viotti, Antonio Stradivari, Cremona, France, and the United Kingdom. The subsequent event was a Slide-type event that introduced the violin, along with a Wikidata reference to the Dresden Staatskapelle, which currently houses it, as well as a video depicting a performance featuring the violin in the chapel (Fig. 4-a). Events 3 through 10 were 3D-type events, which included an overview of the violin (Fig. 3-a) and explanations of its distinctive components (Fig. 4-b). The text incorporated insights provided by master Bonhoeffer, who elucidated the individual parts and associated facts, complemented by multimedia links to listen to the luthier's commentary. Wikidata entities associated with each event represented key concepts mentioned in the descriptions. The story alternated between the textured model and the structural scan, maintaining the same pose to create the impression of a sudden transition from a textured to a structural perspective. This alternation allowed visitors to engage interactively with the two models, addressing both general stakeholders (who might navigate the events following the proposed narrative) and luthiers (who may wish to examine details of the instrument). Collectively, the events provided all visitors with comprehensive insights into the components of the violin, including the scroll, the f-holes, the sound post, and even Stradivari's signature located within the internal portion of the sound box (Fig. 4-b). The story concluded with a Slide-type event that described the violin-making workshop in Montecastelli Pisano and its connection to the original and replicated violins (Fig. 3-c).

The evaluation was structured into two distinct phases addressing different research objectives based on the presented case study. The first phase (*story map creation*, Section 3.7) focused on the usability and technical effectiveness of the story creation environment and involved evaluators with academic, technical backgrounds. The second phase (*story map visualisation*) evaluated narrative visualisation and exploration and included participants with varying degrees of domain expertise, including instrument makers, musicologists, and users with knowledge of ancient music. This distinction reflected the different competencies required to assess authoring workflows versus narrative consumption and interpretation.

3.7 Evaluating the story map creation

The evaluation presented in this section aimed to assess the usability of the case study story map and its effectiveness in fulfilling the objectives of the ViolTwin initiative.

As the initial evaluation phase, we focused on the user interaction experience with the story map creation tool. Specifically, we assessed the challenges encountered and feedback provided by non-expert users when building a story map related to a specified topic. The insights gathered were crucial in identifying areas for enhancement in the interface design.

Five doctoral students from the University of Pisa (Italy) were recruited to construct story maps incorporating the three types of events delineated in the preceding sections. These candidates possessed diverse academic backgrounds, from Computer Science to Digital Humanities. Specifically, two hold degrees in Digital Humanities, one in Computer Engineering, and two in Cultural Heritage. Such varied academic profiles represent potential user groups, including humanities scholars, museum practitioners, cultural heritage specialists, and engineers focused on applications within cultural heritage.

We identified four primary tasks that each student needed to complete:

1. Creating a 3D-type event in which the sound post of the violin (an internal component of the instrument) was shown;
2. Creating a presentation about Antonio Stradivari through a Map-type event;
3. Creating a Slide-type event that showcased the Montecastelli violin workshop;
4. Publishing a comprehensive story map that incorporated the aforementioned events.

The tasks were designed to be concise and specific to enhance the granularity of the evaluation. The rationale behind the task definitions was to encourage users to recreate pivotal events in the history of the 1737 Stradivari violin central to the case study. We communicated this rationale to the students to elucidate the purpose behind their actions, fostering more authentic and insightful interactions. Notably, all students had undergone pre-training on utilising the tool. Furthermore, they were provided with comprehensive texts and URLs for populating the entity forms.

Five subtasks were associated with the four main tasks:

1. The event creation tasks necessitated that users added event titles, textual descriptions, and relevant Wikidata entities (selected from those suggested by the tool);
2. For the 3D-type event, users were required to input the Sketchfab identifier of the 3D scene representing the Stradivari violin;
3. For the Map-type event, users had to incorporate an image through an external URL and delineate a polygon around Stradivari's birthplace in Cremona;
4. For the Slide-type event, users were to insert a relevant background image via an external URL and include the web site "Il Poggio Montecastelli" as a digital object;
5. During the publication step, users were tasked with previewing the story map and using the designated buttons to create and publish the associated KB.

After the story creation sessions, we systematically collected data regarding user interactions throughout the principal tasks and subtasks. We evaluated the following performance metrics: (i) the percentage of subtasks completed correctly, (ii) the time taken to finalise each subtask, and (iii) the percentage of errors committed during subtask execution. Based on our observations of the users' ease of interaction with the interface, and given that users were equipped with the necessary materials to fulfil the entry forms, we classified a subtask as *failed* if not completed within one minute, and as *successfully completed* otherwise.

3.8 Evaluating the story map visualisation

As the second evaluation step, we concentrated on the user interaction experience with the case study story map. In this instance, the users were not story creators and represented a heterogeneous group, which included expert instrument makers as well as individuals with varying levels of interest in the subject matter.

To simulate this diversity, we invited a cohort of 20 participants spanning different ages and competencies to interact with the story. The participants included computer scientists and humanists, all with varying degrees of interest in violin making and chamber music; specifically, 55% were aged 18–30 years, 35% were aged 31–50 years, and 10% were aged 51–65 years. The majority of participants possessed a background in Digital Humanities (45%) or Computer Science (35%), while a smaller proportion had backgrounds in Humanities (10%) or other fields (10%). Furthermore, 15% identified as instrument makers or chamber music enthusiasts. A significant majority (90%) accessed the story map via a computer, whereas 10% used smartphones.

Participants interacted with the Stradivari violin story map and subsequently completed an anonymous survey created using Google Forms to collect feedback on specific evaluation aspects. The survey began by gathering general information from participants, including age, educational background, and device usage. Additionally, it inquired whether participants were instrument makers and their level of interest in chamber music. This question aimed to identify respondents with specialised knowledge or interest in the subject, given that their insights could provide deeper and more valuable feedback regarding the visualisation's effectiveness.

Following this part, the form collected quantitative data regarding the quality of interaction with the story map by requesting participants to rate various aspects on a scale from 1 to 5: (i) The initial question evaluated the overall usability of the story map, specifically whether navigating and visualising the events within the map was intuitive or convoluted. Subsequently, three questions focused on 3D model interaction, whereby participants were asked to rate (ii) the ease of interaction with the 3D model, (iii) the usefulness of navigating across model annotations to enhance story comprehension, and (iv) the effectiveness of the 3D models in improving knowledge about the violin. The survey then progressed to evaluate other event types by requesting ratings on (v) the usefulness of background images (in Slide-type events) to aid in event-content understanding, and (vi) the usefulness of displaying a polygon on a geographic map (in the Map-type event) to add geospatial context to an event. Lastly, the survey allowed users to submit a free-text response to justify their ratings.

The overall numeric assessments from the six evaluation assets served to assess the quality of the user experience, and consequently, the effectiveness of our story map-building approach in illustrating the historical context of the violin. The free-text responses were subsequently analysed to elucidate possible reasons for any weaknesses identified.

4 Results

4.1 Knowledge Base

The knowledge base associated with the case study comprised a total of 1716 triples. For the present evaluation, we focus on the KB corresponding to a 3D-event type representing the "Scroll" component of the violin in our case study (Fig. 5). This graph was derived from an event form compiled via the story map building interface, wherein the event creator specified a title, a description, one related Wikipedia entity (specifically, "Violin"), an associated audio content, in addition to a corresponding 3D model and annotation sourced from Sketchfab ("Scroll").

The knowledge graph was formally consistent and interoperable with other knowledge bases due to its reuse of publicly available reference ontological classes and properties. Moreover, we verified that it was publicly accessible. This characteristic ensured compliance with the Linked Open Data paradigm and fostered interoperability with large knowledge networks [86].

Within the graph, the "Stradivarius 3D model", represented as a "Visual Item", was linked to its annotation ("The Scroll", another "Visual Item") via the P106 property "is composed of". The model and the annotation were connected to a narrative "Event" through the P67 property "refers to". This same event was further linked to a textual description of the "Scroll" annotation using the "hasDescription" property. An MP3 audio file featuring master luthier Bonhoeffer elucidating details about the "Scroll" was represented as an E73 "Information Object" and connected to the event through the P67 property. The Wikidata entity was represented as an E19 "Physical Object" and linked to the event via the "hasWikidataEntity" property.

4.2 Story map creation

Evaluating the user interaction experience with the story map creation tool underscored its generally *good* usability (Table 1). All participants completed the macro tasks successfully, achieving a 100% completion rate. All subtasks were executed correctly except for two subtasks: The subtask about "adding an image using an external link" was accurately performed by only 2 out of 5 users. The predominant error identified was that users instinctively opted to use the button for image upload (after downloading the image) rather than directly copying and pasting the external URL link.

The subtask about "adding the Il Poggio Montecastelli web site as a digital object" proved to be more critical, with only 1 out of 5 users executing it correctly. The main error observed was that most users failed to press the confirmation button necessary to accept the insertion of the digital object.

Errors were due to the fact that users took more than one minute to complete the tasks. However, these errors did not impede the overall completion of the macro tasks, as all users were able to autonomously correct their mistakes after a few attempts.

In summary, this evaluation illustrates that the tasks were largely immediate and intuitive, given that only 2 out of 16 subtasks (12.5%) exhibited errors. Furthermore, the nature of these errors provides clear direction for future interventions to enhance the usability of our tool.

4.3 Story map visualisation

The story map visualization assessment, conducted over the six evaluation assets delineated in Section 3.8, highlighted strengths and weaknesses of our story map visualisation (Fig. 6).

Regarding the overall usability of the story map (Fig. 6-a), 55% (11 out of 20 participants) assigned the maximum rating of 5, 35% (7 users) assigned a very high rating of 4, and 10% (2 users) assigned a medium rating of 3. Notably, no participants assigned lower ratings, indicating a generally positive perception of the story map's intuitiveness and ease of navigation through the events.

