

EventKG - the Hub of Event Knowledge on the Web - and Biographical Timeline Generation

Simon Gottschalk^a and Elena Demidova^{a,*}

^a *L3S Research Center, Leibniz Universität Hannover, Hannover, Germany*

E-mails: gottschalk@L3S.de, demidova@L3S.de

Abstract. One of the key requirements to facilitate the semantic analytics of information regarding contemporary and historical events on the Web, in the news and in social media is the availability of reference knowledge repositories containing comprehensive representations of events, entities and temporal relations. Existing knowledge graphs, with popular examples including DBpedia, YAGO and Wikidata, focus mostly on entity-centric information and are insufficient in terms of their coverage and completeness with respect to events and temporal relations. In this article we address this limitation, formalise the concept of a temporal knowledge graph and present its instantiation - EventKG. EventKG is a multilingual event-centric temporal knowledge graph that incorporates over 690 thousand events and over 2.3 million temporal relations obtained from several large-scale knowledge graphs and semi-structured sources and makes them available through a canonical RDF representation. Whereas popular entities often possess hundreds of relations within a temporal knowledge graph such as EventKG, generating a concise overview of the most important temporal relations for a given entity is a challenging task. In this article we demonstrate an application of EventKG to biographical timeline generation, where we adopt a distant supervision method to identify relations most relevant for an entity biography. Our evaluation results provide insights on the characteristics of EventKG and demonstrate the effectiveness of the proposed biographical timeline generation method.

Keywords: Events, Knowledge Graph, Biographical Timelines

1. Introduction

Motivation: The amount of event-centric information regarding contemporary and historical events of global importance, such as the US elections, the 2018 Winter Olympics and the Syrian Civil War, constantly grows on the Web, in the news sources and within social media. In the literature, an event is typically described as something that happens at a specific time and location [1]. Events considered in this work are real-world happenings of societal importance, with examples including military conflicts, sports tournaments and political elections. In particular, we consider events, entities they involve and temporal relations - i.e. real-world relations between events and entities valid over a time period.

Efficiently accessing and analysing large-scale event-centric and temporal information is crucial for a variety of real-world applications in the fields of Semantic Web, NLP and Digital Humanities. In Semantic Web and NLP, these applications include timeline generation [2, 3] and Question Answering [4]. In Digital Humanities, multilingual event repositories can facilitate cross-cultural studies analysing language-specific and community-specific views on historical and contemporary events (examples of such studies can be seen in [5, 6]). Furthermore, event-centric knowledge graphs can facilitate the reconstruction of histories as well as networks of people and organisations over time [7, 8]. One of the pivotal pre-requisites to facilitate effective analytics of events is the availability of knowledge repositories providing reference information regarding events, involved entities and their temporal relations.

*Corresponding author. E-mail: demidova@L3S.de.

Limitations of the existing sources of event-centric and temporal information: Currently, event representations and temporal relations are spread across heterogeneous sources. First, large-scale knowledge graphs (KGs) (i.e. graph-based knowledge repositories [9] such as Wikidata [10], DBpedia [11] and YAGO [12]) typically focus on entity-centric knowledge. Event-centric information included in these sources is often not clearly identified as such, can be incomplete and is mostly restricted to named events and encyclopedic knowledge.

For example, as discussed later in Section 5.1, out of 322,669 events included in EventKG V1.1, only 18.7% are classified using the `dbo:Event` class in the English DBpedia as of 12/2017. Furthermore, event descriptions in the existing knowledge graphs often lack the key properties, i.e. times and locations. For example, according to our analysis, only 33% of events in Wikidata provided temporal and 11.7% spatial information at that time.

Second, a variety of manually curated semi-structured sources (e.g. the Wikipedia Current Events Portal (WCEP) [13] and multilingual Wikipedia event lists) contain information on contemporary events. However, the lack of structured representations of events and temporal relations in these sources hinders their direct use in real-world applications, e.g. through semantic technologies. Overall, a comprehensive integrated view on contemporary and historical events and their temporal relations is still missing. EventKG will help to overcome these limitations.

An additional source of event-centric information on the Web are the recently proposed knowledge graphs containing events obtained from unstructured news sources using Information Extraction methods (such as [7, 14–17]). These knowledge graphs are potentially highly noisy (e.g. [7] reports an extraction accuracy of only 0.55). Due to significant differences in quality and event granularity, the integration of events from these sources with the information in the established knowledge repositories such as DBpedia or Wikidata within a common knowledge graph does not appear meaningful. These event sources as well as the corresponding Information Extraction methods for unstructured news articles are out of scope of this work.

