

Towards User Interfaces for Semantic Storytelling

Julián Moreno-Schneider^(✉), Peter Bourgonje, and Georg Rehm

German Research Center for Artificial Intelligence (DFKI),
Alt-Moabit 91c, 10559 Berlin, Germany
{julian.moreno_schneider,peter.bourgonje,georg.rehm}@dfki.de
<http://digitale-kuratierung.de>

Abstract. Digital content and online media have reached an unprecedented level of relevance and importance. In the context of a research and technology transfer project on Digital Curation Technologies for online content we develop a Semantic Storytelling prototype. The approach is based on the semantic analysis of document collections, in which, among others, individual analysis results are, if possible, mapped to external knowledge bases. We interlink key information contained in the documents of the collection, which can be essentially conceptualised as automatic hypertext generation. With this semantic layer on top of the set of documents in place, we attempt to identify interesting, surprising, eye-opening relationships between different concepts or entities mentioned in the document collection. In this article we concentrate on the current state of the user interfaces of our Semantic Storytelling prototype.

Keywords: Digital Curation Technologies · Linked data · Natural language processing · Semantic storytelling

1 Introduction

Digital content and online media have reached an unprecedented level of relevance and importance, especially with regard to commercial, political and societal aspects, debates and collective decisions. In that regard, one of the many technological challenges refers to better support and smarter technologies for knowledge workers, i.e., persons, who work primarily at and with a computer, who are facing an ever increasing stream of heterogeneous incoming information and who create, based on the specific requirements and expectations of the sector they work in and their job profile, in a rather general sense, new information. For example, experts in a digital agency build mobile apps or websites for clients who provide the digital agency with documents, data, pictures, videos and other assets. Knowledge workers in a library digitise a specific archive, augment it with additional information and publish the archive online. Journalists need to stay on top of the news stream including blogs, microblogs, newswires, websites etc. in order to produce a shorter or longer article on a breaking topic. All these different professional environments and contexts can benefit immensely from semantic

technologies that support the knowledge workers in their respective activities: finding relevant information, highlighting important concepts, sorting incoming documents in multiple different ways, translating articles in foreign languages, suggesting interesting topics. We call these different semantic services, that can be applied in multiple different professional environments that all have to do with the processing, analysis, synthesis and production of digital information, *Curation Technologies*.

In the context of our research and technology transfer project Digital Curation Technologies we develop a Curation Platform that offers language- and knowledge-aware services such as semantic analysis, search, analytics, recombination and generation (e.g., thematic, chronological and spatial) for the curation of various types of digital content. Our goal is to shorten the time it takes knowledge workers to familiarise themselves with a potentially large collection of documents by extracting relevant data and presenting it in a way that allows them to be more efficient, especially in the situation when the respective knowledge worker is not a domain expert with regard to the topics discussed in the document collection. In the project we work with data sets provided by the project partners, four SME companies active in different sectors.¹ The data sets cover multiple domains, some are publicly available, others are not.

In addition to semantic text analytics there has been, in recent years, a renewed interest in Natural Language Generation applications, especially with regard to content generation based on structured data sets. At the same time there is a growing need for semantic technologies that help content curators and knowledge workers to make the most of their time, to support them by automating or semi-automating some of the more time-consuming tasks and to accomplish them in a more efficient way, ideally also by providing content products that are of higher quality. There are multiple potential use cases for these technologies in many different sectors [3, 4, 18, 20]. In this paper we focus upon the semantic analysis of document collections in order to suggest interesting story paths between different entities mentioned in the collection. We call this approach *Semantic Storytelling*.

The immediate availability of large amounts of semantic data and knowledge bases makes this approach possible by applying semantic analysis methods, most of them based on Natural Language Processing (NLP). The annotation of semantic information in the source documents provides the necessary additional knowledge needed to tackle the ambitious task of computing and recommending potential story paths. Let us consider the example of the Panama Papers.² It is virtually impossible for one or even a team of investigative journalists to go through 11.5 million documents (2,6 TB of data) to find all the relevant information needed to prepare a report, news article or book on the investigation of the collection. Especially in this but also in many other use cases it would save a lot of time and resources to make use of a system that can, ideally automatically,

¹ The four SME partners are ART+COM AG, Condat AG, 3pc Neue Kommunikation GmbH and Kreuzwerker GmbH.

² <https://panamapapers.icij.org>.

process the whole collection and offer the investigators potential story paths in order to use them as starting points for their own articles.

The remainder of this paper is structured as follows. Section 2 covers related work. Section 3 provides a description of the current state of our curation platform. In Sect. 4 we discuss our results. Section 5 concludes the article and presents suggestions for future work.

2 Related Work

There are several approaches that are closely related to our Semantic Storytelling concept, all of them concentrating on their own specific objectives and providing solutions for their respective challenges. Our key challenge is to generate, ideally with limited or no human intervention at all, a story based on semantically annotated natural language text collections.

Some systems focus on providing content or applications for entertainment purposes. For example, Wood [32] uses a collection of pictures, videos and other media to generate albums, Gervás et al. [9] focus on storytelling in the gaming domain. Other groups use story structuring methods as part of therapy programs such as Kybartas et al. [13]. Other approaches focus on storytelling in a particular domain, typically recipes [7, 8] or weather reports [1, 11, 22, 30], requiring knowledge about characters, actions, locations, events, or objects that exist in this particular domain [10, 23, 31].