Regarding the ease of interaction with the 3D models (Fig. 6-b), 50% (10 users) assigned a 5-rating, 35% (7 users) assigned a 4-rating, 10% (2 users) assigned a 3-rating, and 1 user assigned a 2-rating. Although most users rated the interaction highly, finding it sufficiently intuitive, one participant expressed confusion due to the ability to orbit around the model, which could lead to a distraction from the main subject of the annotation.

Regarding the usefulness of navigating across the 3D model annotations to enhance story comprehension (Fig. 6-c), 65% (13 users) assigned a rating of 5, while 15% (3 users) assigned a rating of 4, another 15% (3 users) a rating of 3, and 1 user a rating of 2. The single dissenting participant preferred a more static focus on each annotation and judged the 3D visualisation marginally superior to a static image.

Regarding the usefulness of the 3D models for enhancing the understanding of the Stradivari violin (Fig. 6-d), 70% (14 users) assigned a rating of 5, 20% assigned a rating of 4, and 10% a rating of 3. These findings indicate that the 3D visualisations were overwhelmingly recognised as providing an advantage over static images in comprehending the instrument.

Regarding the usefulness of background images to support content understanding (Fig. 6-e), 45% (9 users) assigned a rating of 5, 40% (8 users) assigned a rating of 4, while 3 users assigned ratings of 3, 2, and 1, respectively. This scenario reflects a consensus on the utility of background images, although some users expressed a preference for the quality and impact of 3D models.

Finally, concerning the usefulness of the polygon overlayed on the geographic map for geospatial contextualisation (Fig. 6-f), 75% (15 users) assigned a rating of 5, 15% (3 users) assigned a rating of 4, and 10% (2 users) assigned a rating of 3. Thus, the geographical information was acknowledged as an added value, yielding a high consensus among participants.

As for the free-text responses, most participants emphasised that the opportunity to explore the models in detail enriched their learning experience. Some users suggested the inclusion of additional annotations on the 3D models to provide supplementary context and information. Furthermore, participants generally found the background images beneficial in supporting the understanding of the event content; however, some recommended enhancing the contrast between the text and background images or adjusting text placement to improve readability. All participants highlighted the utility of visualising the contextual polygon on the geographic map, noting that it complemented the narrative and bolstered their spatial understanding. Additionally, several users advocated for incorporating multiple geometries, points, and polygons within the events.

5 Discussion and conclusion

Our results demonstrate that users generally appreciated the added value of 3D models and geospatial information in the case-study story map. This indicates that integrating such information into narratives improved the accessibility and interpretability of static visuals, providing more immersive experiences and enhancing user engagement, consistent with findings from other studies [87, 88]. The narrative presented as a case study enriched the traditional story map by bridging the gap between 3D representations, the spatiotemporal unfolding of the narrative, the broader context, and human perception. Compared to alternative methods, we have shown that our approach could effectively manage 3D model inspection, map visualisation, and traditional events to establish a novel type of story map. Using a KB as the backend, based on widely used ontologies, was a further distinguishing feature. This KB facilitated the connection of heterogeneous events in our stories to broader contexts, which may be advantageous for cultural and touristic applications.

In summary, the story map presented (i) allowed users to appreciate the details of a rare ancient instrument, (ii) offered auditory experiences through audio-video samples, (iii) enabled the inspection of the violin components at varied levels of competency (from overviews to structural details), and (iv) improved the understanding of the instrument's spatiotemporal pathways and cultural context. While this experience cannot replace direct, live interaction and auditory engagement, it provided a more comprehensive experience than traditional indirect methods (e.g., through technical sheets and static image sequences).

Another significant point is that the backend KB allows for the interconnection of different narratives created by various authors, potentially leading to the discovery of new relationships among events and objects [29, 30]. This aspect holds substantial potential for uncovering relationships over time and across geographies among the life histories of different instruments. Finally, releasing our tool as an open-source and freely accessible web service empowers users to overcome the limitations of commercial software. This is particularly advantageous for public research and non-profit organisations.

Despite the positive usability feedback from the creators and visualisers of our story map, several areas for improvement were identified. These aspects will be the focus of our future enhancements and form general best practices for developing subsequent narratives. Regarding story creation, we will explore usability improvements for digital objects and image importing. In terms of story visualisation, we observed that increasing the number of annotations and events for the 3D models would enhance user experience in two primary ways: first, it would provide additional technical details for instrument makers; second, it would offer less interactive users the opportunity to examine various parts and views of the model. Finally, while the geospatial representation was deemed effective overall, it will require incorporating multiple geometry management tools within the story map building interface.

As a subsequent development to showcase and evaluate our tool, we narrated the stories of other instruments from the Associazione il Poggio Montecastelli (Fig. 7 and Supplementary Material). Specifically, we detailed a "Viola" crafted by master luthier Gasparo da Salò in 1580, and a "Viola da Gamba" created by master Francesco Ventura di Linarol in 1585. The insights and feedback received from the luthiers regarding these narratives enlightened us on their practical utility for instrument reproduction, in addition to serving as a means of preserving cultural heritage. For instance, the "Viola da Gamba" has recently disappeared due to legal sequestration [89], making our story the only means to appreciate this ancient masterpiece structurally and culturally. This narrative has been instrumental in assisting the luthiers from il Poggio in recreating a copy. As an ancillary story, we assembled a contextual story map of the Montecastelli village itself (Supplementary Material), featuring touristic points of interest elaborated through story events and enriched with links to Wikidata/Wikipedia pages for additional information. Ultimately, we compiled the three ancient instrument narratives along with the story of Montecastelli into a website to promote the cultural heritage of the territory. These examples illustrate the potential our methodology holds for other museums and workshops - especially those equipped with scanning tools - to reach broader audiences.

Although the case study presented in this paper focuses on a single class of artefacts (violins), the scalability of our tool across heterogeneous cultural heritage domains is primarily ensured by the design of the Narrative Ontology and its separation of narrative modelling from domain knowledge representation. This ontology is intentionally domain-agnostic and models general narrative primitives - such as events, agents, temporal relations, and places - rather than domain-specific semantics. As a result, the core narrative structure remains stable when applied to different artefact types or even to complex heritage contexts such as sites and landscapes. Domain-specific concepts introduced by new artefact classes are not embedded into the narrative model itself, but are instead represented through external domain ontologies and linked to the Narrative Ontology via alignment and semantic linking strategies. This approach preserves interoperability while minimising the need for structural changes to the narrative framework. The conceptual scalability of this design has been validated through the reuse of the Narrative Ontology in several projects spanning markedly different knowledge domains, including mountain value chains in the Horizon 2020 project MOVING (bioeconomy) [28, 29], medieval travel narratives in the Italian National Research Project IMAGO (Latin philology) [90], and CNR research on ecosystem-based models for predicting the Giant Squid diffusion (computational biology) [91]. Therefore, the presented tool can accommodate diverse narrative domains while maintaining a consistent narrative backbone.

As a final note, it is essential to emphasise that our methodology only partially satisfies the criteria of a digital twin for a three-dimensional musical instrument. Specifically, it primarily addresses the historical and cultural context of the instrument alongside its three-dimensional visualisation, while excluding its mechanical and acoustic characteristics. Consequently, our approach should be regarded as a complementary component intended to be integrated with a comprehensive simulation of the instrument. Such a simulation can be embedded within a narrative framework to enable users to engage with sound production processes or specialists to analyse sound generation mechanisms. A story map can indeed incorporate samples and outputs from simulations to represent the internal functioning of the instrument succinctly. The narrative environment should therefore serve as an adjunct to external simulation tools designed for virtual sound reproduction [92, 93], specifically addressing expert luthiers aiming to replicate sound mechanics in other instruments. An integration between our tool and a mechanical-acoustic model is underway in the PlaceMUS XR European project [94], through the creation of seven musical pathways that will embed the physical and acoustic characteristics of ancient instruments into narratives to produce richer and more immersive representations of musical heritage.

One interesting potential application case is the description of ancient instruments, whose history and iconography are intertwined with their functioning, materials, and craftsmanship, which should be studied and represented together within a unified view [95–97]. The combination of story maps and simulation tools holds broad implications for cultural heritage preservation and museums because it facilitates explanations of the instrument for diverse audiences, ranging from general visitors to expert practitioners [98, 99]. Moreover, our tool offers significant potential benefits for specialist users such as musicologists and experts in historical luthiery. The use of semantically grounded entities, explicit provenance links, and ontology-aligned knowledge graphs enables fine-grained representations of instruments, makers, materials, and historical contexts. This structure supports advanced scholarly activities, including cross-source comparison, temporal reasoning, and the exploration of alternative interpretative narratives. Furthermore, the separation between story creation and semantic publication allows domain experts to engage with enriched narrative outputs without being exposed to the technical complexity of semantic modelling, while still benefiting from its expressive power and interoperability.

Acknowledgments

This paper is the outcome of a non-profit, zero-budget collaboration within the ViolTwin initiative, a scientific partnership between CNR-ISTI and the Associazione Il Poggio Montecastelli (protocol number 417591/2023, dated 27/12/2023). The authors would like to express their gratitude to everyone who participated in the user tests, as well as to Master Philipp Bonhoeffer for his valuable input.

Supplementary Material

The Supplementary Material is attached to the present publication as an additional file. It contains links to (i) the new 3D SMBVT software developed, (ii) the free-to-use service instance to build story maps, (iii) the models on Sketchfab, (iv) the ViolTwin initiative, (v) the link to the SPARQL to query the knowledge base, and (vi) all story maps we produced for historical musical instruments.