A temporal knowledge graph & EventKG: In this article we formalise the concept of a temporal knowledge graph that interconnects real-world entities and events using temporal relations valid over a time period. Furthermore, we present an instantiation of a temporal knowledge graph - EventKG. EventKG

takes an important step to facilitate a global view on events and temporal relations currently spread across entity-centric knowledge graphs and manually curated semi-structured sources. EventKG integrates this knowledge in an efficient light-weight fashion, enriches it with additional features such as indications of relation strengths and event popularity, adds provenance information and makes all this information available through a canonical RDF representation. Through the light-weight integration and fusion of event-centric and temporal information from different sources, EventKG enables to increase coverage and completeness of this information. For example, EventKG increases the coverage of locations and dates for Wikidata events it contains by 14.43% and 17.82%, correspondingly (see Table 9 in Section 5.1 for more detail). Furthermore, relation strengths and event popularity provided by EventKG are the characteristics that gain the key relevance given the rapidly increasing amount of event-centric and temporal data on the Web and the resulting information overload.

EventKG was first introduced in [18]. Compared to [18], in this article we formally introduce the concept of a temporal knowledge graph, provide more details on the algorithms adopted for the EventKG generation and the corresponding evaluation results. Furthermore, we present a method that facilitates an application of EventKG to biographical timeline generation. We make EventKG, including the dataset, a SPARQL endpoint, the code and evaluation data, as well as the benchmarks created for the biographical timeline generation available online¹.

Generation of Biographical Timelines using a Temporal Knowledge Graph: A popular entity such as an influential person, a city or a large organisation can impose hundreds of temporal relations within a temporal knowledge graph. For example, the entity Barack Obama possesses 2,608 temporal relations in EventKG. Identifying the most important temporal relations within the temporal knowledge graph to provide a concise overview for a given entity becomes a challenging task in these settings.

Timelines are an effective method to provide a visual overview of entity-centric temporal information, such as temporal relations in a knowledge graph [2]. In particular, biographical timelines describe significant happenings in a person's life and typically include events of major relevance from the personal perspec-

¹<http://eventkg.13s.uni-hannover.de/>

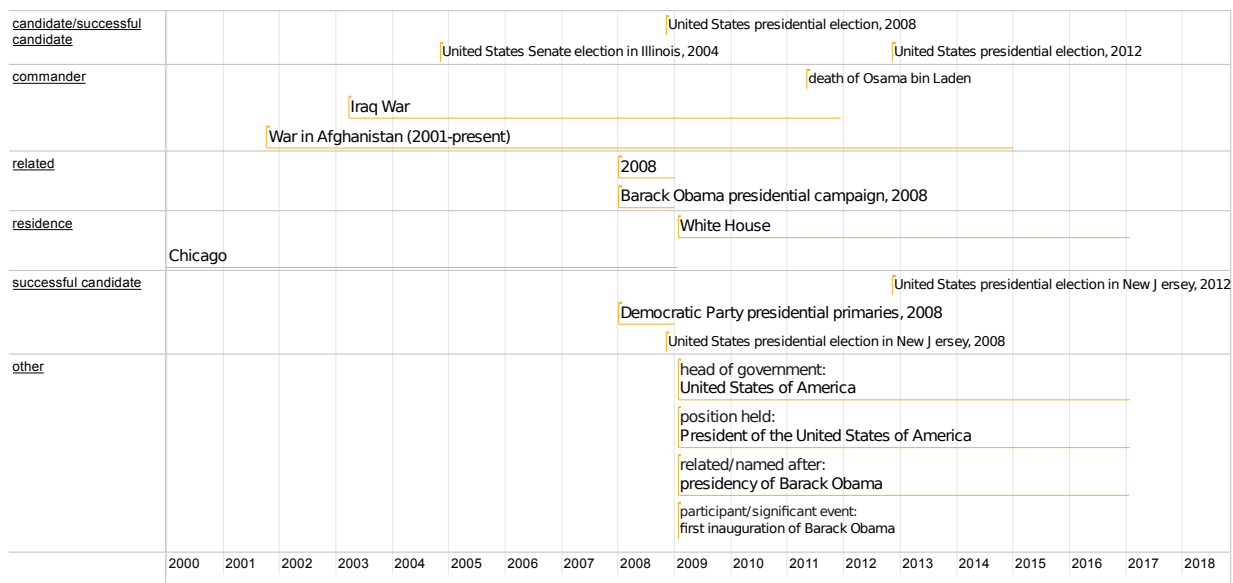


Figure 1. An excerpt of the biographical timeline for the entity Barack Obama, generated from the EventKG knowledge graph using a proposed model trained on the Wikipedia abstracts of other entities (BS-ENC benchmark). Orange lines represent the temporal validity of the relations. Each row corresponds to a predicate characterising the relation (e.g. *commander*) to the specific event or entity (e.g. *Iraq War*).

