



From unstructured texts to semantic story maps

Valentina Bartalesi, Gianpaolo Coro, Emanuele Lenzi, Pasquale Pagano & Nicolò Pratelli

To cite this article: Valentina Bartalesi, Gianpaolo Coro, Emanuele Lenzi, Pasquale Pagano & Nicolò Pratelli (2023) From unstructured texts to semantic story maps, International Journal of Digital Earth, 16:1, 234-250, DOI: [10.1080/17538947.2023.2168774](https://doi.org/10.1080/17538947.2023.2168774)

To link to this article: <https://doi.org/10.1080/17538947.2023.2168774>



© 2023 The Author(s). Published by Informa UK Limited, trading as Taylor & Francis Group



Published online: 18 Jan 2023.



Submit your article to this journal [↗](#)



View related articles [↗](#)



View Crossmark data [↗](#)



From unstructured texts to semantic story maps

Valentina Bartalesi[†], Gianpaolo Coro[†], Emanuele Lenzi[†], Pasquale Pagano[†] and Nicolò Pratelli[†]

Istituto di Scienza e Tecnologie dell'Informazione 'A. Faedo' del Consiglio Nazionale delle Ricerche, Pisa, Italy

ABSTRACT

Digital maps greatly support storytelling about territories, especially when enriched with data describing cultural, societal, and ecological aspects, conveying emotional messages that describe the territory as a whole. Story maps are interactive online digital narratives that can describe a territory beyond its map by enriching the map with text, pictures, videos, and other multimedia information. This paper presents a semi-automatic workflow to produce story maps from textual documents containing territory data. An expert first assembles one territory-contextual document containing text and images. Then, automatic processes use natural language processing and Wikidata services to (i) extract key concepts (entities) and geospatial coordinates associated with the territory, (ii) assemble a logically-ordered sequence of enriched story-map events, and (iii) openly publish online story maps and an interoperable Linked Open Data semantic knowledge base for event exploration and inter-story correlation analyses. Our workflow uses an Open Science-oriented methodology to publish all processes and data. Through our workflow, we produced story maps for the value chains and territories of 23 rural European areas of 16 countries. Through numerical evaluation, we demonstrated that territory experts considered the story maps effective in describing their territories, and appropriate for communicating with citizens and stakeholders.

ARTICLE HISTORY

Received 27 October 2022

Accepted 11 January 2023

KEYWORDS

Story maps; semantic web; ontology; natural language processing; narratives; e-infrastructures; virtual research environments

1. Introduction

A narrative is the production of imaginative projects and experiences shown in movements and vocal expressions, and is a conceptual basis of collective human understanding (Wertsch and Roediger, 2008). Humans use stories to represent characters' intentions, feelings, ambitions and the attributes of objects and events (Delafield-Butt and Trevarthen 2015). A widely-held thesis in psychology to justify the centrality of narrative in human life is that humans make sense of reality by structuring events into narrative (Bruner 1991; Taylor 1992). Therefore, narratives are central to human activity in cultural, scientific, and social areas. They are also used between different scientific communities to create shared meanings (McInerney et al. 2014).

Maps have always geographically supported storytelling, stimulated people's imagination, and supported the creation of tools for developing arguments in scientific research. They can represent the geospatial components of narratives, e.g. a geographic space and a territory's influence on the

CONTACT Valentina Bartalesi ✉ valentina.bartalesi@isti.cnr.it  Istituto di Scienza e Tecnologie dell'Informazione 'A. Faedo' del Consiglio Nazionale delle Ricerche, Via Moruzzi 1, Pisa 56124, Italy

 Supplemental data for this article can be accessed online at <http://dx.doi.org/10.1080/17538947.2023.2168774>.

[†]These authors equally contributed to this work.

© 2023 The Author(s). Published by Informa UK Limited, trading as Taylor & Francis Group

This is an Open Access article distributed under the terms of the Creative Commons Attribution License (<http://creativecommons.org/licenses/by/4.0/>), which permits unrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited.

author, or can themselves be the story (Caquard and Cartwright 2014). Despite the importance of integrating emotional dimensions in maps is now widely recognised (Aitken and Craine 2006; Cartwright et al. 2008; Iturrioz and Wachowicz 2011), a map in itself is a rationalised geospatial representation that is very limited in conveying emotions. There is a perceptual gap between the territory and its map in narratives, which has been long analysed within Korzybski's General Semantics (Korzybski 1933): A *perceptive* cartographic challenge for a map is when it tries to represent also the life, emotions, reality, fiction, legends, and expectations associated with the described territory. As for digital maps, this challenge could be met by enriching geographic locations with media that communicate emotional messages, e.g. digital audio/video material to describe the overall territorial complexity. Story maps are a technical solution to fill the gap between a territory and its map in narratives. They are computer science realisations of narratives based on interactive online maps enriched with text, pictures, videos, data, and other multimedia information.

In this paper, we study how online story maps can fill the gap between a map and a territory in narratives. We start from a specific domain (mountain territories and value chains) and propose a general software solution for a semi-automatic transformation from text to narrative. In particular, we built a workflow to transform the textual descriptions of the value chains and territories of 23 rural European areas from 16 countries involved in the MOVING (MOUNTAIN Valorisation through INterconnectedness and Green growth) European project (MOVING "The MOVING European" 2020). MOVING is a project running from 2020 to 2024, gathering 23 organisations and companies that monitor, support, and conduct value chains in mountain areas. The project aims at building co-development policy frameworks across Europe that can contribute to the resilience and sustainability of mountain areas. The heterogeneous MOVING community of practice has produced textual descriptions of 23 territories and 455 value chains, which summarise the value chains' performance and subsistence. The 23 territory descriptions are summaries of the principal value chains of their corresponding rural areas (from the set of the 455 value chains), enriched with landscape, geographic, socio-economic, and touristic information. All descriptions were the synthesis of interviews with citizens and the territories' climatic, meteorological, and geographic data. As usual in ecological communities (Blue Cloud 2020; EcoScope 2022), data are associated with reference locations, environmental data, and unstructured data in textual formats. The textual representations summarise emotional, cultural, and societal contents to attract tourism and customers. Digitalising this information and assigning it to map locations can be complex (Peterle 2019). However, it helps extend the information's reachability, effectively narrate the content, and harmonise the data to produce one overall narrative of the territory out of many. Community members (usually ecologists) require IT solutions to build effective digitalisations. Practical help can come from IT workflows that translate textual information into structured information, which is finally organised into an interconnected knowledge base.