References

[1] Piccialli, F.: The internet of things supporting the cultural heritage domain: analysis, design and implementation of a smart framework enhancing the smartness of cultural spaces. Universita degli Studi di Napoli Federico II (2016)

[2] Nilson, T., Thorell, K.: Cultural heritage preservation: The past, the present and the future. Halmstad University Press (2018)

[3] Ranjgar, B., Sadeghi-Niaraki, A., Shakeri, M., Rahimi, F., Choi, S.-M.: Cultural heritage information retrieval: Past, present and future trends. IEEE Access (2024)

[4] Catalano, C.E., Vassallo, V., Hermon, S., Spagnuolo, M.: Representing quantitative documentation of 3d cultural heritage artefacts with cidoc crmdig. International Journal on Digital Libraries **21**(4), 359–373 (2020)

[5] Malik, U., Tissen, L.N., Vermeeren, A.P.: 3d reproductions of cultural heritage artifacts: evaluation of significance and experience. Studies in digital heritage **5**(1), 1–29 (2021)

[6] Pietroni, E., Ferdani, D.: Virtual restoration and virtual reconstruction in cultural

heritage: Terminology, methodologies, visual representation techniques and cognitive models. Information **12**(4), 167 (2021)

[7] Basu, A., Paul, S., Ghosh, S., Das, S., Chanda, B., Bhagvati, C., Snasel, V.: Digital restoration of cultural heritage with data-driven computing: A survey. IEEE Access **11**, 53939–53977 (2023)

[8] Sachs, C.: The History of Musical Instruments. Courier Corporation, Garden City, NY, US (2012)

[9] Bates, E.: The social life of musical instruments. Ethnomusicology **56**(3), 363–395 (2012)

[10] Bellia, A., Pavone, D.P.: Computed tomography and handcrafting processes of an ancient musical instrument: The aulos from poseidonia. Archeologia e Calcolatori **1** (2021)

[11] Turchet, L., Fischione, C., Essl, G., Keller, D., Barthet, M.: Internet of musical things: Vision and challenges. IEEE Access **6**, 61994–62017 (2018) https://doi.org/10.1109/ACCESS.2018.2872625

[12] Howard, K.: Musical instruments as tangible cultural heritage and as/for intangible cultural heritage. International Journal of Cultural Property **29**(1), 23–44 (2022)

[13] Fang, Y., Mitoulis, S.-A., Boddice, D., Yu, J., Ninic, J.: Scan-to-bim-to-sim: Automated reconstruction of digital and simulation models from point clouds with applications on bridges. Results in Engineering **25**, 104289 (2025) https://doi.org/10.1016/j.rineng.2025.104289

[14] Nezzi, C., Gufler, V., Vidoni, R., Rauch, E.: Kinematic and dynamic modeling of mechanical systems towards digital twins. Results in Engineering **26**, 104874 (2025) https://doi.org/10.1016/j.rineng.2025.104874

[15] Bernardini, F., Tuniz, C., Zanini, F.: X-ray computed microtomography for paleoanthropology, archaeology, and cultural heritage. In: Nanotechnologies and Nanomaterials for Diagnostic, Conservation and Restoration of Cultural Heritage, pp. 25–45. Elsevier, Amsterdam (2019)

[16] Bardelli, S., Coro, G., Scaramuzzo, R.T., Ciantelli, M., Cuttano, A.: An open data collection of 3d tool and equipment models for neonatology. Results in Engineering **25**, 104236 (2025) https://doi.org/10.1016/j.rineng.2025.104236

[17] Rehan, M., He, T., Tahir, D., Khalil, A.K., Yip, W.S., To, S.S.: 3d finite element modeling and experimental investigation of micro milling of laser powder bed fusion fabricated ti6al4v. Results in Engineering **26**, 104635 (2025) https://doi.org/10.1016/j.rineng.2025.104635

[18] Khan, N.H.R., Kumar, S.V.: Terrestrial lidar derived 3d point cloud model, digital elevation model (dem) and hillshade map for identification and evaluation of pavement distresses. Results in Engineering **23**, 102680 (2024) https://doi.org/10.1016/j.rineng.2024.102680

[19] Niccolucci, F., Felicetti, A., Hermon, S.: Populating the data space for cultural heritage with heritage digital twins. Data **7**(8) (2022) https://doi.org/10.3390/data7080105

[20] Zou, C., Rhee, S.-Y., He, L., Chen, D., Yang, X.: Sounds of history: A digital twin approach to musical heritage preservation in virtual museums. Electronics **13**(12) (2024) https://doi.org/10.3390/electronics13122388

[21] Cairoli, M., Tagliabue, L.C.: Digital twin for acoustics and stage craft facility management in a multipurpose hall. Acoustics **5**(4), 909–927 (2023) https://doi.org/10.3390/acoustics5040053

[22] Turchet, L.: Musical metaverse: vision, opportunities, and challenges. Personal and Ubiquitous Computing **27**(5), 1811–1827 (2023)

[23] Takada, T., Nakazato, J., Orsholits, A., Tsukada, M., Ochiai, H., Esaki, H.: Design of

digital twin architecture for 3d audio representation in ar. In: 2024 IEEE International Conference on Metaverse Computing, Networking, and Applications (MetaCom), pp. 222–230 (2024). https://doi.org/10.1109/MetaCom62920.2024.00044

[24] Bartalesi, V., Coro, G., Lenzi, E., Pratelli, N., Pagano, P., Moretti, M., Brunori, G.: A semantic knowledge graph of european mountain value chains. Scientific Data **11**(1), 978 (2024)

[25] Korzybski, A.: Science and Sanity: An Introduction to non-Aristotelian Systems and General Semantics. Institute of GS, n.p. (1958)

[26] MacEachren, A.M.: How Maps Work: Representation, Visualization, and Design. Guilford Press, New York (2004)

[27] Roth, R.E.: Cartographic design as visual storytelling: synthesis and review of map-based narratives, genres, and tropes. The Cartographic Journal **58**(1), 83–114 (2021)

[28] Bartalesi, V., Coro, G., Lenzi, E., Pratelli, N., Pagano, P.: Towards digital twins of territories through semantic story maps. In: BUILD-IT 2023 Workshop-BUILding a Digital Twin: Requirements, Methods, and Applications, Rome, Italy, pp. 41–45 (2023)

[29] Bartalesi, V., Coro, G., Lenzi, E., Pagano, P., Pratelli, N.: From unstructured texts to semantic story maps. International Journal of Digital Earth **16**(1), 234–250 (2023)

[30] Bartalesi, V., Coro, G., Lenzi, E., Pratelli, N., Pagano, P., Felici, F., Moretti, M., Brunori, G.: Using semantic story maps to describe a territory beyond its map. Semantic web **14**(6), 1255–1272 (2023) https://doi.org/10.3233/SW-233485

[31] Wilkinson, M.D., Dumontier, M., Aalbersberg, I.J., Appleton, G., Axton, M., Baak, A., Blomberg, N., Boiten, J.-W., Silva Santos, L.B., Bourne, P.E., *et al.*: The fair guiding principles for scientific data management and stewardship. Scientific data **3**(1), 1–9 (2016)

[32] Mai, G., Huang, W., Cai, L., Zhu, R., Lao, N.: Narrative cartography with knowledge graphs. Journal of Geovisualization and Spatial Analysis **6**(1), 4 (2022)

[33] Ranade, P., Dey, S., Joshi, A., Finin, T.: Computational understanding of narratives: A survey. IEEE Access **10**, 101575–101594 (2022) https://doi.org/10.1109/ACCESS.2022.3205314

[34] Seneviratne, O., Hendler, J. (eds.): Linking the World's Information: Essays on Tim Berners-Lee's Invention of the World Wide Web vol. 52, 1st edn. Association for Computing Machinery, New York, NY, USA (2023)

[35] Malakan, Z.M., Anwar, S., Hassan, G.M., Mian, A.: Sequential vision to language as story: A storytelling dataset and benchmarking. IEEE Access **11**, 70805–70818 (2023) https://doi.org/10.1109/ACCESS.2023.3293646

[36] Bartalesi, V., Lenzi, E., Pratelli, N.: A web tool to create and visualise semantic story maps. In: Text2Story at ECIR, pp. 163–169 (2023)

[37] Schnürer, R.: Storytelling with Animated Interactive Objects in Real-time 3D Maps. ETH Zurich, Switzerland (2023). https://doi.org/10.3929/ethz-b-000663385

[38] Wang, X.: 3D Interactive Spatio-temporal and Semantic Visualization for Discrete Data Associated with Multiple Categories and Long Historical Events. Technical University of Munich, Munich (2021)

[39] Li, J., Chen, S., Chen, W., Andrienko, G., Andrienko, N.: Semantics-space-time cube: A conceptual framework for systematic analysis of texts in space and time. IEEE transactions on visualization and computer graphics **26**(4), 1789–1806 (2018)

[40] Antoniou, V., Nomikou, P., Konstantinos, P., Karatzaferis, O., Vlasopoulos, O., Stentoumis, C., Kalisperakis, I.: A Journey to Salamis Island (Greece) using a GIS Tailored Interactive Story Map Application. In:

7th International Conference on Geographical Information Systems Theory, Applications and Management, pp. 187–194 (2020). https://doi.org/10.5220/0010440701870194

[41] Antoniou, V., Ragia, L., Nomikou, P., Bardouli, P., Lampridou, D., Ioannou, T., Kalisperakis, I., Stentoumis, C.: Creating a story map using geographic information systems to explore geomorphology and history of methana peninsula. International Journal of Geo-Information (2018) https://doi.org/10.3390/ijgi7120484

[42] Walshe, N.: Using arcgis online story maps. Teaching Geography **41**(3), 115–117 (2016)

[43] Bajjali, W.: Working with arcgis online and storymap app. In: ArcGIS Pro and ArcGIS Online: Applications in Water and Environmental Sciences, pp. 357–392. Springer, Gewerbestrasse, Cham (2023)