A notable exception to relying on domain knowledge is Li et al.'s approach [15], who attempt to construct plot graphs from a set of stories annotated using a crowd sourcing approach. The resulting plot graphs are used to generate alternative versions of the annotated stories, which are then evaluated again using a crowd sourcing approach. The evaluation task consists of either accepting a story as a logical piece of content, or making changes to the ordered list of events to make it a logical sequence (e.g., deleting events or re-ordering them). The authors report no statistically significant difference in the number of changes made to an automatically generated story compared to the number of changes made to a human-generated story. Important notions in the construction of a story plot are the logical ordering of events (a criminal has to be arrested before a conviction can occur) and the mutual exclusiveness of events (one cannot simultaneously be dead and alive). While some authors include the order of events [6], we currently do not take it into account.

In this paper, we focus on the user interface for story generation systems. Most similar interfaces focus on the final visualisation of the story as a game or a visual element but not on a graphical user interface for the, in a general sense, content management. Ma et al. [16] present several alternative visualisation scenarios for stories. Similarly, Segel and Heer [26] demonstrate different visualisation options for data in a narrative way. What they do not present, however, is an interface for the management of the content for these stories. Tanahashi and Ma [29] provide a detailed description of an approach for arranging events

onto time lines, not taking into account the management of the content (events) itself.

Kybartas and Bidarra [14] present two different interfaces for semantic storytelling, both different to our approach. They concentrate on novel visual interactivity methods for end users. Their first interface, Smart Storybook, is a storytelling prototype designed specifically to enable children to experience storytelling in a creative and exploratory way, while the second, Improv Game is the prototype of a casual improvisational multi-player storytelling game, in which players take turns in building segments of a story.

In addition to these research activities, the area of Content Management Systems is of relevance to our project. With regard to our goals, a web-based solution for curating smaller content pieces or atomic content modules would be ideal. There are, of course, many Content Management Systems available on the market (e.g., Joomla, Drupal, Wordpress, Typo3 and many others), that all have different sets of features. What they lack, however, is the possibility of handling and managing semantic information. The core of our project, curation technologies and curation services, are not yet a typical feature or set of features of Content Management Systems.

The results presented by Mulholland et al. [17] are similar to the idea of story content management that we have developed. They describe an ontology for the definition and curation of stories in museum exhibits. Besides the ontology, it also presents a graphical interface, based on the CMS Drupal, in which the events that make up a story can be managed. Events can be added, deleted or modified. In our approach, we manage events in a similar way, but we cannot import them from external sources yet (this functionality is under development). However, we offer users the option of ranking events (positive, negative) and make the underlying engine aware of these positive or negative scores.

The most closely related approach is the one developed by Poulakos et al. [19], which presents “an accessible graphical platform for content creators and even end users to create their own story worlds, populate it with smart characters and objects, and define narrative events that can be used by existing tools for automated narrative synthesis”. Its functionality is divided in three parts: Story World Creation, Smart Object Creation and Event Creation. The first part defines the states and relationships in the story world, the second defines the objects and characters of the story, and the third defines the events that create the story. The generation of new events is similar to our approximation, because it is also done in a new window where values are assigned to the properties of the event. Conversely, their consideration of event is wider than ours, because they regard an event as an action happening in the story, while we currently still consider events as semantic triples.

3 Architecture and System Description

In this section we briefly describe the current prototype of our interactive, semi-automatic semantic storytelling system. It can be used to analyse document

collections of almost arbitrary sizes, ranging from a few files to collections comprising millions of documents, in order to generate story paths that can be further explored interactively. The system is still under development. We re-arrange and present the information extracted through the semantic analysis in a structured way. This is why we refer to our approach as semantic storytelling. Its foundation is the digital curation platform also developed in the project [2,5,25]. The platform offers multiple curation services [18,20], linked data extraction [5], machine translation [28] and natural language processing [3,4].

In addition to the storytelling system based on several NLP components (Named Entity Recognition, temporal expression analysis, linked data integration etc.), a second key contribution of this article is a user interface (UI) that allows the interaction with potential story paths to modify their content and structure and the user interaction mechanisms. The system architecture is shown in Fig. 1. It consists of three main parts, the NLP engine for analysing the documents, the storytelling module, that is responsible for creating story paths out of collections and the user interface (UI), which is handling the interaction between users and the system.

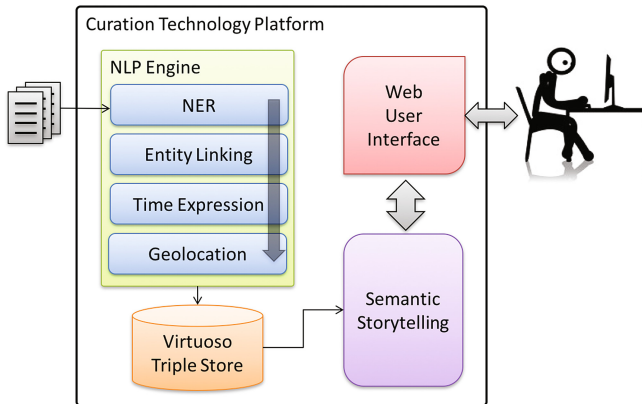


Fig. 1. Architecture of the interactive semantic storytelling engine

3.1 Natural Language Processing Engine

Our platform consists of a set of NLP modules to analyse documents and extract information to be used in several digital curation scenarios, among others, for generating stories. Interoperability between the modules is realised through the NLP Interchange Format (NIF) [24], i.e., all modules accept NIF as input and return NIF as output. The shared usage of NIF allows for the combination of web services in a decentralised way without hard-wiring specific workflows or

pipelines. In the following we describe the modules to put together a semantic analysis workflow.³