In this paper, we introduce a semi-automatic workflow to transform the unstructured information about a territory produced by the MOVING community of practice into a story map. The workflow uses natural language processing algorithms (Coro et al. 2017; Coro, Panichi, and Pagano 2019; Coro et al. 2021) to extract terms (persons, locations, organisations, and keywords) with high importance to understand the text (named entities). Successively, it associates the extracted terms with Wikidata entries and geographic coordinates. Then, it produces a sequence of story events enriched with titles, descriptions, entities, coordinates, and multimedia data. Finally, the workflow feeds a semantic knowledge base, based on an ontology for representing narratives (Meghini, Bartalesi, and Metilli 2021), to enable data discovery and integration. The workflow was designed to be independent of the specific application context. It is only constituted by open-source components (see Supplementary Material). The input is the textual description of a territory, organised as a sequence of events described through text, hyperlinks, and images. The output is an online, interactive, accessible story map of the territory, organised into geolocalised events, each with several Wikidata entries and multimedia information (image, video, hyperlinks) associated. We also released an editing environment for each story map that the community experts can use for revisions. Our workflow requires the human intervention to define an *a priori* structure for the story (a subdivision

into events) that the automatic processes enrich with multimedia and geospatial information. Different from other story-map creation tools (Walshe 2016; Alemy, Hudzik, and Matthews 2017; ArcGIS 2022; Timescape 2020), our tools are free-to-use, open-source, and oriented to the Open Science directives of transparency and reproducibility of data, results, and processes. Moreover, different from other story-visualisation software (Knightlab “StoryMap JS tool for narrative” 2020; Frenvik Sveen 2020; Map Box 2020; Becker, Köbben, and Blok 2009; Odyssey 2020), it offers a semantic knowledge base to automatically inter-connect the created stories, extract and discover knowledge from them, and comply with the Linked Open Data paradigm.

We conducted a qualitative assessment of our workflow after producing 23 story maps for the MOVING community of practice, i.e. one for each project partner and managed European sub-region. We numerically evaluated the feedback collected by the community experts who revised the story maps. Our workflow produced meaningful stories with appropriate contents for most evaluated aspects (titles, descriptions, images, coordinates, entities, digital objects), and the experts judged that the stories were able to reproduce the content expressed in the original data. Therefore, the produced stories were effective to fill the gap between unstructured data describing mountain territories and their representations on a map.

The paper is organised as follows: Section 2 describes the case study on the MOVING story maps and our workflow. Section 3 reports the evaluation results. Finally, Section 4 discusses the results and concludes.

2. Methodology

This section describes the components of our workflow (modules) that allowed us to manage the reported case study regarding the description of 23 mountain-related value chains and territories (Figure 1). This section first describes the case study (Section 2.1). Then, it explains the workflow module (story-structure building module) that transforms unstructured raw text into a structured story (Section 2.2). Moreover, it explains the module (story map creation module) that creates a structured version of the story maps in JSON format -- later stored on a PostgreSQL-JSON database -- and populates the semantic knowledge base (Section 2.3). Finally, it describes the Story Map Building and Visualising Tool (SMBVT) that hosts, publishes, and visualises the story maps (Section 2.4), and the evaluation strategy used (Section 2.5). The Supplementary Material contains references to all produced data, source code, services, and interfaces mentioned in this paper.




Figure 1. The architectural schema of our workflow for semi-automatically creating story maps.

2.1. Case study

The MOVING European project aims to support mountain value chains and their resilience to climate change through a bottom-up participatory process. Mountain ecosystems, cultural heritage, and society are highly relevant worldwide (Koohafkan and Altieri 2011; Egan and Price 2017). In Europe, mountain areas cover 36% of the territory. However, their management for sustainability needs improvement. The MOVING project produced data for 455 value chains of 16 European countries (MOVING “The MOVING European” 2020) and 23 European sub-regions (Table 1), which include information on economic, meteorological, climatic, ecological, cultural, and societal aspects. These sub-regions represent the vast diversity of mountain areas in Europe and neighbouring countries. Global descriptions of the sub-regions also summarise the contents of the principal value chains and their territories through 23 online general descriptions (MOVING “The MOVING Project” 2020), which serve to engage relevant stakeholders. We used these documents as the input for our analysis.

The data associated with each sub-region can be distinguished into three categories (Table 2): (i) textual descriptions of the region’s natural characteristics, (ii) quantitative descriptions of the region in terms of geography, population, income, tourism, and employment, (iii) key attributes of the regional products and value chains. These data were available to the project members, under textual, Web page, and MS Excel formats.

We pre-processed these data to prepare one new textual document (in MS Excel format) for each region. We organised this MS Excel, for each story, by describing one story event in each row. Therefore, each row reported (i) a title, (ii) a description, (iii) one representative image (optionally), (iv) hyperlinks to online multimedia material (optionally), and (v) the event *type*. Event type could be selected among Natural (i.e. natural aspects of the region), Historical (i.e. important historical events), or Valorisation (value chain and regional key attributes). The row sequence of an MS Excel file represented the event sequence of a sub-region story. An event sequence reported the same concepts expressed by the paragraphs of the corresponding sub-region description on the MOVING Web site (MOVING “The MOVING Project” 2020), prepared by sub-region experts. The data associated with each sub-region from the online descriptions and original data

Table 1. The MOVING project sub-regions with the indication of their belonging countries, regions, and associated principal value chains.

Country	Region	Value chain
Austria	Austrian Alps	Lamb from the region of Weiz
Bulgaria	Stara Planina	Public goods from High Nature Value Farming
Czechia	Šumava -- Cesky Les	Cattle farms
France	Corsica	Chestnut flour
France	Drome Valley	Sheep meat
Greece	Crete	Crete Carob
Hungary	Transdanubian Mountains	agroecological knowledge
Italy	Central Apennines	Alto-Molise dairy
Italy	Eastern Alps	Trento Doc wine
Italy	Northern Apennines	Chestnut flour
North Macedonia	Maleshevski mountains	Rural tourism
Portugthe al	Cordilheira central	Serra da Estrela Protected Designation of Origin cheese
Portugal	Maciço Noroeste	Douro wine
Romania	Southern Romanian Carpathian mountains	Certified ecotourism
Serbia	Dinaric Mountains	Dinaric Alps – Sjenica lamb Protected Designation of Origin
Slovakia	Slovak Carpathian mountains	Bio-honey
Spain	Betic Systems	Organic mountain olive oil
Spain	Sierra Morena	Los Pedroches Protected Designation of Origin Iberian Ham
Spain	Spanish Pyrenees	Mountain wine
Switzerland	Swiss Alps	Mountain grain
Switzerland	Swiss Jura	Tête de Moine Protected Designation of Origin cheese
Turkey	Beydaglari	Greenhouse tomato
UK-Scotland	Highlands and Islands	Speyside Malt Whisky

Table 2. The information associated to each sub-region and value chain (VC), organised into three categories. We use the notation P for province-scale and R for regional-scale.