[44] ArcGIS: ArcGIS CityEngine, a modern 3D city design solution. https://www.esri.com/en-us/arcgis/products/arcgis-cityengine/overview (2025)

[45] Sugianto, E., Hosea, J.F., Jabar, B.A., Irwansyah, E., Fitrianah, D.: 3d modelling building of district johar baru using arcgis pro and cityengine. Procedia Computer Science **227**, 623–631 (2023)

[46] Naul, B., Walt, S., Crellin-Quick, A., Bloom, J.S., Pérez, F.: cesium: Open-source platform for time-series inference. arXiv preprint arXiv:1609.04504 (2016)

[47] The Cesium Platform: CesiumJS. https://cesium.com/platform/cesiumjs/ (2025)

[48] Google Inc.: 3D storytelling getting started. https://developers.google.com/maps/architecture/3d-storytelling-getting-started?hl=en (2025)

[49] Matterport Discover: Matterport Virtual Tours. https://discover.matterport.com/tag/maps (2025)

[50] Hariffin, N.A.B., Razali, M.H.B., Luh, L.C., Sulaiman, S.A., *et al.*: The suitability of matterport for building parcel dimension survey. In: 2023 IEEE 13th International Conference on System Engineering and Technology (ICSET), pp. 239–244 (2023). IEEE

[51] The Omeka S Web platform: Omeka S. https://omeka.org/s/ (2025)

[52] Zanoni, M., Setragno, F., Sarti, A., *et al.*: The violin ontology. In: Proc. of the 9th Conference on Interdisciplinary Musicology (CIM14), p. 6 (2014). CIM14

[53] Kolozali, S., Barthet, M., Fazekas, G., Sandler, M.B.: Knowledge representation issues in musical instrument ontology design. In: ISMIR, pp. 465–470 (2011)

[54] Achichi, M., Bailly, R., Cecconi, C., Destandau, M., Todorov, K., Troncy, R.: Doremus: Doing reusable musical data. In: ISWC: International Semantic Web Conference, pp. 1–5 (2015)

[55] Berardinis, J., Carriero, V.A., Jain, N., Lazzari, N., Meroño-Peñuela, A., Poltronieri, A., Presutti, V.: The polifonia ontology network: Building a semantic backbone for musical heritage. In: International Semantic Web Conference, pp. 302–322 (2023). Springer

[56] Renda, G., Daquino, M., Presutti, V.: Melody: A platform for linked open data visualisation and curated storytelling. In: Proceedings of HT 2023, pp. 1–8. Association for Computing Machinery, New York, NY, USA (2023). https://doi.org/10.1145/3603163.3609035

[57] Meghini, C., Bartalesi, V., Metilli, D.: Representing narratives in digital libraries: The narrative ontology. Semantic Web **12**(2), 241–264 (2021)

[58] Bartalesi, V., Pratelli, N.: Representing geospatial knowledge in narratives. J. Comput. Cult. Herit. (2025) https://doi.org/10.1145/3703918

[59] Bartalesi, V., Pratelli, N.: Representing geospatial knowledge in narratives. ACM

Journal on Computing and Cultural Heritage (2024)

[60] Lenzi, E., Bartalesi, V., Pratelli, N.: Story Map Building and Visualising Tool software repository. https://github.com/EmanueleLenzi92/SMBVT (2025)

[61] Vrandečić, D., Pintscher, L., Krötzsch, M.: Wikidata: The making of. In: Companion Proceedings of the ACM Web Conference 2023, pp. 615–624 (2023)

[62] Europeana: European digital cultural heritage. https://www.europeana.eu (2025)

[63] Macrì, E., Cristofaro, C.L.: The digitalisation of cultural heritage for sustainable development: The impact of europeana. Cultural Initiatives for Sustainable Development: Management, Participation and Entrepreneurship in the Cultural and Creative Sector, 373–400 (2021)

[64] Silva, A.L., Terra, A.L.: Cultural heritage on the semantic web: The europeana data model. IFLA journal **50**(1), 93–107 (2024)

[65] Sketchfab: Sketchfab - The best 3D viewer on the web. https://sketchfab.com (2025)

[66] Wikidata: The Wikidata Query Service. https://query.wikidata.org/ (2025)

[67] Bielefeldt, A., Gonsior, J., Krötzsch, M.: Practical linked data access via sparql: The case of wikidata. In: LDOW@ WWW, pp. 1–10 (2018)

[68] Open Geospatial Consortium: Well-known text representation of coordinate reference systems spec. version 2.0.6. https://docs.ogc.org/is/18-010r7/18-010r7.html (2025)

[69] The Apache Software Foundation: Apache Tomcat. https://tomcat.apache.org/ (2025)

[70] ECMA-404: JavaScript Object Notation. https://www.json.org/json-en.html (2025)

[71] The PostgreSQL Global Development Group: PostgreSQL: The World's Most Advanced Open Source Relational Database. https://www.postgresql.org/ (2025)

[72] CNR-ISTI (AI and Digital Humanities group at AIMH lab): SMBVT Java triplifier software. https://github.com/AIMH-DHgroup/SMBVT-Triplifier (2025)

[73] McGuinness, D.L.: Owl web ontology language overview. W3C Member Submission (2004)

[74] The Apache Software Foundation: Apache Jena Fuseki. https://jena.apache.org/documentation/fuseki2/ (2025)

[75] Sketchfab: ViewerAPI. https://sketchfab.com/developers/viewer (2025)

[76] Bartalesi, V., Coro, G., Lenzi, E., Pagano, P., Pratelli, N.: From unstructured texts to semantic story maps. International Journal of Digital Earth **16**(1), 234–250 (2023) https://doi.org/10.1080/17538947.2023.2168774 https://doi.org/10.1080/17538947.2023.2168774

[77] Knight Lab: StoryMapJS: Maps that tell stories. storymap.knightlab.com/ (2025)

[78] Doerr, M.: The cidoc conceptual reference module: an ontological approach to semantic interoperability of metadata. AI magazine **24**(3), 75–75 (2003)

[79] W3C: RDF specifications. https://www.w3.org/TR/PR-rdf-syntax/Overview.html (2025)

[80] Manola, F., Miller, E., McBride, B., *et al.*: Rdf primer. W3C recommendation **10**(1-107), 6 (2004)

[81] Apache Software Foundation: Apache Jena Fuseki. https://jena.apache.org/documentation/fuseki2/ (2024)

[82] DuCharme, B.: Learning SPARQL: Querying and Updating with SPARQL 1.1. " O'Reilly Media, Inc.", California (2013)

[83] Pérez, J., Arenas, M., Gutierrez, C.: Semantics and complexity of sparql. In: International Semantic Web Conference, pp. 30–43

(2006). Springer

[84] Associazione Culturale Il Poggio Montecastelli: Science and Music in Montecastelli Pisano. https://www.scienceandmusic.com/ (2024)

[85] Coro, G., Masetti, G., Bonhoeffer, P., Betcher, M.: Distinguishing violinists and pianists based on their brain signals. In: Artificial Neural Networks and Machine Learning–ICANN 2019: Theoretical Neural Computation: 28th International Conference on Artificial Neural Networks, Munich, Germany, September 17–19, 2019, Proceedings, Part I 28, pp. 123–137 (2019). Springer

[86] Thanos, C., Meghini, C., Bartalesi, V., Coro, G.: An exploratory approach to data driven knowledge creation. Journal of Big Data **10**(1), 29 (2023)

[87] Harris, T.M., Bergeron, S., Rouse, L.J.: Humanities gis: Place, spatial storytelling, and immersive visualization in the humanities. In: GeoHumanities, pp. 226–240. Routledge, New York (2011)

[88] Thöny, M., Schnürer, R., Sieber, R., Hurni, L., Pajarola, R.: Storytelling in interactive 3d geographic visualization systems. ISPRS International Journal of Geo-Information **7**(3), 123 (2018)

[89] Ministero della Cultura Italiano: I Carabinieri del Nucleo per la Tutela del Patrimonio Culturale di Bologna bloccano una collezione di strumenti antichi alla frontiera. https://cultura.gov.it/comunicato/i-carabinieri-del-nucleo-per-la-tutela-del-patrimonio-culturale-di-bologna-bloccano-una-collezione-di-strumenti-antichi-alla-frontiera (2021)

[90] Bartalesi, V., Coro, G., Lenzi, E., Pratelli, N.: A web tool to create and visualise semantic story maps. In: Proceedings of Text2Story — Sixth Workshop on Narrative Extraction From Texts (Text2Story 2023), Held in Conjunction with the 45th European Conference on Information Retrieval (ECIR 2023), vol. 3370, pp. 163–169. CEUR Workshop Proceedings, Dublin, Ireland (2023)

[91] Coro, G., Magliozzi, C., Ellenbroek, A., Pagano, P.: Improving data quality to build a robust distribution model for architeuthis dux. Ecological Modelling **305**, 29–39 (2015) https://doi.org/10.1016/j.ecolmodel.2015.03.011

[92] Camargo Caicedo, Y., Montoya Arango, J.A., Acosta Acosta, R.: Study on compressive strength of bamboo (guadua amplexifolia j. presl) in san jorge river basin, colombia. Results in Engineering **24**, 102926 (2024) https://doi.org/10.1016/j.rineng.2024.102926

[93] Mia, R., Bakar, M.A., Islam, M.R., Ahmed, T.: Eco-friendly coloration from mahogany wood waste for sustainable dyeing of organic nonwoven cotton fabric. Results in Engineering **17**, 101032 (2023) https://doi.org/10.1016/j.rineng.2023.101032