Named Entity Recognition. First we convert every document to NIF and then perform Named Entity Recognition (NER). The NER module currently consists of two different approaches (additional ones to be added as part of future work). If annotated data is available, we train models to use in the NER task to deal with disambiguation and spotting of previously unseen entities. If we have a lexicon, dictionary or list of words (as is the case in one of our SME partner's use cases and data sets), we can use the dictionary for lexicon-based proper noun identification (with limited mechanisms for disambiguation). Both approaches are based on OpenNLP modules.⁴ For entities that were spotted using the trained model approach, we proceed to look up the entity on its (language-specific) DBpedia page using DBpedia spotlight [27]. This allows us to retrieve relevant information using a SPARQL query. For our current use cases and domains we use domain-specific dictionaries with names and URIs, i.e., we upload this dictionary in a key-value format, enabling the annotation of a spotted entity with the URI as additional information. In the case of locations the system points to Geonames URIs.⁵ We use a SPARQL query against the Geonames ontology to retrieve the latitude and longitude of entities of the type location that we identify in the text. In the case of persons and organisations, our system points to URIs at Deutsche Nationalbibliothek⁶. As document-level statistics, we calculate the average latitude and longitude points and corresponding standard deviations.

Entity Linking. The next step within the workflow is Entity Linking. As input it receives the text annotated with entities. These entities are looked up in an ontology to retrieve, if available, a URI, depending on the presence of the entity in the ontology. This module was also implemented using DBpedia Spotlight.

Temporal Expression Analysis. The temporal expression analyser uses a regular expression based grammar that is currently able to process German and English natural language text. After the identification of temporal expressions, these are normalised to a shared machine readable format and added to the NIF representation. In this step we also add document-level statistics based on the normalised temporal values. The specific value is the mean date range; it is a range consisting of the average date in the document minus the standard deviation as the starting point, and the average date in the document plus the standard deviation as the end point of the range.

³ Interested readers are invited to experiment with the API: <http://dev.digitale-kuratierung.de/api/e-nlp/namedEntityRecognition>. The documentation is available on Github: <https://github.com/dkt-projekt/e-OpenNLP/tree/master-architecture-update>.

⁴ <https://opennlp.apache.org>.

⁵ <http://www.geonames.org>.

⁶ <http://www.dnb.de>.

Geographical Localisation Module. The geographical location module uses a SPARQL query to retrieve the latitude and longitude of the location as specified in the DBpedia entry (see above). This information is then added to the entity annotation in NIF. After the identification of geolocation entities, the module calculates the mean and standard deviation value for latitude and longitude of all location entities of a document for which a DBpedia entry could be found. This information is also added to the NIF representation.

Annotated Final Document. An example of an annotated document after processing with the NLP engine is shown below. The red text denotes the annotated entities, the green text denotes the annotated temporal expressions, the blue text denotes the information added by the entity linking module and the orange text shows the information added by the geolocation module

```
@prefix rdf: <http://www.w3.org/1999/02/22-rdf-syntax-ns#>.
@prefix xsd: <http://www.w3.org/2001/XMLSchema#>.
@prefix itsrdf: <http://www.w3.org/2005/11/its/rdf#>.
@prefix nif: <http://persistence.uni-leipzig.org/nlp2rdf/
  ontologies/nif-core#>.
@prefix rdfs: <http://www.w3.org/2000/01/rdf-schema#>.
@prefix geo: <http://www.w3.org/2003/01/geo/wgs84_pos/>.
<http://link.omitted/documents/document1#char=0,26>
  a      nif:RFC5147String , nif:String , nif:Context;
  nif:beginIndex      "0"^^xsd:nonNegativeInteger;
  nif:endIndex      "26"^^xsd:nonNegativeInteger;
  nif:isString      "Welcome to Berlin in 2016."^^xsd:string;
  dfkinif:averageLatitude      "52.516666666666666"^^xsd:double;
  dfkinif:averageLongitude      "13.383333333333333"^^xsd:double;
  dfkinif:standardDeviationLatitude      "0.0"^^xsd:double;
  dfkinif:standardDeviationLongitude      "0.0"^^xsd:double;
  nif:meanDateRange      "201601010000_201701010000"
    ^^xsd:string.
<http://link.omitted/documents/document1#char=21,25>
  a      nif:RFC5147String , nif:String;
  nif:anchorOf      "2016"^^xsd:string;
  nif:beginIndex      "21"^^xsd:nonNegativeInteger;
  nif:endIndex      "25"^^xsd:nonNegativeInteger;
  nif:entity      <http://link.omitted/ontologies/nif#date>;
  itsrdf:taIdentRef      <http://link.omitted/ontologies/nif#date=
    20160101000000_20170101000000>.
<http://link.omitted/documents/#char=11,17>
  a      nif:RFC5147String , nif:String;
  nif:anchorOf      "Berlin"^^xsd:string;
  nif:beginIndex      "11"^^xsd:nonNegativeInteger;
  nif:endIndex      "17"^^xsd:nonNegativeInteger;
  itsrdf:taClassRef      <http://dbpedia.org/ontology/Location>;
  nif:referenceContext      <http://link.omitted/documents
    /#char=0,26>;
```

```

geo:lat           "52.51666666666666"^^xsd:double;
geo:long          "13.383333333333333"^^xsd:double;
itsrdf:taIdentRef <http://dbpedia.org/resource/Berlin>.

```

3.2 Towards Semantic Storytelling

The storytelling module makes use of the analysis information to generate potential story paths. This generation is based on templates. Templates are JSON files defining the fields that compose every event involved in the story. These fields also define the type of the entities involved in the events. These can be constructed in two different ways, i.e., through manual definition or through the semi-automatic generation based on events from previously annotated documents. The manually defined templates are used as an initial point for the semantic storytelling and they include the most frequent events (entities and relations) mentioned in the collections of documents that we had available from our SME partners. The semi-automatic approach, which is currently under development, takes a collection of documents annotated with entities, relations (verbs) and events and infers a template structure. This inference is made using clustering algorithms to group the events and filtering based on frequency of appearance. Using clusters we can classify those events belonging to different parts of a story type.