Natural description	Quantitative description	Valorisation of VC
Member state responsible for the region	Local Administrative Unit (LAU)	Name of the VC
Description of the mountain landscape	Size of the area	Description of the VC
Presence of Protected Areas	Min and max altimetry of LAU	Challenges of the VC
Other interesting information	Total population	Type of VC
Reference mountain chain	Population density	Type of VC innovation
Reference mountain landscape	Population change rate in the last 10 years	Description of the local assets
	Regional population density	
	Average per capita income (P/R)	
	Gross Value Added (GVA) (P/R)	
	Primary sector share of GVA/year (P/R)	
	Secondary sector share of GVA/year (P/R)	
	Tertiary sector share of GVA/year (P/R)	
	Primary sector share of employment/year (P/R)	
	Secondary sector share of employment/year (P/R)	
	Tertiary sector share of employment/year (P/R)	
	Bed places in tourist accommodations/year (LAU)	
	Regional share of Bed places/year	
	Number of agricultural holdings	

(Table 2) were therefore imported into our MS Excel documents by attaching them to the most appropriate events. The 23 newly prepared documents were sequentially passed to our workflow as input data. This organisation phase of the sparse original data into events would be required also for other case studies and contexts.

2.2. Story-structure building module

The story-structure building module creates a structured story by processing the plain text of the input MS Excel rows (events). The module was developed in Java and is open source (see Supplementary Material). It characterises each story event through the text fragments associated with automatically-detected named abstract or physical objects in the text (*named entities*). The module automatically extracts named entities from the story titles and event descriptions. To this aim, it uses the NLPHub service (Coro et al. 2021), a distributed cloud computing system that orchestrates, interconnects, and combines the outputs of different state-of-the-art text mining processes hosted or integrated by the D4Science e-Infrastructure (Assante et al. 2019, 2022). These processes detect named entities and keywords from the text. Keywords are single or compound words with a recognised meaning within their context that help contextualise and understand the text. For simplicity, we will use the generic expression ‘named entity’ also for keywords. A cloud computing platform behind the scenes (Coro et al. “Parallelizing the execution of native” 2015; Coro et al. 2017) speeds up the processing by parallelising the named-entity extraction requests. It offers NLPHub a standardised service interface based on the Web Processing Service (WPS) standard of the Open Geospatial Consortium (Schut and Whiteside 2007), which makes it compliant with the Open Science paradigm (Hey, Tansley, and Tolle 2009). Our workflow invokes the NLPHub on each event description to extract named entities associated with *locations*, *persons*, *organisations*, and *keywords*.

As a further step, the process uses the Wikidata semantic service SPARQL endpoint (Vrandečić 2013; Wikidata 2022) to retrieve, for each extracted named entity, a possibly corresponding Wikidata entry. If correspondence is found, the Wikidata Internationalized Resource Identifier (IRI) and the links to the related Wikipedia pages are retrieved from the JSON answer and associated to the named entity. For example, for the ‘Olive oil’ named entity, the process would associate the Wikidata ‘olive oil’ entry, whose Wikidata-assigned IRI is ‘<https://www.wikidata.org/wiki/Q93165>’. A named entity is considered *valid* if it corresponds to a Wikidata entry not linked to a Wikipedia

disambiguation page. Moreover, the principal Wikipedia page should have a title perfectly corresponding to queried named entity. These rules improve the probability that the semantics of an extracted named entity was correctly identified. We defined all valid extracted named entities as the event-associated *entities*. For valid *location* and *keyword* entities, our algorithm also searches for spatial coordinates associated with the corresponding Wikidata entry and transforms them into decimal representation. It is worth noting that the Wikidata knowledge base changes over time. However, this change is unlikely to affect our Workflow because the story-structure building module collects IRIs that identify the Wikidata resources definitively. The descriptions and names of the Wikidata resources can change over time, but the IRIs remain the same. After the story-map publication, the entity's IRI link will correspond to the latest Wikidata page available. This way, the Wikidata updates are independent of the story map.

In summary, the first step of our story extraction module uses the NLPHub and Wikidata services to extract valid entities, possibly with associated spatial coordinates. The algorithm can be summarised as follows:

Algorithm 1 Entity extraction algorithm

```

for each event description and title
  invoke the NLPHub to extract all location, person, organisation, and keyword named entities
  for each extracted named entity
    check validity (i.e. no ambiguity exists) with respect to Wikidata/Wikipedia entries
    if the entity is location or keyword, try to retrieve associated coordinates from the Wikidata entry
  collect the list of entities and coordinates associated to the event.

```

The extracted spatial coordinates need revision to guarantee that they focus on the narrated territory. For example, a story map on the Austrian Alps might cite a cow breed also present in America. Thus some United States regions might be mentioned in an event and detected as named entities. Including too far points could be dispersive for the narrative and might result in a story jumping from one continent to another. In order to keep the extracted locations focused on a geographically consistent region, our module traces a bi-variate log-normal distribution on the longitude-latitude pairs. It sets the boundaries to the upper and lower log-normal confidence limits over the axes. All coordinates lying outside of these boundaries are considered outliers. If most coordinates refer to the same region, this process automatically sets the boundaries not too far from the region. Otherwise, if coordinates are uniformly distributed worldwide (e.g. in the case of a global-scale narrative), the boundaries will likely include all coordinates. The correct approximation of the events' coordinate distributions with a bi-variate log-normal was verified through sample Pearson's chi-square tests with 5% significance level in R (Ricci 2005).