tive such as birth, education and career. Figure 1 illustrates a biographical timeline for Barack Obama, which includes places where Barack Obama lived (first Chicago and then the White House), important events he was involved in (e.g. the Iraq War) and the major political positions he held (e.g. the President of the United States). This timeline also indicates the temporal validity of these relations.

In this article we present an approach for the generation of biographical timelines from a temporal knowledge graph. To generate such timelines, we propose a distant supervision method, where we train the relevance model using external sources containing biographical and encyclopedic texts. With that model, we extract the most relevant biographical data from the temporal knowledge graph concisely describing a person's life, while using features such as relation strength and event popularity information contained in EventKG, as well as predicate labels. The results of our user evaluation demonstrate that this approach is able to generate high quality biographical timelines while significantly outperforming a state-of-the-art baseline for timeline generation: our timelines were preferred over the baseline's timelines in approximately 68% of the cases.

Overall, our contributions in this article are as follows:

- 1 We formally define the concept of a temporal knowledge graph TKG that incorporates entities, events and temporal relations.
- 2 We present an instantiation of TKG : EventKG - a multilingual RDF knowledge graph that incorporates over 690 thousand events and over 2.3 million temporal relations in V1.1. We provide insights on the extraction and fusion methods adopted to generate the EventKG knowledge graph and their quality.
- 3 We define the problem of biographical timeline generation from a temporal knowledge graph and present our method based on distant supervision.
- 4 We demonstrate the effectiveness of the proposed timeline generation method in a user study.

The remainder of this article is organised as follows: First, in Section 2 we motivate the need for a temporal knowledge graph and introduce a running example. In Section 3, we formally define the concepts of a temporal knowledge graph and a biographical timeline. Then, in Section 4, we describe EventKG, including its RDF data model and the extraction pipeline. In Section 5, we provide statistics and evaluation results of the data contained in EventKG. Our approach towards biographical timeline generation using temporal knowledge graph is presented in Section 6. The experimental setup and evaluation of the biographical timelines generated with our approach using EventKG is

provided in Section 7. Related work is discussed in Section 8. Finally, we discuss our findings and provide a conclusion in Section 9.

2. Motivation

Our society faces an unprecedented number of events that impact multiple communities across language and community borders. In this context, the efficient access to event-centric multilingual information originating from different sources, as facilitated by EventKG, is of utmost importance for several scientific communities, including Semantic Web, NLP and Digital Humanities and a variety of applications, including timeline generation, question answering, as well as cross-cultural and cross-lingual event-centric analytics.

Timeline generation is an active research area [2, 3], where the focus is to generate a timeline (i.e. a chronologically ordered selection) of events and temporal relations for entities from a knowledge graph. In this article we focus on the application of EventKG to the automated generation of timelines representing people biographies. In this task, information regarding event popularity and relation strength available in EventKG in a combination with a benchmark extracted from external biographical sources can enable the selection of the most relevant timeline entries.

EventKG facilitates the generation of detailed timelines containing complementary information originating from different reference sources, potentially resulting in more complete timelines and event representations. For example, Table 1 illustrates an excerpt from the timeline for the query “*What were the events related to Barack Obama between November 4 and November 16, 2011?*” generated using EventKG. The last event in the timeline in Table 1 about Obama visiting Australia extracted from an English Wikipedia event list (“2011 in Australia”²) is not contained in any of the reference knowledge graphs used to populate EventKG (Wikidata, DBpedia, and YAGO). The reference sources of the other two events include complementary information. For example, while the “2011 White House shooting” is assigned a start date in Wikidata, it is not connected to Barack Obama in that source.