In the following we present, as an example, the Biography template in a triple-based format, omitting some information for readability.

```

biographyTemplate_V1
  mainEntity: PERSON
  Events:
    PERSON(met)PERSON
    PERSON(born)LOCATION
    PERSON(study)LOCATION
    PERSON(joined)UNKNOWN
    PERSON(moved)LOCATION
    [...]
    PERSON(employed)LOCATION
    PERSON(marry)PERSON
    PERSON(lived)LOCATION
    PERSON(died)LOCATION

```

A story is defined as a structured set of events which contain entities and relations. In the current version of the system, relations are constituted by specific groups of verbs that relate entities with one another.⁷ This generic template is filled using information semantically extracted from the document collection. The first step of the template filling process is searching the most common entity that matches the mainEntity type of the template. For example, in the case of

⁷ This approach will be augmented with a more rhetorical or discourse-oriented approach as part of future work.

Biography the mainEntity is PERSON, so the system retrieves all the entities of type PERSON appearing in the collection and selects the most frequent one. After selecting the main entity, the system iterates over the events in the template and looks for other events matching the main entity and the relation (verb). If there is more than one suitable event, it selects the most frequent. Once the system has tried to fill all the events of the template, it returns the story path containing the information. This approach cannot be considered full Natural Language Generation because the typical last step of surface realisation is not included. Our main goal is to increase the efficiency in knowledge acquisition for human users, not the fully automatic generation of natural language text.

3.3 User Interfaces for Semantic Storytelling

During the runtime of the project we have implemented several different semantic storytelling interfaces. The initial prototype is shown in Fig. 2.

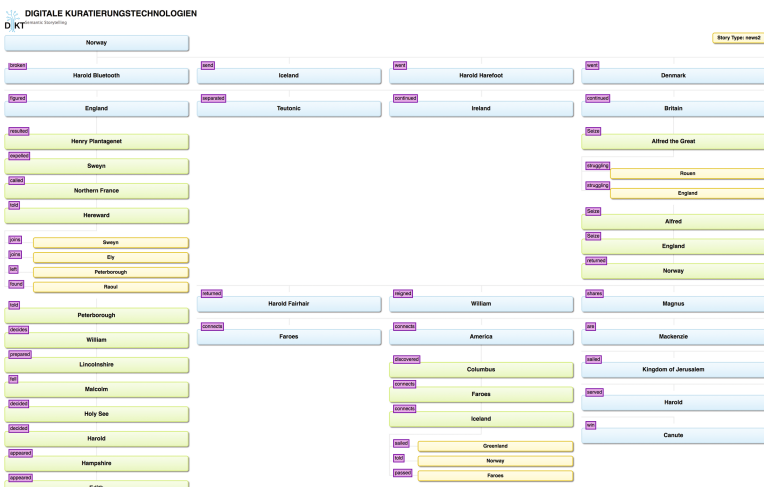


Fig. 2. The initial semantic storytelling prototype

The prototype presents a semantic story using a tree structure. The blue, green and yellow boxes are entities, which are interlinked by relations (specific classes of verbs), shown in the purple boxes. A detailed example can be seen in Fig. 3, which shows three events: Odin \rightarrow *traced* \rightarrow Loki, Loki \rightarrow *chagrined* \rightarrow Balder, Balder \rightarrow *found* \rightarrow Nanna.

This initial prototype interface lacked interactivity, i.e., users could not modify the content or the structure of stories. In order to solve the lack of interactivity, we developed a second version. The new UI was designed to be simple and intuitive. It is composed of three main parts.

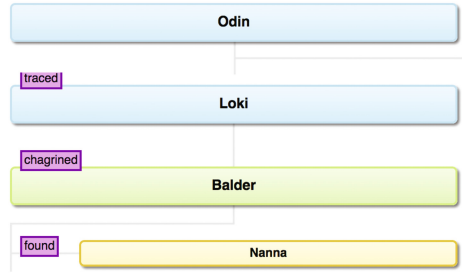


Fig. 3. Detail of the previous figure showing the relations $\text{Odin} \rightarrow \text{traced} \rightarrow \text{Loki}$, $\text{Loki} \rightarrow \text{chagrined} \rightarrow \text{Balder}$ and $\text{Balder} \rightarrow \text{found} \rightarrow \text{Nanna}$

Collection 1
Template 2
Story 3

Step 1: Select a collection

Select a collection where the information will be extracted from.

Available Collections		Files in Collection: mendelsohn			
#	Name	#	Name	Size	Modified
1	collection1	1	mendelsohn_1.nif	4076 bytes	2017-02-06 14:08:07
2	mendelsohn	2	mendelsohn_10.nif	9443 bytes	2017-02-06 14:08:07
3	news	3	mendelsohn_100.nif	7237 bytes	2017-02-06 14:08:07
4	wikings	4	mendelsohn_1000.nif	6975 bytes	2017-02-06 14:08:07
4 collections		5	mendelsohn_1001.nif	6975 bytes	2017-02-06 14:08:07
		6	mendelsohn_1002.nif	9469 bytes	2017-02-06 14:08:07

Copyright © DFKI 2016 Select Collection





Fig. 4. Collection interface showing the content of a collection

1. **Collection Selection Interface** (Fig. 4): this interface allows the selection of an information source (i.e., a collection of documents) to use for generating stories. Users can navigate through the available collections, visualise documents and inspect their content and annotations. Users can also modify the annotated information.
2. **Template Selection Interface** (Fig. 5): this interface allows users to select a story template to be filled with the information of the collection previously selected. Templates are associated with story types. Templates cannot be modified directly in the interface; this feature will be added in a future iteration of the interface.
3. **Story Management Interface** (Fig. 6): this interface allows users to visualise the automatically generated story in a list or tree version.