After outlier removal, an event is assigned the best representative coordinates among its entities' coordinates. The assigned coordinates should spatially characterise the event and distinguish it from other events. For example, an event with references to Italy and a town in Tuscany, should better have the town's coordinates associated because the event is probably reporting information about the town, especially if previous events have already defined Italy as the narrative context. Our module realises this operation by selecting, for each event, the event's entity-coordinates that are either unique in the entire narration or has the lowest occurrence frequency. The selection also gives a higher priority to the coordinates that are not associated with previous events. In summary, this strategy assigns the most characteristic coordinates to each event. Still, some events might not have coordinates associated (e.g. events about employment) but cannot be left unassigned to give dynamism to the story while browsing between the events. Our module sets the coordinates of these events as generic, by assigning the coordinates with the highest occurrence frequency, i.e. the ones characterising the entire story. Optionally, a small amount of random noise can be given to distributing the equal-coordinate events around the assigned location.

The second part of the story-building module can therefore be summarised as follows:

Algorithm 2 Story-structure building algorithm

```

identify coordinate outliers through bi-variate log-normal confidence intervals
remove all outlier coordinates from the data attached to the extracted entities
for each event
    if the event has geospatial entities associated: assign the coordinates (among those of the associated entities) with the lowest
    frequency across the entire story. Prefer those that have not been assigned so far
    if the event has no geospatial entity associated: assign the coordinates with the highest frequency across all story entities.

```

This process produces a sequence of events for each story, each associated with a title, description, entities (with Wikidata IRIs associated), and coordinates. The sequence is saved as a Comma Separated (CSV) file and represents a structured version of the raw input text format that is later enriched for finalising the story map.

2.3. Story map creation module

Each story structured file is post-processed through a Python script to finalise the story map representation. Each story event is associated with the event type and multimedia hyperlinks specified in the original input text. All acronyms are expanded through a reference domain-specific dictionary to make descriptions less technical. Images are linked to the events if referred in the original input text. Otherwise, the first image associated with the event entities on Wikidata, ordered by their position in the text, is retrieved and linked. The event sequence with all associated entities, images, and links is described in JSON format, according to the schema used by the Story Map Building and Visualising Tool (SMBVT) (Bartalesi et al. 2022). This JSON document is an offline realisation of the story map. The Python script stores it on a PostgreSQL-JSON database used by the SMBVT for fast online visualisation (Section 2.4). Finally, the script invokes a JAVA-triplifier software that translates the JSON document into a Web Ontology Language (OWL) graph and stores it in an Apache Jena Fuseki triple store (Jena 2014). This graph complies with the Narrative Ontology model (Meghini, Bartalesi, and Metilli 2021), and its scope is to populate the SMBVT knowledge base and enable semantic queries for knowledge extraction.

2.4. Story map building and visualising tool

SMBVT (Bartalesi et al. 2022) is open-source software that represents narratives as a network of spatiotemporal events related by semantic relations (part-of, temporal, spatial, and causal relations) enriched by event components, i.e. the entities that take part in the event (e.g. persons, objects, places, concepts). For the present experiment, we used a free-to-use SMBVT instance hosted on the D4Science e-Infrastructure accessible after free registration to the platform (see Supplementary Material).

Following the Semantic Web approach (Berners-Lee, Hendler, and Lassila 2001), SMBVT assigns to each event and event component an IRI (Dürst and Suignard 2005). These IRIs are mainly extracted from Wikidata (Vrandecic 2013), which SMBVT uses as an external reference knowledge base. SMBVT retrieves events, entities and all associated data from a PostgreSQL-JSON database for visualisation. It synchronises this database with an OWL-graph representation of the stories stored on an Apache Jena Fuseki server, compliant with the Narrative Ontology model. The Fuseki server provides a SPARQL endpoint to query the complete graph of collected stories. This server organises the stories as the sub-graphs of one overall story graph, still compliant with the Narrative Ontology model. The graph automatically connects the stories through the entities shared between the events. The server allows executing SPARQL data-extraction queries on the entire story graph within or across the stories. This feature allows connecting the stories to other knowledge bases published as Linked Open Data (Bizer, Heath, and Berners-Lee 2009) (e.g.




Figure 2. Interface of the event search functionality. The example retrieves two stories containing the *Castanea sativa* entity.

Europeana, Isaac and Haslhofer (2013)). In particular, based on the SPARQL server, SMBVT offers an entity-search functionality through a Web interface (Figure 2) that allows querying the entire knowledge base in a user-friendly way. The interface allows users to search for an entity (helped by automatic completion) and retrieve the following information based on predefined SPARQL queries:

- (1) All stories in which the entity appears;
- (2) All the events of the stories in which the entity appears;
- (3) The number of entity occurrences across all stories;
- (4) The co-occurring entities across all events of all stories.

This feature is crucial to explore story inter-connections, for example, to retrieve (i) the Valorisation events involving a specific entity (e.g. sheep, beer, etc.), (ii) the nations sharing the same entities (e.g. products, export locations, etc.), (iii) the correlation between events through their co-occurrence, and (iv) the most frequent entities across the stories (e.g. the most common products). The necessity for quickly creating cross-story event connections and exploration facilities justifies the use of Semantic Web technologies.




Figure 3. The Story Map Building and Visualising Tool interface to create and modify story maps.




Figure 4. Examples of story map events. Each event is composed by the story-map title (in the top bar); a map (on the left-hand side), where the event-related pin is highlighted; an image (on the right-hand side); an event title (e.g. Organic Olive Oil Value Chain); an event description (under the title); the Wikidata entries associated (Entities section); and a hyperlink to a digital object (Digital Objects section, when present).

SMVBT provides an online graphical interface to create and manage the stories (Figure 3). This interface facilitates story-event building and event sequencing and contextualisation. SMBVT visualises the produced stories as story maps placing the narrative events on an interactive map that respects an event browsing order based on the user-defined plot (Figure 4). Each story event is associated with one positional pin, one image/video, a colour and pin style (depending on the event type), one title and descriptive text, several Wikidata entries (representing persons, locations, organisations, and other entities that occur in the event), and external digital objects (e.g. Web pages and Europeana objects). The software uses a customised version of the StoryMapJS library (Knightlab “StoryMap JS tool for narrative” 2020) for map interaction, event browsing, and visualisation. StoryMapJS allows managing large background maps and images associated with the events and can visualise stories represented as JSON documents. At the story-map loading time, SMBVT on-the-fly translates its PostgreSQL-JSON story representation into a StoryMapJS JSON-compliant document. Together with the story map visualisation, SMBVT also supports story visualisation as an event timeline, when temporal information is available for the events. Timeline visualisation is based on a customised version of the TimelineJS library (Knightlab “Timeline JS library” 2020).