An important application of EventKG is cross-cultural and cross-lingual analytics. Such analytics can

provide insights on the differences in the event perception and interpretation across communities. For example, event popularity and relation strength between events and entities varies across different cultural and linguistic contexts. These differences can be observed and analysed using information provided by EventKG. For example, Table 2 presents the top-4 most popular events in the English vs. the Russian Wikipedia language editions as measured by how often these events are referred, i.e. linked to in the respective Wikipedia language edition. Whereas both Wikipedia language editions mention events of global importance, here the two World Wars, most frequently, the other most popular events (e.g. “October Revolution” and “American Civil War”) are language-specific. The relation strength between events and entities in specific language contexts can be inferred by counting their joint mentions in Wikipedia. For example, Table 3 lists the persons most related to the financial crisis in the years 2007 and 2008 in different Wikipedia language editions. This information is directly provided by EventKG. An EventKG application to cross-lingual timeline generation was presented in [3]. In this context, EventKG-empowered interfaces can be used as a starting point to identify events controversial in their cross-cultural aspects. Such events can then be analysed in more detail using tools such as MultiWiki [19] proposed in our previous work.

Another intended future application of EventKG is semantic event-centric question answering. With the provision of EventKG, it becomes possible to answer questions such as “*Which events related to Bill Clinton happened in Washington in 1980?*” and “*What are the most important events related to Syrian Civil War that took place in Aleppo?*” that are of interest for both cross-cultural and cross-lingual event-centric analytics (e.g. illustrated in [6, 20]) as well as question answering and semantic search applications (e.g. [4, 21, 22]).

2.1. Running Example: A Biographical Timeline of Barack Obama

As a running example throughout this article, we will use the task of biographical timeline generation for the entity Barack Obama. First, we will illustrate the heterogeneity of data about Barack Obama available in the reference knowledge graphs used to populate EventKG (Wikidata, DBpedia, YAGO and Wikipedia), and the extraction and integration of this data into a canonical RDF representation in EventKG. As mentioned above, this process leads to 2, 608 temporal

²https://en.wikipedia.org/wiki/2011_in_Australia

Table 1

All events connected with Barack Obama in EventKG that started between November 4 and November 16, 2011.

| Start Date | Sources | Description |
|------------|---|---|
| Nov 4 | YAGO, Wikidata, DBpedia _{EN} , DBpedia _{FR} , DBpedia _{RU} | 2011 G20 Cannes summit |
| Nov 11 | YAGO, Wikidata, DBpedia _{EN} | 2011 White House shooting |
| Nov 16 | Wikipedia _{EN} | The President of the United States Barack Obama visits Australia to commemorate the 60th anniversary of the ANZUS alliance. |

Table 2

Most linked events in the English (EN) and the Russian (RU) Wikipedia.

| Rank | Event (EN) | #Links (EN) | Event (RU) | #Links (RU) |
|------|--------------------|-------------|--------------------|-------------|
| 1 | World War II | 189,716 | World War II | 25,295 |
| 2 | World War I | 99,079 | World War I | 22,038 |
| 3 | American Civil War | 37,672 | October Revolution | 7,533 |
| 4 | FA Cup | 20,640 | Russian Civil War | 7,093 |

Table 3

Top-4 persons mentioned jointly with the financial crisis (2007–2008) per language.

| Rank | EN | FR | DE | RU | PT |
|------|-----------------|--------------|-------------------|-------------------|-------------------|
| 1 | Barack Obama | Kevin Rudd | Barack Obama | Michael Moore | Barack Obama |
| 2 | George W. Bush | John Howard | Geir Haarde | Roman Abramovich | José Sócrates |
| 3 | Joseph Stiglitz | Don Cheadle | George W. Bush | Adam McKay | Pope Benedict XVI |
| 4 | Ben Bernanke | Ben Bernanke | Wolfgang Schäuble | Mikhail Prokhorov | Gordon Brown |

relations involving Barack Obama. In order to generate a biographical timeline of Obama, the relevance of these relations to his biography needs to be assessed. We will describe the distant supervision approach and the features adopted to this task, which finally leads to the timeline depicted in Figure 1.

3. A Temporal Knowledge Graph and Biographical Timelines

A temporal knowledge graph TKG connects real-world entities and events using temporal relations, i.e. relations valid over a time period.

Definition 1. A temporal knowledge graph $TKG : \langle E_t, R_t \rangle$ is a directed multigraph. The nodes in $E_t = E \cup \mathcal{V}$ are temporal entities, where E is a set of real-world entities and \mathcal{V} is a set of real-world events. The directed edges in R_t represent temporal relations of the temporal entities in E_t .