Collection 1 Template 2 Story 3

Step 2: Select a Template

Available Templates

#	Name	Size	Modified
1	 biographyDE.json	1784 bytes	2017-02-09 11:24:43
2	 biographyEN.json	1616 bytes	2017-02-09 11:24:43
3	 museum_location.json	1002 bytes	2017-02-09 11:29:23
4	 news.json	1284 bytes	2017-02-09 11:24:43

4 templates

Using template: biographyDE.json;

Person

- Person → heiraten → Person [1]
- Person → heiraten → Location [1]
- Person → geboren werden → Location [1]
- Person → sterben → Location [1]
- Person → treffen → Person [1]
- Person → lieben → Person [1]
- Person → arbeiten → Location [1]
- Person → arbeiten → Organisation [1]
- Person → beitreten → Organisation [1]
- Person → besuchen → Location [1]
- Person → besuchen → Person [1]
- Person → kommen → Location [1]
- Person → kommen → Person [1]
- Person → wohnen → Location [1]
- Person → gehen → Person [1]
- Person → gehen → Location [1]



Copyright © DFKI 2016 Select Template






Fig. 5. Template interface showing the content of a template

Collection 1 Template 2 Story 3

Step 3: Review Generated Story

Story generated using Collection: mendelsohn and Template: biographyDE.json

Luise Mendelsohn(PERSON)  

Erich Mendelsohn(PERSON) → treffen(ACTION) → Luise Mendelsohn(PERSON)
 Luise Mendelsohn(PERSON) → lieben(ACTION) → Erich Mendelsohn(PERSON)
 Luise Mendelsohn(PERSON) → besuchen(ACTION) → Hedwig Witkop(PERSON)
 Luise Mendelsohn(PERSON) → besuchen(ACTION) → Lanzkirch(LOCATION)
 Luise Mendelsohn(PERSON) → kommen(ACTION) → Potsdamer Platz(LOCATION)
 Erich Mendelsohn(PERSON) → kommen(ACTION) → Luise Mendelsohn(PERSON)
 Luise Mendelsohn(PERSON) → wohnen(ACTION) → Villa San Giorgio Firenze(LOCATION)     

Copyright © DFKI 2016 Save Download

Fig. 6. Story interface showing the content of an automatically generated story using a specific collection and template

Enabling Interactivity. When the knowledge worker is presented with the generated story, we also provide interactive mechanisms to allow the user to modify the story efficiently. The interactivity was implemented as a set of feedback tools that are attached to every event in the storyline (shown in Fig. 7).

Cassady(PERSON) → born(ACTION) → Salt Lake City(LOCATION)
 Kerouac(PERSON) → joined(ACTION) → United States Navy(UNKNOWN)
 Kerouac(PERSON) → moved(ACTION) → College Park(LOCATION)



Fig. 7. Detail of the feedback tools attached to every event of a generated story

The feedback tools allow the following five interactions.

1. **Confirm** the event in the structure of the story (green thumbs-up button).
2. **Reject** the event of the story and its content (orange thumbs-down button).
3. **Edit** the type or content of the entities (blue pencil button).
4. **Delete** the event and all sub-events (red times-circle button).
5. **Add** a new event to the story (gray plus button)

When one of the feedback buttons is clicked, a new window appears (Fig. 8). This window allows users to modify the type and content (text) of the entities and the relation (verb) involved in the event. Currently, the user has to write manually the entities and relation text and types. We plan to replace these text fields with drop down menus to make available all annotated entities from a collection.

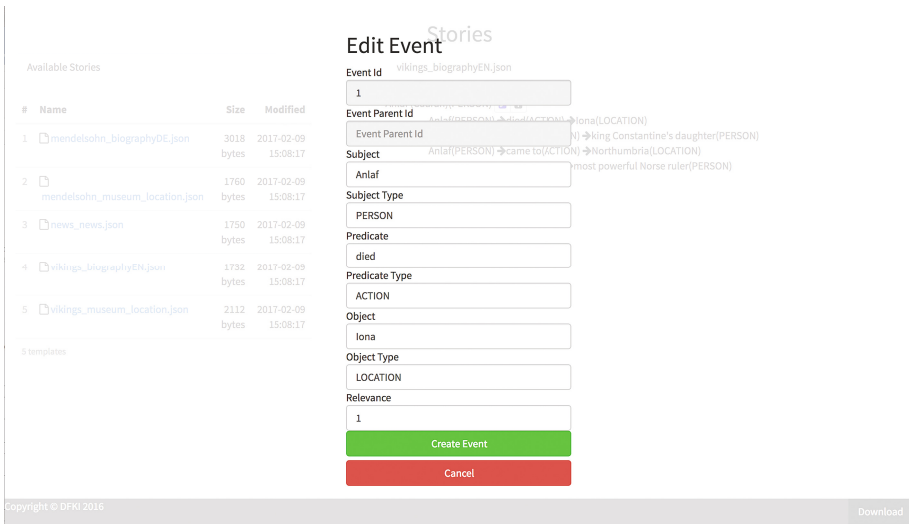


Fig. 8. Detail of the “Edit event” window

4 Experiments

We evaluated the current version of the dashboard using users. We have not yet applied a user-centric evaluation because we are still in the process of defining and shaping the technology behind this new type of interaction interface

and visual metaphor [21]. The goal was to measure the usability of the new visual metaphor. We applied the methods described by Kelly [12], who measures usability in three dimensions: *effectiveness*, *efficiency* and *satisfaction*. In this preliminary study we focus upon satisfaction.