The SMBVT story map publication process automatically generates a Web application that embeds all required JavaScript libraries, instructions and styles, story-related JSON documents, and event images. This application is transferred to a public-access Apache Tomcat Web service running in the D4Science e-Infrastructure, invisible to external Web crawlers. Story map publication can be invoked through the story map editing interface and via an HTTP-REST request. The publication process returns a public link to the Tomcat application. Each publication operation overwrites the previously published application so that the public link always points to the latest story-map version. This operation is necessary to support the continuous updating and enrichment of the story while offering the users always the latest version. Therefore it guarantees a long-term story’s maintenance, usability, and accessibility.

2.5. Evaluation methodology

We evaluated the quality of 23 story maps of the MOVING sub-regions, written in English and published through our workflow, by making 23 region-experts revise them. Each MOVING sub-region was assigned a reference expert who was asked to study the story map and use the tool to change the story. Live support was offered to use SMBVT to reduce interaction and technical hindrances. The experts were identified after a preliminary experiment, detailed in Bartalesi et al. (2022), conducted during the design of the story map visualisation interface: We called all MOVING project partners for participating to an anonymous survey about one story map describing the Apuan Alps. The survey contained 22 questions about the story map’s usability and usefulness in representing the territory. Forty-three members accepted to participate and returned impressions and suggestions. They were the primary contact points we asked for identifying the 23 voluntary experts who reviewed the story maps analysed in the present paper.

We evaluated how many aspects of the story the experts decided to modify. In particular, we concentrated on the following aspects (*assessment dimensions*):

- (1) Titles
- (2) Descriptions
- (3) Images
- (4) Geographic coordinates
- (5) Wikidata entries
- (6) Digital objects

Table 3. Total events and user-added events for the study-case regions, ordered by the number of user-events added.

Selected region and value chain	Expert-added events	Total events
Austrian Alps: Sheep farmers from the region of Weiz	2	13
Italy: Eastern Alps Trento Doc Wine Value Chain	2	13
France: Corsican Chestnut Flour Value Chain	1	12
Greece: Crete Central Rethymno Carob Value Chain	1	12
Portugal: Cordilheira Central (Serra da Estrela PDO Cheese)	1	12
Romania: Ecotourism Southern Romanian Carpathian Mountains	1	12
Serbia: Dinaric Mountains Sienica Lamb	1	12
Scotland: Speyside Malt Whisky Value Chain	1	12
Spain: Betic Systems Organic Olive Oil Value Chain	1	12
Spain: Sierra Morena -- Jamón Ibérico Value Chain	1	12
Spain: Spanish Pyrenees Mountain Wine Value Chain	1	12
Turkey: Beydaglari Tomatoes	1	12
Italy: Central Apennines Alto-Molise Dairy Value Chain	1	12
Bulgaria: Stara Planina High Nature Value Farming Value Chain	0	11
Czechia: Sumava – Cesky Les Cattle Value Chain	0	11
France: Drome Valley Sheep Meat Value Chain	0	11
Hungary: Transdanubia Mountains – Agroecological Knowledge Value Chain	0	11
Italy: Northern Apennines Chestnut Flour	0	11
North Macedonia: Rural Tourism in the Maleshevski Mountains	0	11
Portugal: Maciço Noroest	0	11
Slovakia: Bio-Honey Value Chain, Slovak Mountains	0	11
Swiss Alps Grain Value Chain	0	11
Swiss Jura Tête de Moine Cheese Value Chain	0	11

The percent modifications to these dimensions were used to identify the most sensitive ones and those our workflow correctly represented. The total modification percentage was used to assess the overall quality of the workflow results, i.e. to approximate the confidence in offering the automatically produced story maps to experts and stakeholders without posterior intervention.

3. Results

The 23 generated story maps contained 268 events overall, with 11-to-13 events for each story (Table 3). The experts added events to 13 stories (57% of the total) but did not delete events. The experts indeed judged the event-sequence lengths of the story maps sufficient for describing the value chains and territories. In the 13 modified stories, one added event always served to provide more information about the value chain stakeholders. In two stories, an additional event added more information on the value chain characteristics, geographical extension, and related economy. After the event additions, 17 stories (74% of the total) were judged ready for publication for the citizens and potential value-chain stakeholders and were published among the MOVING official results (Moving European Project 2022; Bartalesi et al. 2022). The experts also confirmed that these stories contained sufficient information to describe the value chains and the related territories concisely and effectively.

Six stories (26% of the total) required event modifications before being published. These stories contained 143 events overall (Figure 5). Two Spanish story maps (*Organic mountain olive oil* and *Los Pedroches Protected Designation of Origin Iberian Ham*) required changes principally to descriptions and geographic coordinates. In particular, descriptions were modified to change the English text for better understanding by (English-speaking) stakeholders. Geographic coordinates were modified to specify the production locations more precisely. Descriptions were also changed in the other 4 modified story maps. Coordinates required modification in two of these stories (from Portugal and Czech Republic) to specify 5 production locations more precisely. An aggregated overview of the modifications (Figure 6) highlights that the most frequently modified dimension across the revised stories was Descriptions (28%), followed by Geographic coordinates (20%), Images (8%), Titles (2%), Wikidata entries (1%), and Digital objects (1%). The 6 stories were accepted and published after the experts' modifications.




Figure 5. Bar charts reporting the numbers of expert-made changes to the assessment dimensions. The charts show all the story maps that were modified by the experts.

In summary, the fraction of stories over the total that could be accepted directly by the experts was ~40% (entire-story-generation accuracy). The fraction of events over the total that did not require modification was ~80% (event-generation accuracy). Therefore, the entity-extraction and assignment algorithm was overall satisfactory. The main modifications consisted in changing the event descriptions to better meet the stakeholders' and citizens' understanding. Across all stories, only 29 coordinates over the total 268 coordinates required a revision (~89% coordinate-assignment accuracy), indicating our coordinate-assignment algorithm's reasonably good reliability. The complete 23-story maps' creation, after the data pre-processing stage, required ~30 minutes, with an average of 1.5 minutes for each story. We asked an expert to build a story map on Apuan Alps starting from the same pre-processed data, and it took ~1 hour to achieve the same result. Therefore, our workflow would likely reduce the time to build a story map by over 97%.