A temporal entity $e \in E$ represents a real-world entity such as a person, a location, an organisation or a concept. A temporal entity $e \in \mathcal{V}$ represents a real-world historical or contemporary event. Examples of

events include cultural, sporting or political happenings. The temporal entities in TKG are characterised through their existence time (for real-world entities) or happening time (for events).

Definition 2. A temporal entity $e \in E_t$ represents a real-world entity or event. e is annotated with a tuple $\langle e_{uri}, e_{time} \rangle$, where e_{uri} is the unique entity identifier and $e_{time} = [e_{start}, e_{end}]$ denotes the existence time of the entity (for $e \in E$) or the happening time of the event (for $e \in \mathcal{V}$).

A temporal entity $e \in E_t$ can be assigned further properties, such as an entity type, a label and a textual description.

A temporal relation is a binary relation of the temporal entities valid over a certain period of time. More formally:

Definition 3. A temporal relation $r \in R_t$ represents a binary relation between two temporal entities. r is annotated with a tuple $\langle r_{uri}, r_{time}, e_i, e_j \rangle$, where r_{uri} is a unique relation identifier, e_i and e_j are the temporal entities participating in the relation r and $r_{time} = [r_{start}, r_{end}]$ denotes the validity time interval of the temporal relation.

The relation identifier r_{uri} reflects the semantics of the temporal relation and is typically specified as a vocabulary term.

Given a temporal knowledge graph $TKG : \langle E_t, R_t \rangle$, we denote the temporal entity of user interest $e \in E_t$ for which the biographical timeline is generated as a *timeline entity*.

A biographical timeline is a chronologically ordered list of temporal relations involving the timeline entity and relevant to that entity's biography.

Definition 4. A biographical timeline $TL(e, bio) = (r_1, \dots, r_n)$ of a timeline entity e is a chronologically ordered list of timeline entries (i.e. temporal relations involving e), where each timeline entry r_i is relevant to the entity biography bio .

In this article, we assume a binary notion of relevance, i.e. $\forall r_i \in TL(e, bio) : relevance(e, r_i, bio) = 1$.

The list of timeline entries in $TL(e, bio)$ is ordered chronologically by their start time: $\forall r_i, r_j \in TL(e, bio) : i \leq j \Leftrightarrow r_{i_{start}} \leq r_{j_{start}}$.

An entity connected to e via a timeline entry r_i is denoted as a *connected entity* in the following.

4. EventKG Knowledge Graph

EventKG is a knowledge graph that instantiates the temporal knowledge graph defined in Definition 1, and at the same time facilitates the integration and fusion of a variety of heterogeneous event representations and temporal relations extracted from several reference sources.

A *reference source* is a semantic source such as a knowledge graph (e.g. Wikidata or YAGO) or a collection of articles (e.g. the French Wikipedia) used to populate EventKG.

In the following, we present the RDF data model of EventKG in Section 4.1 and its transformation into a TKG in Section 4.2. Following that we present the EventKG generation pipeline in Section 4.3 and illustrate the pipeline steps with our running example of Barack Obama in Section 4.4.

4.1. EventKG RDF Data Model

The goals of the EventKG RDF data model are to facilitate a light-weight integration and fusion of heterogeneous event representations and temporal relations extracted from the reference sources, as well as to make this information available to real-world appli-

cations through an RDF representation. The EventKG data model is driven by the following objectives:

- Define the key properties of events through a canonical representation.
- Represent temporal relations between events and entities (including event-entity, entity-event and entity-entity relations).
- Include information quantifying and further describing these relations.
- Represent relations between events (e.g. in the context of event series).
- Support an efficient light-weight integration of event representations and temporal relations originating from heterogeneous sources.
- Provide provenance for the information included in EventKG.

EventKG schema and the Simple Event Model: In EventKG, we build upon the Simple Event Model (SEM) [23] as a basis to model events in RDF. SEM is a flexible data model that provides a generic event-centric framework. The main rationale of SEM is to provide a simple model that can represent events and their key properties. Events within EventKG come from heterogeneous sources where they can be described at a different level of detail. SEM provides the lowest common denominator for event-centric information, whereas it still includes the key properties of events and their relations. The properties of events in the EventKG data model are not mandatory, such that we can also include under-specified events in EventKG, e.g. in case the corresponding temporal or geospatial information is missing in the reference sources. In addition to SEM, within the EventKG schema, we adopt additional properties and classes to adequately represent the information extracted from the reference sources, to model temporal relations and event relations as well as to provide provenance information. The schema of EventKG is presented in Figure 2 and the used RDF namespaces are listed in Table 4.