Users were asked to apply the system to create stories based on a variety of document collections and story types. Users were also asked to use the interactive controls: accepting or rejecting information, deleting or adding fields. Finally, users were asked to fill in a questionnaire. We recruited 15 subjects, aged between 23 and 47, with different educational backgrounds (engineering: 7, linguistics: 5, other: 3; 8 female, 7 male). All participants had been using computers for five to ten years. None of them had ever used a system like or similar to our semantic storytelling prototype. The users were asked to apply the system to create stories based on various combinations of a template and one of the following document collections:

- *Mendelsohn Letters*: a collection of 5144 personal letters written by the well known architect Erich Mendelsohn and his wife, Luise. Most of the letters are short. Topics range between architecture, travels and personal events. Language: German.
- *Vikings*: a collection containing seven longer documents (ca. 40,000 words on average). The collection is about Vikings, their history and culture. Language: English.
- *News*: a set of 2352 online news articles in German. Language: English.

Users were asked to apply the following templates:

- *Biography*: this template is characterised by having a main entity of type PERSON and a set of events that are related to the main entity. Because of that, this template is composed by a set of events related to the main entity, such as, *born* → LOCATION, *marry* → PERSON or *work* → ORGANISATION among others.
- *Museum Exhibit*: this template is also characterised by a central information, but contrary to the biography, it is not an entity but a whole topic. An example of museum exhibit covers the topic of Vikings. This exhibit does not only present a single Viking, but the entire Viking civilization, including all the famous Vikings, the locations where they lived, etc. In this sense, a template covering this kind of story is much more generic, and needs much structure than a single list. Because of that, this template contains a tree structure where there are nested events. A nested event is an event where the subject of the nested event is the object of a previous event.
- *News*: like the stories it generates, this template is the most generic and in which it is difficult to define a structure. Therefore, we have tried to keep this template open to facilitate news from different domains and topics. Therefore, this template also organises the events in the form of a tree, but with a peculiarity: events do not have to belong to the same topic, in contrast to the museum exhibit.

The detailed characteristics of the templates are shown in Table 1.

Table 1. The templates used in the evaluation

Name	# Events	Lang	Domain	Description
Biography	16	EN/DE	Biography	Specifies the entities and events that make up a biography
Museum exhibit	9	EN/DE	Single topic	Typical entities and events relevant for a museum exhibit
News	11	EN/DE	General	Typical entities and events in news articles

The workflow of the interface starts by selecting a collection and a template. Once these are selected, an automatically generated storyline is presented to the user in the interface shown in Fig. 6. Users are asked to verify the extracted storyline in two different ways: first, to check if the structural elements (entities and events) and their organisation is correct and, second, check if the content (text and URL) of the elements (events) is correct. Users are asked to explore all possibilities offered by the platform. After completing the task, users were asked to fill in a questionnaire with ten questions, allowing to score the respective system feature with a numeric value (between 0 and 10), and also to provide comments (see Table 2).

Table 2. The post-task questionnaire

	Question	Evaluated feature
Q1	How easy was it to generate a story using the interface?	Story generation
Q2	How intuitive was it to select a collection?	Collection selection
Q3	How intuitive was it to select a template?	Template selection
Q4	How easy was it to positively/negatively rank an event?	Event ranking
Q5	How easy was it to add new events to the story?	Add event
Q6	How easy was it to edit an event?	Edit event
Q7	How easy was it to delete an event?	Delete event
Q8	Inside the new event window, how easy was it to select information for filling the event properties (predicate, object, etc.)?	New event window
Q9	Inside the edit event window, how easy was it to change the information contained in the event (predicate, object, etc.)?	Edit event window
Q10	Was it intuitive to use the download story button?	Download button

The numerical scores were used to assess users' perception of a certain interface feature (Fig. 9). The free text comments were analysed manually and qualitatively to draw conclusions about users' perceptions.




Fig. 9. Average and standard deviation of the scores assigned to the ten questions

It can be seen that users value the interface in a positive way (average value of 8.2), although there is also some disparity with a high standard deviation (1.75). The selection of collections and templates (Q2, Q3) have been positively scored (8.25 and 8.375) by all users (standard deviation: 0.7 and 0.74). As regards the event management buttons (Q4–Q7) the average values are high (8.25, 7.75, 9.13, 9.5) and the standard deviations are small (1.6, 0.83, 0.75) with only one exception in the ranking button (3.1). This indicates a high degree of general satisfaction among users. In contrast, the event management window (Q8 and Q9) and the download button (Q10) get lower scores on average (5.9, 6.6, 7.6). The qualitative analysis of the comments of the questions provides interesting insights:

- There was a general feeling that the interface was not finished yet. This makes sense since the full integration of the semantic storytelling prototype into the curation dashboard [21] is still pending.
- The current graphics and visual design of the interface was commented upon (“it is a little 90 s style”).
- Negative remarks arose from the difficulty of including the text of the entities and relations or of the types, by not having access to the content of the documents of the collection.
- The high standard deviation in Q4 is due to a very low value given by a user who specified that “using the button was easy, but the internal functionality of the button was not clear”.
- Another general comment was the lack of information about what the functionality of each button was.

These results will be taken into account in the next iteration of the user interface [21].

5 Conclusions

In this paper we present the first prototype of a novel type of application: semantic storytelling. This curation service is developed in the context of a research and technology transfer project on Digital Curation Technologies. The approach is based on the semantic analysis of document collections. With this semantic layer in place, we attempt to identify interesting, surprising, eye-opening relationships between different concepts or entities mentioned in the document collection.