As a general quality assessment, the MOVING community was involved in giving feedback on the story map appearance (Figure 4 and Supplementary Material) and event search functionality (Figure 2). A think-aloud method (Van Someren, Barnard, and Sandberg 1994) was used to collect free and open opinions and suggestions through the MOVING internal social networking facilities offered by the D4Science e-Infrastructure to support collaboration between the project partners (Moving European Project 2022). Story maps were appreciated for their cross-community and multi-disciplinary capacity to interactively access information and enhance territory understanding. Overall, the interaction was judged sufficiently user-friendly for communicating with local stakeholders and citizens. The event search functionality was more appreciated by data analysts of the MOVING community, who plan to discover general patterns in the European value chains once all 455 value chain story maps are created. They especially appreciated the possibility to seamlessly explore the network of stories and find their shared concepts.

4. Discussion and conclusions

In this paper, we have presented a semi-automatic workflow to transform textual documents containing territory information into story maps. Using Natural Language Processing (NLP) and Semantic Web technologies, our workflow produces online story maps composed of events with associated Wikidata entries and multimedia information. Each story map is published as an online interactive Web application. The stories are published altogether in a semantic knowledge base compliant with the Linked Open Data paradigm. This approach adds potential domain scalability




Figure 6. Summary charts of the total number of expert-made changes to the assessment dimensions.

to the stories because the knowledge base can directly be linked with other knowledge bases (e.g. about tourism, environmental sustainability, and transportation) through shared classes, thus forming larger knowledge bases (Meghini, Bartalesi, and Metilli 2017; Thanos et al. 2022). The internally-used NLP service complies with the Open Science directives of process standardisation and re-usability, and all used software is open source (see Supplementary Material). Therefore, the entire workflow goes towards meeting full Open Science compliance (Coro “Open Science and Artificial Intelligence” 2020). The described experiment was funded by the MOVING European project, which gave us the mandate to create the story-maps and supported the work of the experts (as project partners). The SMBVT long-term sustainability and availability are guaranteed by the ones of Wikidata and the hosting D4Science European e-infrastructure, which relies on shared resource funding from multiple disciplines and projects (Assante et al. 2019). The sustainability of the entire solution allowed us to make the SMVBT freely available as-a-service (Supplementary Material) for every citizen or scientist who may want to generate their own private or public story maps. It is worth noting that SMBVT is not responsible for the uploaded user-provided content and does not own it. Moreover, administrators cannot access private data; only users can access the building

interface of their story maps. Users can also upload their copyrighted material and distribute their story maps to selected people only.

After evaluating 23 story maps of mountain-related value chains, our results indicate that the stories had sufficient lengths and contents. The expert reviewers did not delete events and added relatively few events. Moreover, the stories were declared suitable for promoting citizens' awareness about the value chains and the territories and engaging local stakeholders. Therefore, based on the general and specific feedback received by the experts, the proposed workflow effectively filled the gap between the unstructured value chain and territory descriptions and their digital representation, i.e. it appropriately described these concepts beyond a map. Overall, the generated story maps can enhance the capacity to explore information, foster understanding, and encourage learning between different communities (McInerny et al. 2014).

One of the main advantages of our approach is that, through semantic queries, it contributes to discovering new knowledge from the data; for example, the value chains sharing the same environmental characteristics (e.g. rivers, lakes, vineyards, chestnut trees) and issues (e.g. depopulation, emigration, climate change problems), or providing similar products (e.g. cow or sheep milk cheese). Discovering new knowledge from the data is particularly useful for mountain ecosystems to design sustainable environmental management pathways and contribute to long-term cities' ecological sustainability. The extracted knowledge can indeed help understand and dam the vanishing of essential services in rural areas due to constant depopulation trends in rural communities. The world's rural population, currently at around 47%, will probably fall to 30% by 2050 (United Nations 2018). This trend poses problems to traditional and cultural heritage aspects and human health, welfare, and cities' ecological sustainability. Mass transfer from rural areas to cities generates a quick increase in population, pollution, and power consumption. The United Nations foresee that this trend will bring 7 billion people to populate cities by 2050 (United Nations 2015). Such large cities will likely be unsustainable because they would (i) burden public administration too much, (ii) not guarantee a cost-of-living adequate to salaries, (iii) excessively increase air pollution, power consumption, and greenhouse gas emissions, (iv) threaten people's health and consequently increase health expenditure. Moreover, when production from nearby rural areas decreases, cities necessarily increment resource importation from distant suppliers, and thus long-distance transportation, with an additional negative impact on the environment. In Europe, these aspects are particularly critical because mountains cover 36% of the territory and greatly provision public and private goods. In Italy, this problem is even more severe because rural areas cover up to 90% of the territory and contribute up to 50% of the national income and gross value added (Ministero Italiano delle politiche agricole alimentari e forestali 2020). Therefore, understanding the territorial aspects that might help dam rural (and mountain) area depopulation is crucial for both rural areas and cities' long-term sustainability. It would also meet the United Nations' Sustainable Development Goal 11 (SDG-11) ('sustainable cities and communities'), which calls for strategies to safeguard the economy and vanishing essential services in rural areas (including mountain areas). These strategies should include diversification in economic sectors (to attract multiple businesses and work types) and improving tourism and interest in the territory. In this context, story maps and knowledge discovery can help citizens, investors, and governmental authorities better understand the ecosystem services and value chains that can be improved to support sustainability strategies. The MOVING project has indeed adopted story maps as effective tools for this task because the experts judged them to appropriately convey 'information going beyond the map' for scientists, stakeholders, and the general public.

In future work, we will apply our workflow to all 455 value chain story maps of the MOVING project to automatically create an extensive knowledge base from which we will extract general European patterns and relations among the value chains. Moreover, we will demonstrate the cross-domain applicability of our workflow by seamlessly managing marine science case studies from the Blue Cloud European project (Blue Cloud 2020). Although this paper has presented the story maps created within the MOVING project, the workflow is already used in several other

contexts. For example, other authors have used it to produce stories about (i) a Medieval journey from Verona (Italy) to Konstanz (Germany) (Mele 2022), (ii) the history of the legends, biological investigations, and AI-based modelling for habitat discovery of the giant squid (Coro et al. “Improving data quality to build” 2015; Coro “The Giant Squid: When Myth” 2020), (iii) the possible future scenarios of key oceanic parameters’ change due to greenhouse gas emissions and climate change (InfraScience Lab 2019), and (iv) writers’ and artists’ biographies (DL Narratives 2022).