[94] CNR-ISPC: The PlaceMUS XR European Project. https://www.ispc.cnr.it/en/2025/10/30/placemus-xr/ (2025)

[95] Avanzini, F., CANAZZA TARGON, S., De Poli, G., Fantozzi, C., Pretto, N., Roda, A., Angelini, I., Bettineschi, C., Deotto, G., Faresin, E., *et al.*: Archaeology and virtual acoustics. a pan flute from ancient egypt. In: Proc. of the 12th Int. Conference on Sound and Music Computing, pp. 31–36 (2015)

[96] Bucur, V.: Description of string instruments for classical music. In: Handbook of Materials for String Musical Instruments, pp. 5–91. Springer, Cham, Switzerland (2016)

[97] Nadirova, M., Aliyeva, H.: Evolution of musical instruments: From ancient times to modern innovations. AGRICULTURAL SCIENCES **135**, 12 (2024)

[98] Cui, L., Wei, S., Wang, K., Yuan, S., Lei, M., Cheng, K.: Research and design of cultural and creative industries in the qinba mountains of southern shaanxi. Results in Engineering **22**, 101956 (2024) https://doi.org/10.1016/j.rineng.2024.101956

[99] Guo, X., Ma, J.: Heritage applications of landscape design in environmental art based on image style migration. Results in Engineering **20**, 101485 (2023) https://doi.org/10.1016/j.rineng.2023.101485

Tables and Figures

Task	**Completeness (%)**	**Sub-task**	**Avg. time (s)**	**Errors (%)**
Creating a 3D-type event in which a violin's sound post (internal to the instrument) is shown	100	adding title	7	0
		adding description	10.8	0
		adding Wikidata entities	18.6	0
		adding the ID of the 3D model	16.2	0
Creating a presentation about Antonio Stradivari in the form of a Map-type event	100	adding title	12.8	0
		adding description	27.6	0
		adding Wikidata entities	51	0
		adding an image using an external link	25.4	60
		adding polygon coordinates	38	0
Creating a Slide-type event presenting a violin workshop	100	adding title	4.2	0
		adding description	14.8	0
		adding Wikidata entities	15.4	0
		adding a back-ground image by using an external link	15.6	0
		adding the "Il Poggio Montecastelli" web site as a digital object	134	80
Publishing a story map including the previous events.	100	viewing the story preview	6.8	0
		publishing the story	4.2	0

Table 1: Evaluation of the story map-creation interface. The first column reports the user macro-task names; the second, the percentage of users who eventually completed each macro task; the third, the subtasks required to complete the macro task; the fourth, the average time required by the users to complete each subtask; the fifth, the percentage of users who committed errors during subtask execution.

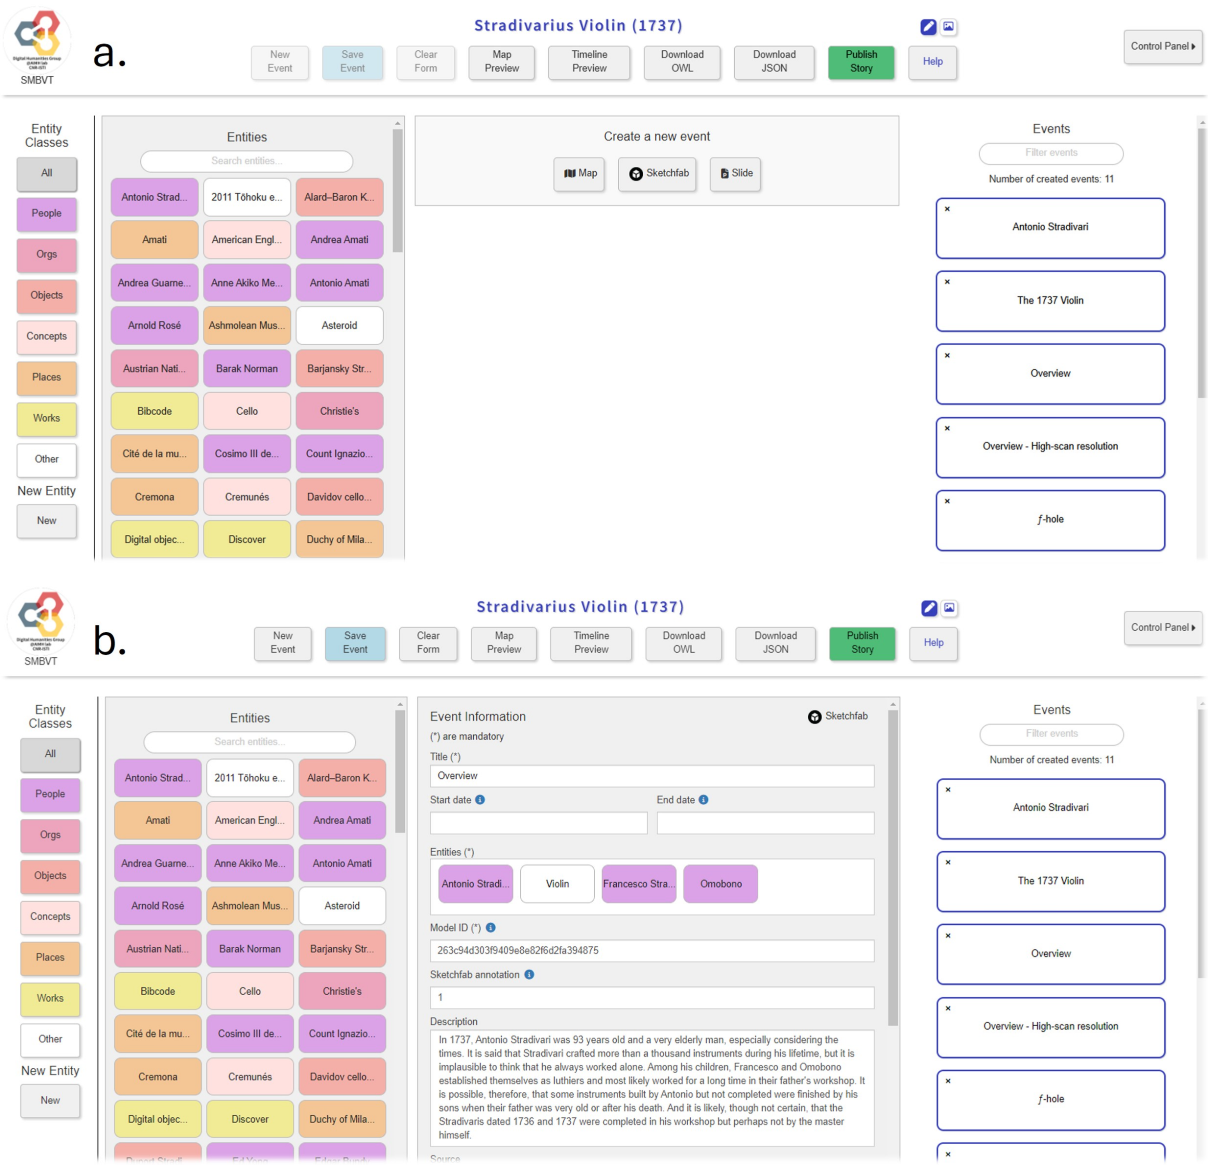

Fig. 1: The interface of our tool to create (a) narrative events and, in particular, (b) 3D-type events.

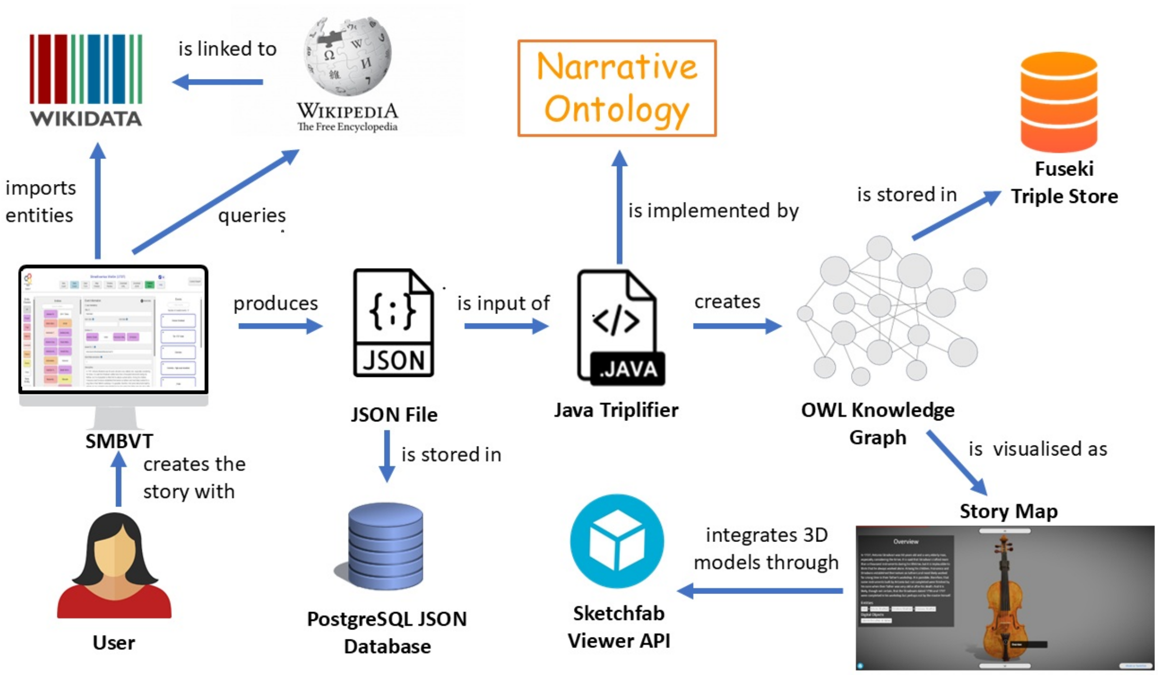

Fig. 2: The software architecture of the newly developed 3D Story Map Building and Visualising Tool (SMBVT).