Especially with regard to the discourse level analytics, the current approach at relation extraction must be considered rather simplistic in nature, especially when compared to more knowledge-intensive or deep Natural Language Generation approaches. This was a conscious technical design decision. One of the key ideas behind our approach is to explore the boundary of what is possible, technically feasible and implementable, in a generic way, in the curation technology platform and what will remain, for the time being, more in the realm of experimental software [21].

With regard to the evaluation, we have gained several conclusions. The three main findings are that users find the semantic storytelling prototype interesting and, generally, useful, but many improvements need to be applied before we can consider the approach a final product, rather than an evaluation interface or experimental prototype. Second, the “edit event” window is difficult to use. It is not intuitive and the information that can be included in entities and relations is not clear, which suggests a redesign of the approach. The visual and graphics design (including colours, typefaces etc.) of the prototype needs to be updated.

The approach to a semantic storytelling interface presented in this paper has created very useful results for defining the next steps and several novel future approaches that we had not considered previously. The most important one is that the storytelling interface can not function as a stand-alone UI, i.e., it has to be integrated into a larger curation framework or dashboard, so that users can develop a much better and deeper understanding of the information that has been or can be included in the storytelling interface and how it was processed.

Acknowledgments. The project “Digitale Kuratierungstechnologien” (DKT) is supported by the German Federal Ministry of Education and Research (BMBF), “Unternehmen Region”, instrument Wachstums-kern-Potenzial (no. 03WKP45). More information: <http://www.digitale-kuratierung.de>.

References

1. Belz, A.: Automatic generation of weather forecast texts using comprehensive probabilistic generation-space models. *Nat. Lang. Eng.* **14**(4), 431–455 (2008). <http://dx.doi.org/10.1017/S1351324907004664>
2. Bourgonje, P., Moreno-Schneider, J., Nehring, J., Rehm, G., Sasaki, F., Srivastava, A.: Towards a platform for curation technologies: enriching text collections with a semantic-web layer. In: Sack, H., Rizzo, G., Steinmetz, N., Mladeníć, D., Auer, S., Lange, C. (eds.) *ESWC 2016. LNCS*, vol. 9989, pp. 65–68. Springer, Cham (2016). doi:[10.1007/978-3-319-47602-5_14](https://doi.org/10.1007/978-3-319-47602-5_14)

3. Bourgonje, P., Schneider, J.M., Rehm, G.: Digital curation technologies for forensic linguistics. In: 13th Biennial Conference of the International Association of Forensic Linguists, Porto, Portugal, July 2017 (in Print)
4. Bourgonje, P., Schneider, J.M., Rehm, G.: Semantically Annotating heterogeneous Document Collections - Curation technologies for digital humanities and text analytics. In: CUTE Workshop, 2017 - CRETA Unshared Task zu Entitätenreferenzen. Workshop bei DHd2017, Berne, Switzerland., February 2017 (in Print)
5. Bourgonje, P., Schneider, J.M., Rehm, G., Sasaki, F.: Processing document collections to automatically extract linked data: semantic storytelling technologies for smart curation workflows. In: Gangemi, A., Gardent, C. (eds.) Proceedings of the 2nd International Workshop on Natural Language Generation and the Semantic Web (WebNLG 2016), pp. 13–16. The Association for Computational Linguistics, Edinburgh (2016)
6. Chambers, N., Jurafsky, D., Manning, C., Ng, A.: Inducing Event Schemas and Their Participants from Unlabeled Text. Computer Science Department, Stanford University (2011). <https://books.google.de/books?id=3cc3yEErDX8C>
7. Cimiano, P., Lüker, J., Nagel, D., Unger, C.: Exploiting ontology lexica for generating natural language texts from RDF data. In: Proceedings of the 14th European Workshop on Natural Language Generation, pp. 10–19. Association for Computational Linguistics, Sofia, Bulgaria, August 2013. <http://www.aclweb.org/anthology/W13-2102>
8. Dale, R.: Cooking Up Referring Expressions. In: Proceedings of the 27th Annual Meeting on Association for Computational Linguistics, ACL 1989, pp. 68–75. Association for Computational Linguistics, Stroudsburg (1989). <http://dx.doi.org/10.3115/981623.981632>
9. Gervás, P.: Stories from games: content and focalization selection in narrative composition. In: I Spanish Symposium on Entertainment Computing. Universidad Complutense de Madrid, Madrid, September 2013
10. Gervás, P., Díaz-Agudo, B., Peinado, F., Hervás, R.: Story plot generation based on CBR. In: Macintosh, A., Ellis, R., Allen, T. (eds.) Applications and Innovations in Intelligent Systems XII, pp. 33–46. Springer, London (2005). http://dx.doi.org/10.1007/1-84628-103-2_3
11. Goldberg, E., Driedger, N., Kittredge, R.I.: Using natural-language processing to produce weather forecasts. *IEEE Expert Intell. Syst. Appl.* **9**(2), 45–53 (1994)
12. Kelly, D.: Methods for evaluating interactive information retrieval systems with users. *Found. Trends Inf. Retr.* **3**(1–2), 1–224 (2009)
13. Kybartas, B., Bidarra, R., Eisemann, E.: Integrating semantics and narrative world generation. In: Proceedings of FDG 2014 - Ninth International Conference on the Foundations of Digital Games, Fort Lauderdale, FL, April 2014. <http://graphics.tudelft.nl/Publications-new/2014/KBE14>
14. Kybartas, B., Bidarra, R.: A semantic foundation for mixed-initiative computational storytelling. In: Schoenau-Fog, H., Bruni, L.E., Louchart, S., Baceviciute, S. (eds.) ICIDS 2015. LNCS, vol. 9445, pp. 162–169. Springer, Cham (2015). doi:10.1007/978-3-319-27036-4_15
15. Li, B., Lee-Urban, S., Johnston, G., Riedl, M.O.: Story generation with crowd-sourced plot graphs. In: Proceedings of the Twenty-Seventh AAAI Conference on Artificial Intelligence, AAAI 2013, pp. 598–604. AAAI Press (2013). <http://dl.acm.org/citation.cfm?id=2891460.2891543>
16. Ma, K.L., Liao, I., Frazier, J., Hauser, H., Kostis, H.N.: Scientific storytelling using visualization. *IEEE Comput. Graph. Appl.* **32**(1), 12–19 (2012)