Acknowledgments

The authors thank the MOVING partners who have contributed to the experiment by revising the produced story maps.

Disclosure statement

No potential conflict of interest was reported by the author(s).

Funding

This work has received funding from the European Union’s Horizon 2020 research and innovation programme under the MOVING project (grant agreement no 862739).

Data availability statement

The authors confirm that the data supporting the findings of this study are available within the article [and/or] its supplementary materials.

References

- Aitken, Stuart, and James Craine. 2006. “Guest Editorial: Affective Geovisualizations.” *Directions Magazine* 7: 2011. doi:10.1109/MCG.2006.80.
- Alemy, Alexis, Sophia Hudzik, and Christopher N Matthews. 2017. “Creating a User-friendly Interactive Interpretive Resource with ESRI’s ArcGIS Story Map Program.” *Historical Archaeology* 51 (2): 288–297. doi:10.1007/s41636-017-0013-7.
- ArcGIS. 2022. “ArcGIS Story Map Building Tool.” <https://storymaps.arcgis.com>.
- Assante, Massimiliano, Leonardo Candela, Donatella Castelli, Roberto Cirillo, Gianpaolo Coro, Andrea Dell’Amico, Luca Frosini, et al. 2022. “Virtual Research Environments Co-creation: The D4Science Experience.” *Concurrency and Computation: Practice and Experience*. doi:10.1002/cpe.6925.
- Assante, Massimiliano, Leonardo Candela, Donatella Castelli, Roberto Cirillo, Gianpaolo Coro, Luca Frosini, Lucio Lelii, et al. 2019. “Enacting Open Science by D4Science.” *Future Generation Computer Systems* 101: 555–563. doi:10.1016/j.future.2019.05.063.
- Bartalesi, V., G. Coro, E. Lenzi, N. Pratelli, and P. Pagano. 2022. *MOVING D3.3 -- Tools for Science-Society-Policy Interfaces. Using Semantic Story Maps to Describe a Territory Beyond Its Map*. Technical Report. ISTI Research Report, MOVING, D3.3, 2022.
- Becker, T., B. Köbben, and C. A. Blok. 2009. “Timemapper: Visualizing Moving Object Data Using WMS Time and SVG SMIL Interactive Animations.” In *Proceedings SVGOpen 2009*, SVGOpen. 7th International Conference on Scalable Vector Graphics (SVGOpen 2009), 02-04-10-2009, Mountain View, United States.
- Berners-Lee, Tim, James Hendler, and Ora Lassila. 2001. “The Semantic Web.” *Scientific American* 284 (5): 34–43. doi:10.1038/scientificamerican0501-34.
- Bizer, C., Heath, T., & Berners-Lee, T. (2009). “Linked Data - The Story So Far.” *International Journal on Semantic Web and Information Systems (IJSWIS)* 5 (3): 1–22. doi:10.4018/jswis.2009081901.
- Blue Cloud. 2020. “The Blue Cloud European Project.” On-Line. Project Web site: <https://blue-cloud.d4science.org/>.
- Bruner, Jerome. 1991. “The Narrative Construction of Reality.” *Critical Inquiry* 18 (1): 1–21. doi:10.1086/448619.
- Caquard, Sébastien, and William Cartwright. 2014. “Narrative Cartography: From Mapping Stories to the Narrative of Maps and Mapping.” *The Cartographic Journal* 51 (2): 101–106. doi:10.1179/0008704114Z.000000000130.