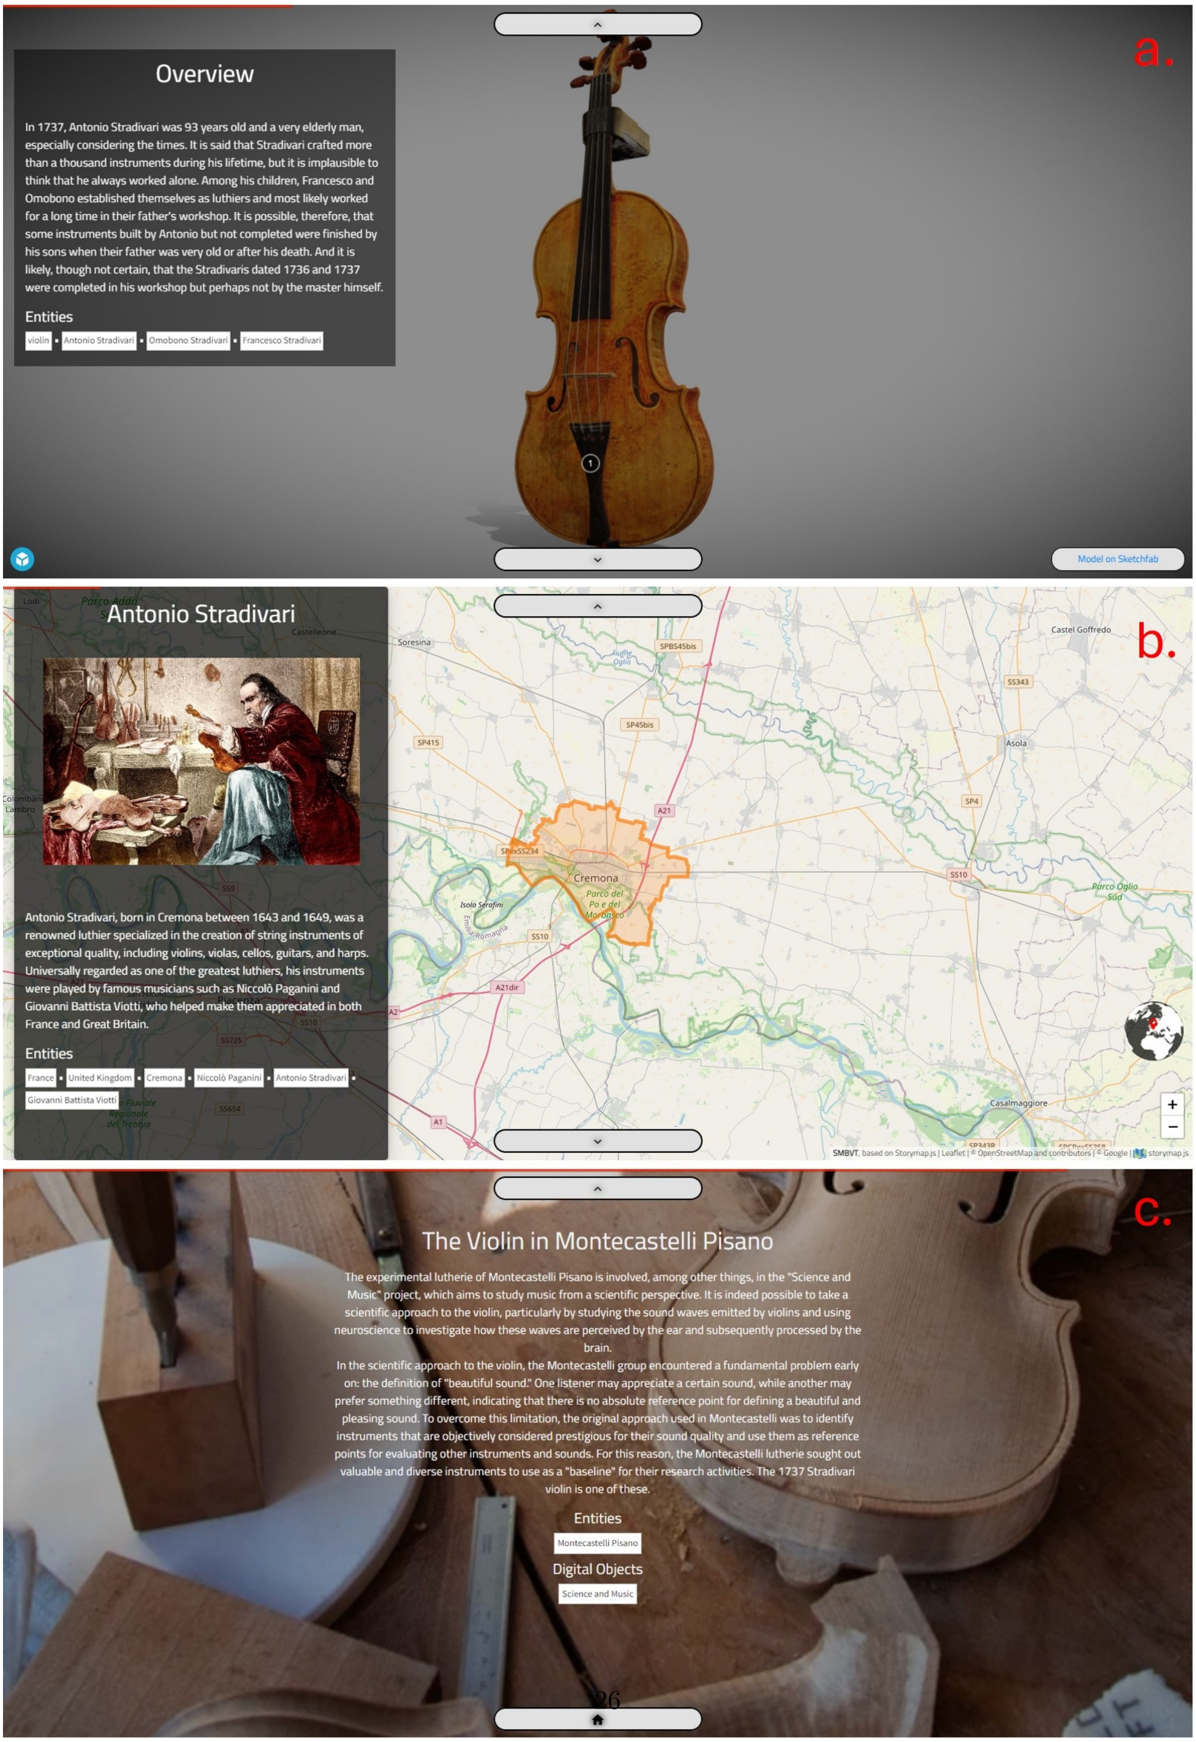

Fig. 3: a. A 3D-type event showing an overview of the 3D model of a Stradivari violin, dated 1737. b. A Map-type event reporting biographic and geographical references about Antonio Stradivari's life. c. A Slide-type event with a background image referring to the violin reproduction in the workshop of Montecastelli Pisano (Tuscany, Italy).