17. Mulholland, P., Wolff, A., Collins, T.: Curate and storyspace: an ontology and web-based environment for describing curatorial narratives. In: Simperl, E., Cimiano, P., Polleres, A., Corcho, O., Presutti, V. (eds.) *ESWC 2012*. LNCS, vol. 7295, pp. 748–762. Springer, Heidelberg (2012). doi:[10.1007/978-3-642-30284-8_57](https://doi.org/10.1007/978-3-642-30284-8_57)
18. Neudecker, C., Rehm, G.: Digitale Kuratierungstechnologien für Bibliotheken. *Zeitschrift für Bibliothekskultur* **4**(2), 48–59 (2016). <http://0277.ch/ojs/index.php/cdrs.0277/article/view/158>
19. Poulakos, S., Kapadia, M., Schüpfer, A., Zünd, F., Sumner, R., Gross, M.: Towards an accessible interface for story world building. In: *AAAI Conference on Artificial Intelligence and Interactive Digital Entertainment*, pp. 42–48 (2015). <http://www.aaai.org/ocs/index.php/AIIDE/AIIDE15/paper/view/11583>
20. Rehm, G.: Flexible Digitale Kuratierungstechnologien für verschiedene Branchen und Anwendungsszenarien. In: Bienert, A., Flessler, B. (eds.) *EVA Berlin 2016 - Elektronische Medien & Medien, Kultur, Historie*, Berlin, Germany, pp. 19–22, 23. Berliner Veranstaltung der internationalen EVA-Serie. 09–11 November (2016)
21. Rehm, G., He, J., Schneider, J.M., Nehring, J., Quantz, J.: Designing user interfaces for curation technologies. In: *19th International Conference on Human-Computer Interaction - HCI International 2017*, Vancouver, Canada, July 2017 (in Print)
22. Reiter, E., Sripada, S., Hunter, J., Davy, I.: Choosing words in computer-generated weather forecasts. *Artif. Intell.* **167**, 137–169 (2005)
23. Riedl, M.O., Young, R.M.: Narrative planning: balancing plot and character. *J. Artif. Int. Res.* **39**(1), 217–268 (2010). <http://dl.acm.org/citation.cfm?id=1946417.1946422>
24. Sasaki, F., Gornostay, T., Dojchinovski, M., Osella, M., Mannens, E., Stoitsis, G., Richie, P., Declerck, T., Koidl, K.: Introducing FREME: deploying linguistic linked data. In: *Proceedings of the 4th Workshop of the Multilingual Semantic Web, MSW 2015* (2015)
25. Schneider, J.M., Bourgonje, P., Nehring, J., Rehm, G., Sasaki, F., Srivastava, A.: Towards semantic story telling with digital curation technologies. In: Birnbaum, L., Popescuk, O., Strapparava, C. (eds.) *Proceedings of Natural Language Processing meets Journalism - IJCAI-16 Workshop (NLPMJ 2016)*, New York, July 2016
26. Segel, E., Heer, J.: *Narrative Visualization: Telling Stories with Data*, vol. 16, pp. 1139–1148. IEEE Educational Activities Department, Piscataway, NJ, USA, November (2010). <http://dx.doi.org/10.1109/TVCG.2010.179>
27. Spotlight, D.: *DBPedia Spotlight Website* (2016). <https://github.com/dbpedia-spotlight/>
28. Srivastava, A., Sasaki, F., Bourgonje, P., Moreno-Schneider, J., Nehring, J., Rehm, G.: How to configure statistical machine translation with linked open data resources. In: *Proceedings of Translating and the Computer 38 (TC38)*, pp. 138–148, London, UK, November 2016
29. Tanahashi, Y., Ma, K.-L.: Design considerations for optimizing storyline visualizations. *IEEE Trans. Visual. Comput. Graphics* **18**(12), 2679–2688 (2012). <http://dx.doi.org/10.1109/TVCG.2012.212>
30. Turner, R., Sripada, S., Reiter, E., Davy, I.P.: Generating spatio-temporal descriptions in pollen forecasts. In: *Proceedings of the Eleventh Conference of the European Chapter of the Association for Computational Linguistics: Posters & #38; Demonstrations, EACL 2006*, pp. 163–166. Association for Computational Linguistics, Stroudsburg (2006). <http://dl.acm.org/citation.cfm?id=1608974.1608998>

31. Turner, S.: *The Creative Process: A Computer Model of Storytelling and Creativity*. Taylor & Francis (2014). <https://books.google.gr/books?id=1AjsAgAAQBAJ>
32. Wood, M.D.: Exploiting semantics for personalized story creation. In: Proceedings of the 2008 IEEE International Conference on Semantic Computing, ICSC 2008, pp. 402–409 (2008). <http://dx.doi.org/10.1109/ICSC.2008.10>