- Cartwright, William, Adrian Miles, Brian Morris, Laurene Vaughan, and Jeremy Yuille. 2008. *Developing concepts for an affective atlas*. 219–234. Berlin Heidelberg: Springer.
- Coro, Gianpaolo. 2020. “Open Science and Artificial Intelligence Supporting Blue Growth.” *Environmental Engineering & Management Journal (EEMJ)* 19 (10): 1719–1729.
- Coro, Gianpaolo. 2020. The Giant Squid: When Myth Encounters Science. <https://tool.dlnarratives.eu/storymaps/gianpaolo.51/N305/?visualization=map>.
- Coro, Gianpaolo, Leonardo Candela, Pasquale Pagano, Angela Italiano, and Loredana Liccardo. 2015. “Parallelizing the Execution of Native Data Mining Algorithms for Computational Biology.” *Concurrency and Computation: Practice and Experience* 27 (17): 4630–4644. doi:10.1002/cpe.v27.17.
- Coro, Gianpaolo, Chiara Magliozzi, Anton Ellenbroek, and Pasquale Pagano. 2015. “Improving Data Quality to Build a Robust Distribution Model for Architeuthis Dux.” *Ecological Modelling* 305: 29–39. doi:10.1016/j.ecolmodel.2015.03.011.
- Coro, Gianpaolo, Giancarlo Panichi, and Pasquale Pagano. 2019. “An Open Science System for Text Mining.” In *CLiC-it*.
- Coro, Gianpaolo, Giancarlo Panichi, Pasquale Pagano, and Erico Perrone. 2021. “NLPHub: An E-Infrastructure-based Text Mining Hub.” *Concurrency and Computation: Practice and Experience* 33 (5): Article ID e5986. doi:10.1002/cpe.v33.5.
- Coro, Gianpaolo, Giancarlo Panichi, Paolo Scarponi, and Pasquale Pagano. 2017. “Cloud Computing in a Distributed E-infrastructure Using the Web Processing Service Standard.” *Concurrency and Computation: Practice and Experience* 29 (18): Article ID e4219. doi:10.1002/cpe.4219.
- Delafield-Butt, Jonathan T, and Colwyn Trevarthen. 2015. “The Ontogenesis of Narrative: From Moving to Meaning.” *Frontiers in Psychology* 6: Article ID 1157. doi:10.3389/fpsyg.2015.01157.
- DL Narratives. 2022. “Examples of Narratives in Digital Libraries.” <https://dlnarratives.eu/narratives.html>.
- Dürst, Martin, and Michel Suignard. 2005. *Internationalized resource identifiers (IRIs)*. Technical Report. W3C.
- EcoScope. 2022. “The EcoScope International Project for Ecocentric Management and Healthy Marine Ecosystems.” <https://ecoscopium.eu/>.
- Egan, Paul A, and Martin F Price. 2017. *Mountain Ecosystem Services and Climate Change: A Global Overview of Potential Threats and Strategies for Adaptation*. Paris: UNESCO Publishing.
- Frenvik Sveen, Atle. 2020. “Leaflet Storymaps Template.” <https://github.com/atlefren/storymap>.
- Hey, Tony, Stewart Tansley, and Kristin M Tolle. 2009. *The Fourth Paradigm: Data-Intensive Scientific Discovery*. Vol. 1. Redmond, WA: Microsoft Research Redmond.
- InfraScience Lab. 2019. “Climate change effects on oceanic parameters.” Accessed January 4 2023 <https://dlnarratives.eu/timeline/climate.html>.
- Isaac, Antoine, and Bernhard Haslhofer. 2013. “Europeana Linked Open Data—data. Europeana. Eu.” *Semantic Web* 4 (3): 291–297. doi:10.3233/SW-120092.
- Iturrioz, Teresa, and Monica Wachowicz. 2011. “An artistic perspective for affective cartography.” In *Mapping Different Geographies*, 74–92. Berlin: Springer. <https://jena.apache.org/documentation/fuseki2/>.
- Jena, Apache. 2014. “Apache Jena Fuseki.” *The Apache Software Foundation* 18:
- Knightlab. 2020. “StoryMap JS tool for narrative building.” <https://storymap.knightlab.com>.
- Knightlab. 2020. “Timeline JS library.” <https://timeline.knightlab.com/>.
- Koohafkan, Parviz, and Miguel A Altieri. 2011. *Globally Important Agricultural Heritage Systems: A Legacy for the Future*. Rome: Food and Agriculture Organization of the United Nations Rome.
- Korzybski, Alfred. 1933. “A Non-Aristotelian System and its Necessity for Rigour in Mathematics and Physics.” In *Science and Sanity: An Introduction to Non-Aristotelian Systems and General Semantics*, Lancaster, PA: Lancaster.
- Map Box. 2020. “Map Box Storymaps Template.” <https://www.mapbox.com/solutions/interactive-storytelling>.
- McInerny, Greg J, Min Chen, Robin Freeman, David Gavaghan, Miriah Meyer, Francis Rowland, David J Spiegelhalter, et al. 2014. “Information Visualisation for Science and Policy: Engaging Users and Avoiding Bias.” *Trends in Ecology & Evolution* 29 (3): 148–157. doi:10.1016/j.tree.2014.01.003.
- Meghini, Carlo, Valentina Bartalesi, and Daniele Metilli. 2017. “Using Formal Narratives in Digital Libraries.” In *Italian Research Conference on Digital Libraries*, 83–94. Modena: Springer.
- Meghini, Carlo, Valentina Bartalesi, and Daniele Metilli. 2021. “Representing Narratives in Digital Libraries: The Narrative Ontology.” *Semantic Web* 12 (2): 241–264. doi:10.3233/SW-200421.
- Mele, Michela. 2022. “Epistola De Peregrinatione Ad Urbem Constantiam.” <https://tool.dlnarratives.eu/storymaps/adminNarra.1/N1238/?visualization=map>.
- Ministero Italiano delle politiche agricole alimentari e forestali. 2020. “Rete Rurale Nazionale 2014–2020.” <https://www.reterurale.it/flex/cm/pages/ServeBLOB.php/L/IT/IDPagina/551>.
- MOVING. 2020. “The MOVING European Project -- Mountain Valorisation Through Interconnectedness and Green Growth.” Accessed January 4 2023. <https://www.moving-h2020.eu/>.
- MOVING. 2020. “The MOVING Project Reference RegionsBlue Cloud European Project.” <https://www.moving-h2020.eu/reference-regions/>.

- Moving European Project. 2022. "The Moving EU Project Virtual Research Environment for Story Map Building." <https://moving.d4science.org/group/moving-gateway> after free registration.
- Odyssey. 2020. "Odyssey Story Maps Template." <http://cartodb.github.io/odyssey.js/>.
- Peterle, Giada. 2019. "Carto-fiction: Narrativising Maps Through Creative Writing." *Social & Cultural Geography* 20 (8): 1070–1093. doi:10.1080/14649365.2018.1428820.
- Ricci, Vito. 2005. "Fitting Distributions with R." *Contributed Documentation Available on CRAN* 96: 1–24.
- Schut, Peter, and A Whiteside. 2007. *OpenGIS Web Processing Service*. Open Geospatial Consortium. Open Geospatial Consortium Technical Reports.
- Taylor, Charles. 1992. *Sources of the Self: The Making of the Modern Identity*. Cambridge, MA: Harvard University Press.
- Thanos, Costantino, Carlo Meghini, Valentina Bartalesi, and Gianpaolo Coro. 2022. "An Exploratory Approach to Archaeological Knowledge Production." *International Journal on Digital Libraries* 23: 1–9. doi:10.1007/s00799-022-00324-3.
- Timescape. 2020. "Timescape Story Map Building Tool." <https://www.recombine.net/projects/timescape.html>.
- United Nations. 2015. "Sustainable Development Goal 11, 'Make Cities and Human Settlements Inclusive, Safe, Resilient and Sustainable'." On-Line. Accessed January 4 2023. <https://sdgs.un.org/goals/goal11>.
- United Nations. 2018. "68% of the World Population Projected to Live in Urban Areas by 2050, Says UN." On-Line. <https://www.un.org/en/desa/68-world-population-projected-live-urban-areas-2050-says-un>.
- Van Someren, Maarten, Yvonne F Barnard, and J Sandberg. 1994. *The Think Aloud Method: A Practical Approach to Modelling Cognitive*. 11, London: AcademicPress.
- Vrandečić, Denny. 2013. "The Rise of Wikidata." *IEEE Intelligent Systems* 28 (4): 90–95. doi:10.1109/MIS.2013.119.
- Walshe, Nicola. 2016. "Using ArcGIS Online Story Maps." *Teaching Geography* 41 (3): 115–117.
- Wertsch, James V, and Henry L Roediger. 2008. "Collective Memory: Conceptual Foundations and Theoretical Approaches." *Memory* 16 (3): 318–326. doi:10.1080/09658210701801434.
- Wikidata. 2022. "Wikidata SPARQL Web Service." <https://query.wikidata.org/>.