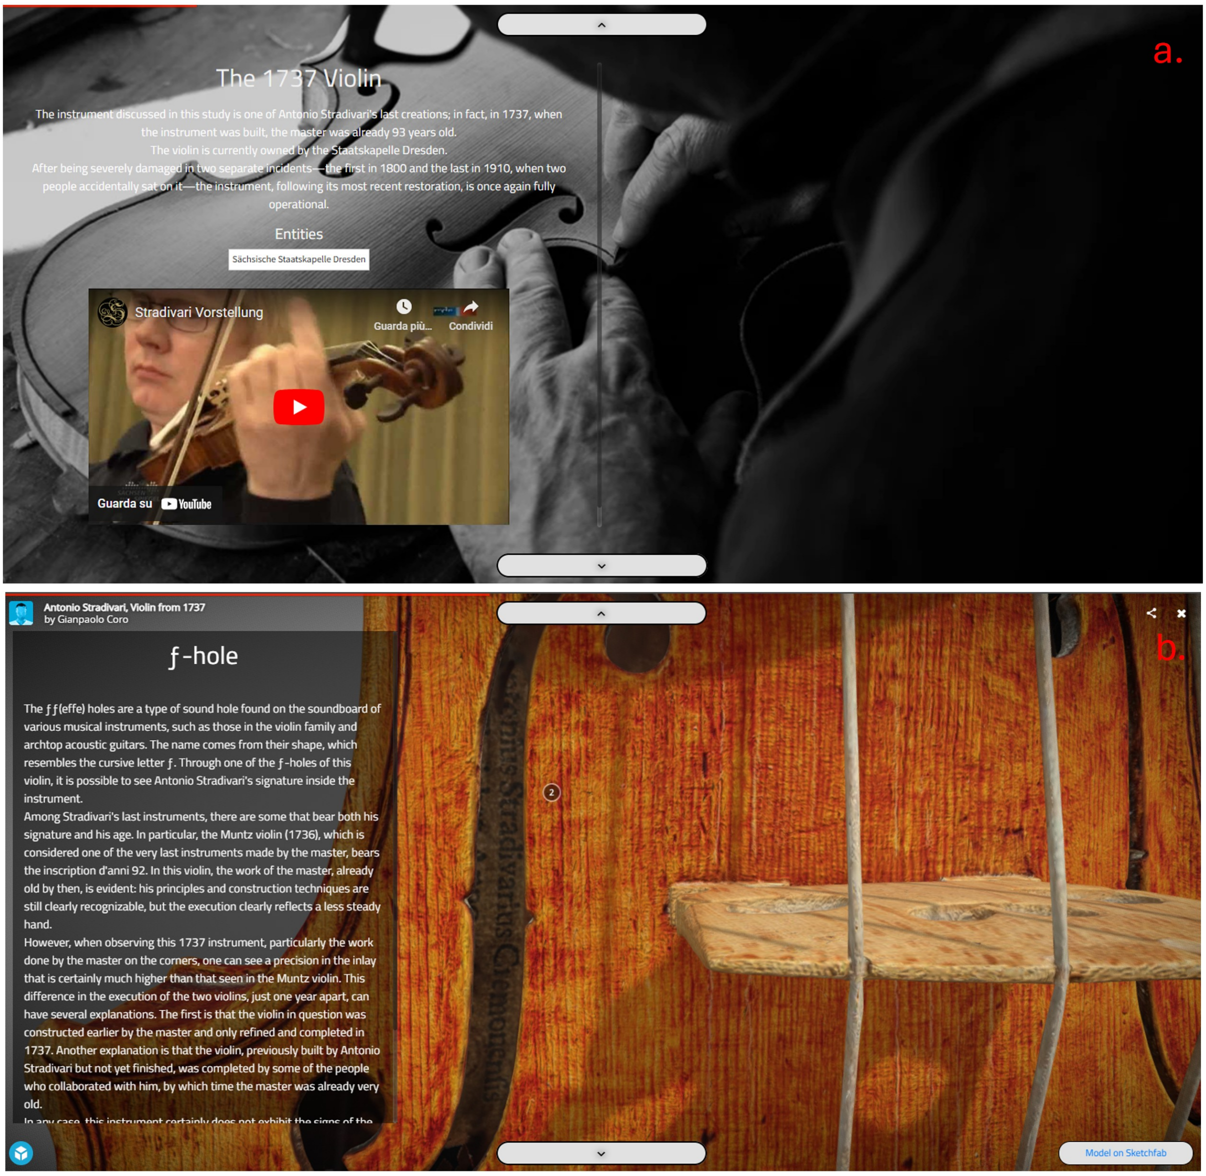

Fig. 4: a. A Slide-type event reporting a brief history of the Stradivari violin in textual format, embedding a YouTube video in which a violinist of the Dresden Staatskapelle plays the violin. b. A 3D-type event showing Stradivari's signature inside the 3D-reproduced violin sound box, with a side explanation of the f-hole function in a violin.

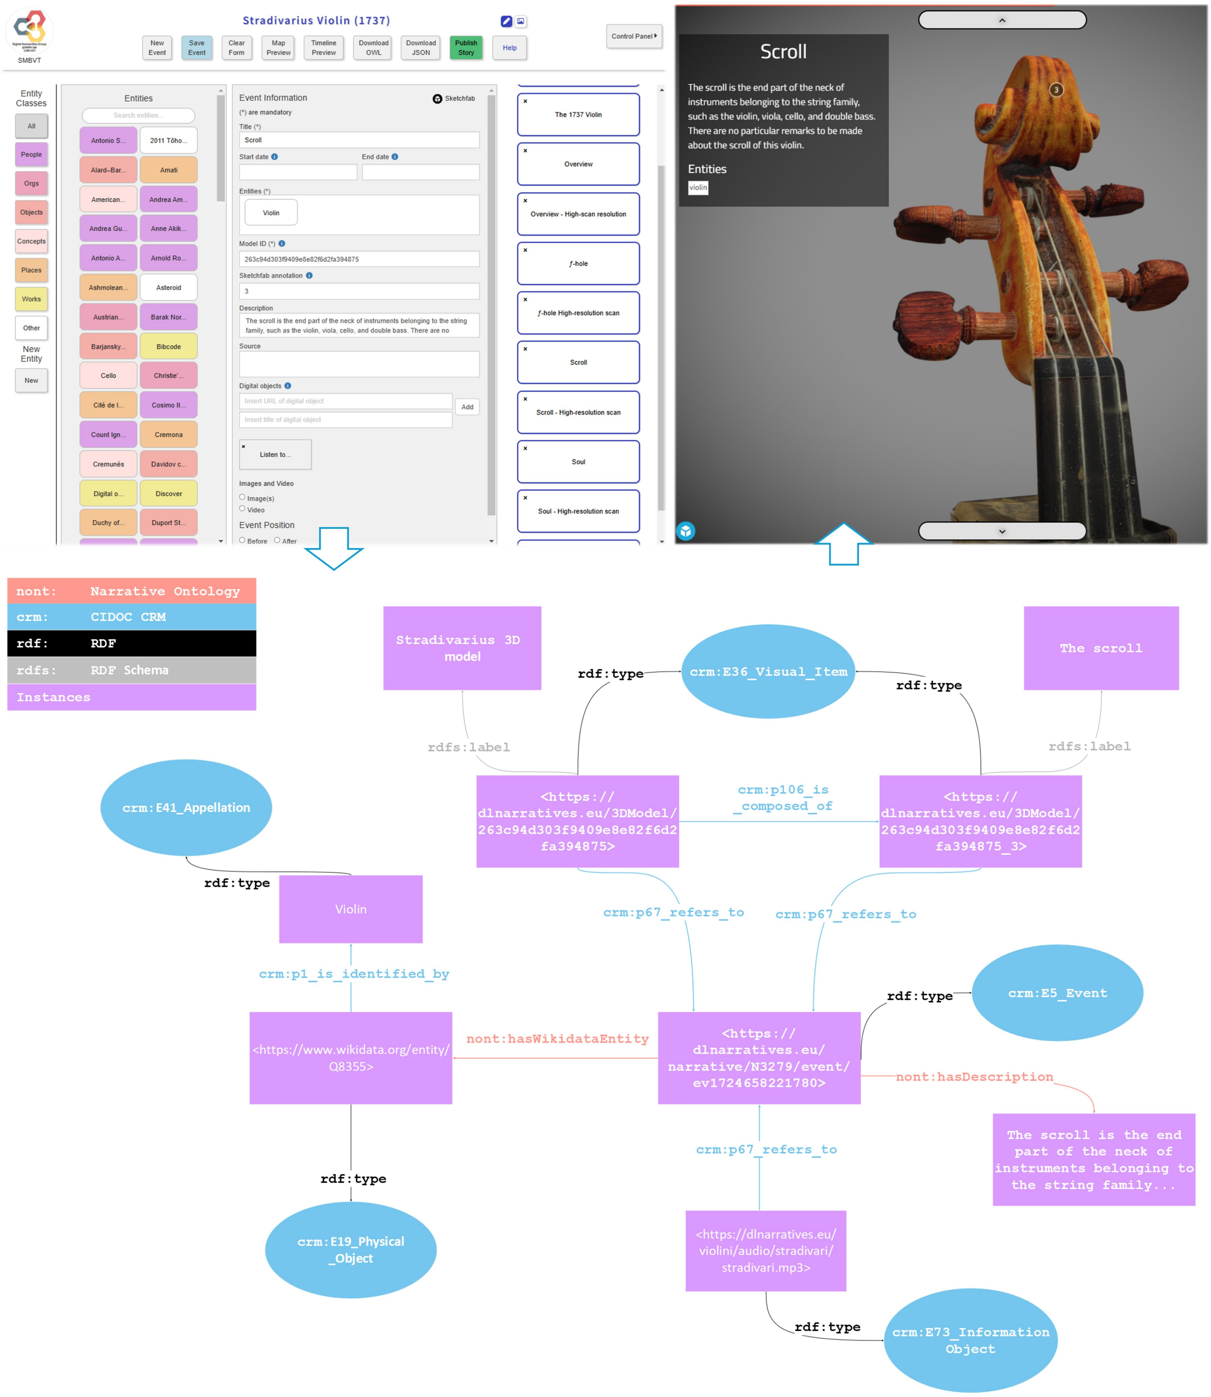

Fig. 5: The knowledge graph of the classes and properties that formally represent a 3D-type event (violin Scroll). The event was created through our tool (top-left image) and contains the identifiers of the 3D model on Sketchfab and the Scroll-related annotation. Contextual information from the graph is shown at the left of the 3D Scroll visual representation (top-right image).

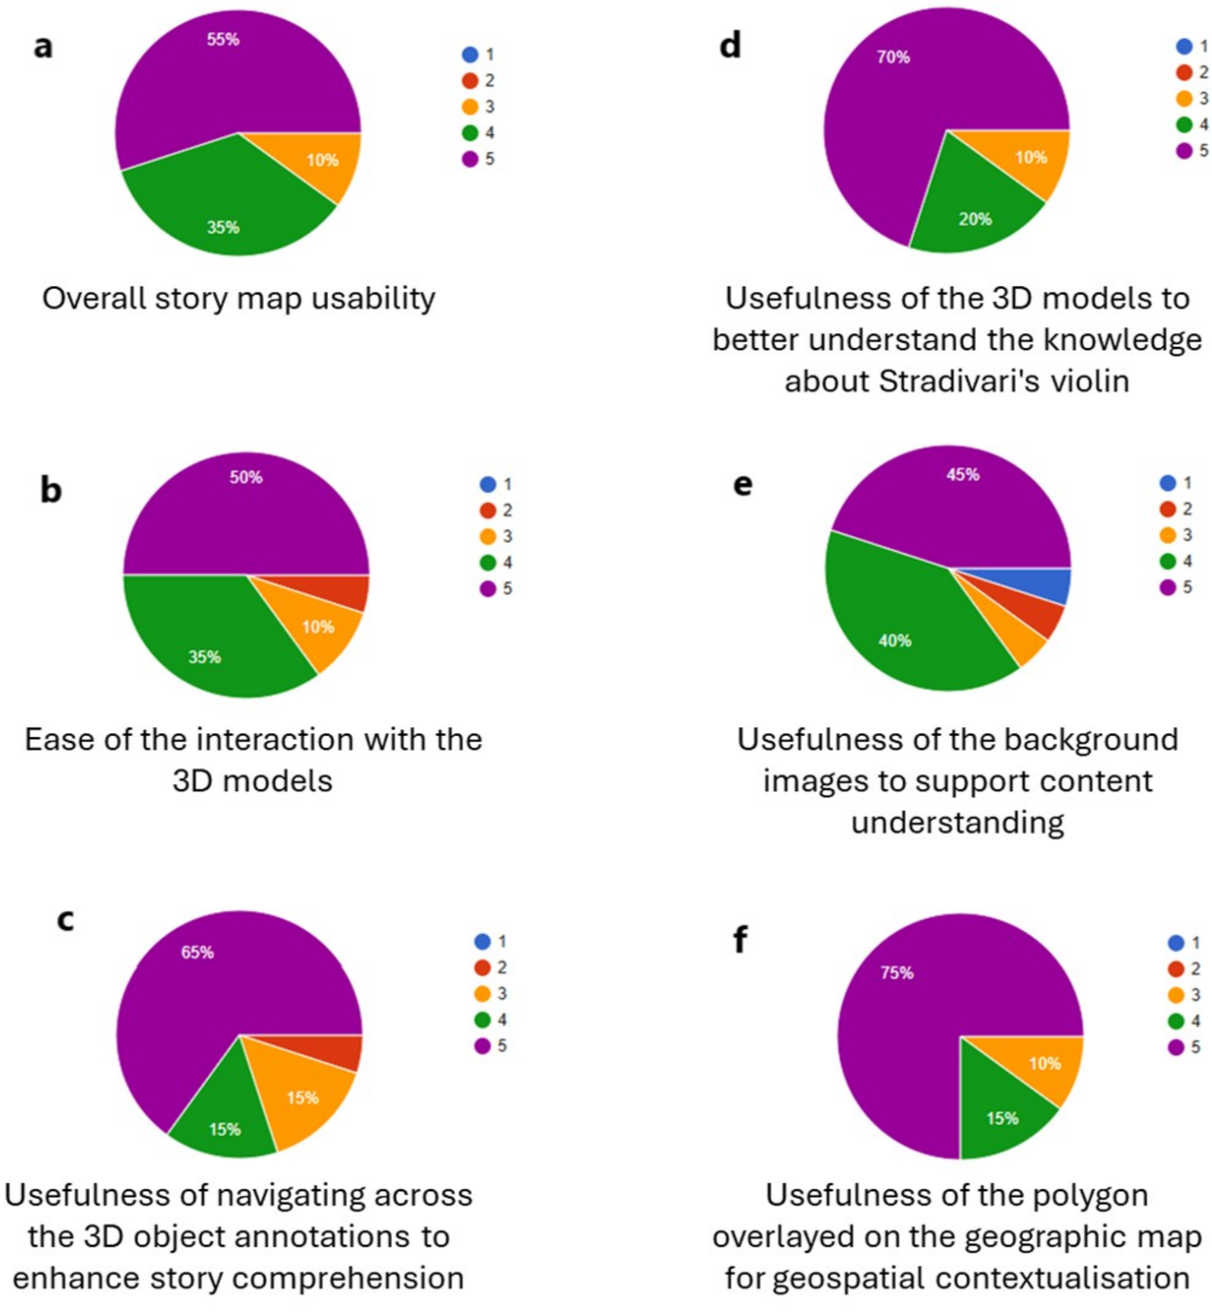

Fig. 6: Evaluation results of the story map visualisation, summarised as pie charts.

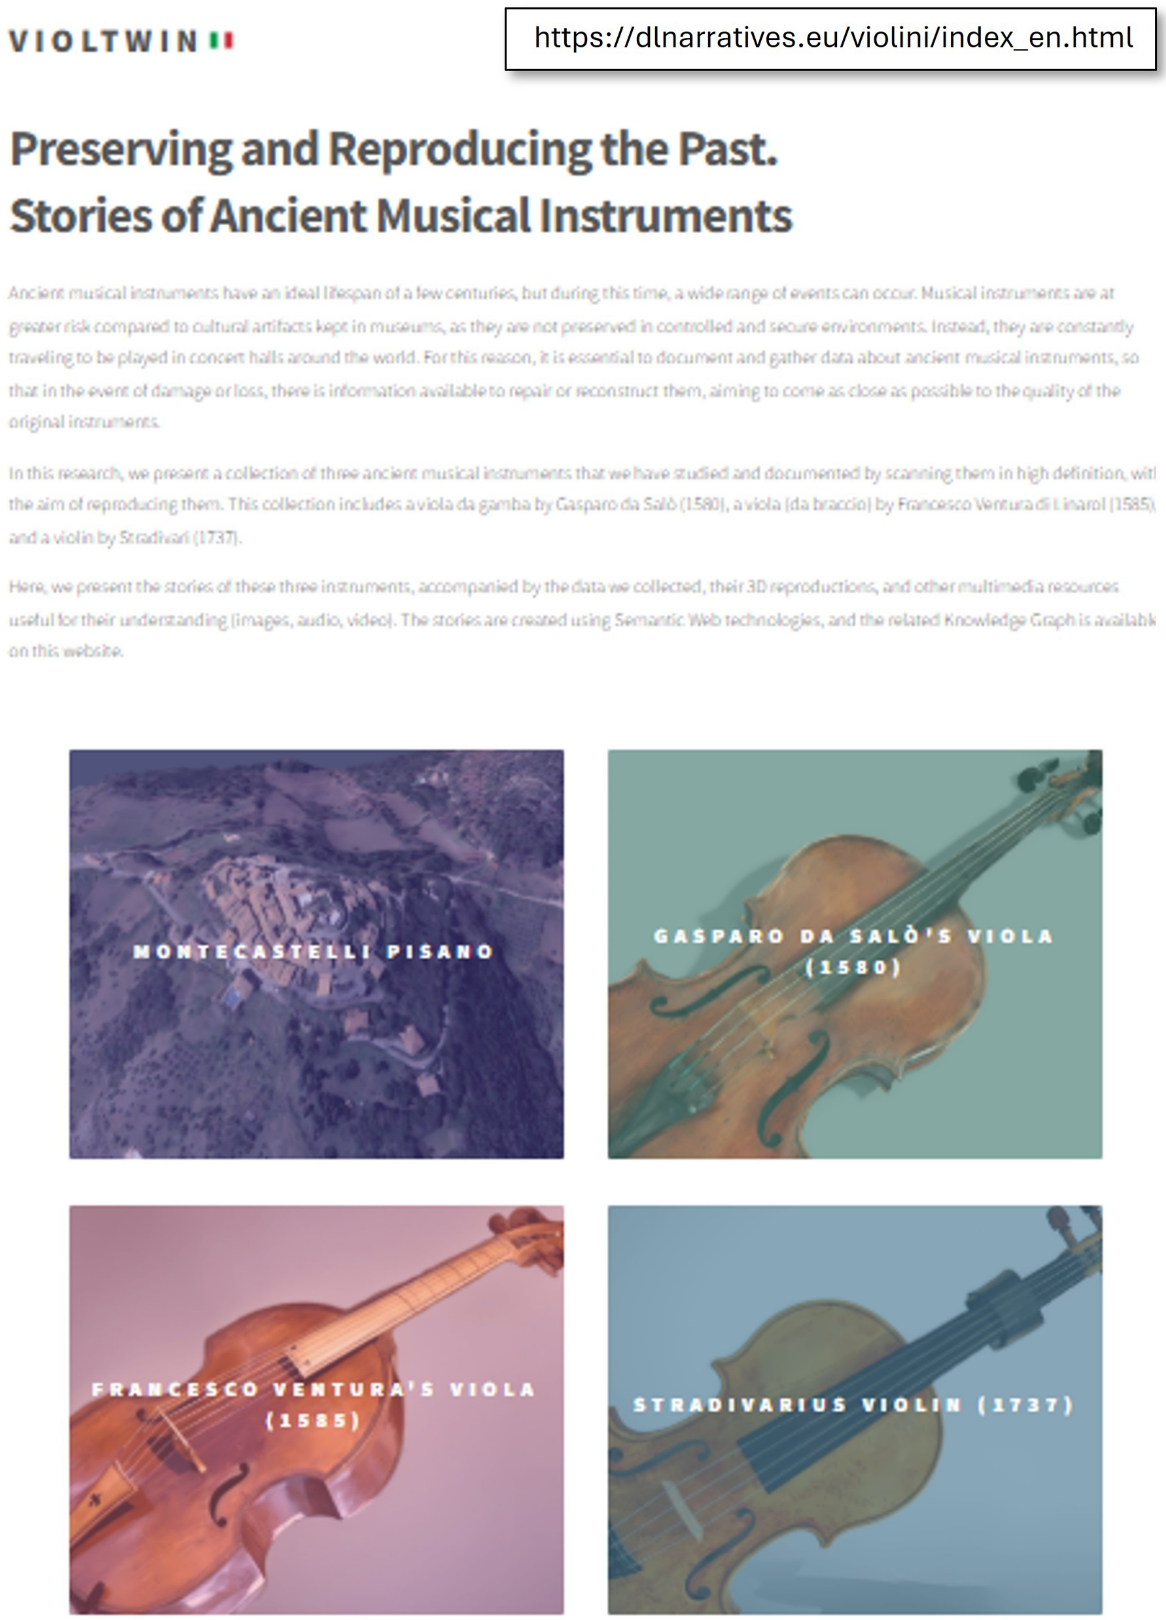

Fig. 7: Presentation of the ViolTwin initiative's web page for the promotion of ancient music instruments and cultural heritage preservation.