EventKG is an RDF-based dataset, such that extensions to its data model are easily possible. In future work, such extensions can be performed to model confidence and uncertainty in the information extraction, integration and fusion, or to provide more fine-granular time information (using e.g. EDTF (Extended Date-Time Format) [24]).

Events and entities: SEM provides a generic event representation including topical, geospatial and temporal dimensions of an event, as well as links to its actors

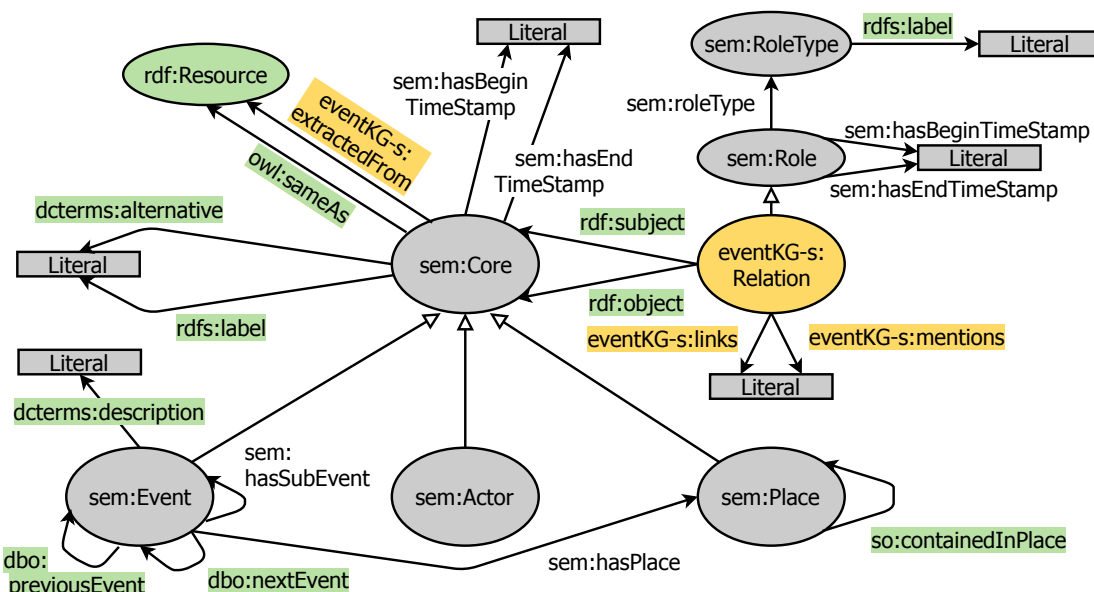


Figure 2. The EventKG schema based on SEM. Arrows with an open head denote `rdfs:subClassOf` properties. Regular arrows visualize the `rdfs:domain` and `rdfs:range` restrictions on properties. Terms from other reused vocabularies are colored green. Classes and properties introduced in EventKG are colored orange.

Table 4
Namespaces used in the EventKG RDF model.

| Namespace prefix | IRI |
|------------------|--|
| so: | http://schema.org/ |
| dbo: | http://dbpedia.org/ontology/ |
| rdf: | http://www.w3.org/1999/02/22-rdf-syntax-ns# |
| rdfs: | http://www.w3.org/2000/01/rdf-schema# |
| dcterms: | http://purl.org/dc/terms/rdfs: |
| sem: | http://semanticweb.cs.vu.nl/2009/11/sem/ |
| eventKG-s: | http://eventKG.l3s.uni-hannover.de/schema/ |
| eventKG-r: | http://eventKG.l3s.uni-hannover.de/resource/ |
| eventKG-g: | http://eventKG.l3s.uni-hannover.de/graph/ |

(i.e. entities participating in the event). Such resources are identified within the namespace `eventKG-r`. Thus, the key classes of SEM and of the EventKG schema are `sem:Event` representing events, `sem:Place` representing locations and `sem:Actor` representing entities participating in the events. Each of these classes is a subclass of `sem:Core`, which is used to represent all entities in the temporal knowledge graph³. Events are connected to their locations through the `sem:hasPlace` property. A `sem:Core` instance can be assigned an existence time denoted via `sem:has-`

³Note that entities in EventKG are not necessarily actors in the events; temporal relations between two entities are also possible.

`BeginTimeStamp` and `sem:hasEndTimeStamp`. In addition to the SEM representation, EventKG provides textual information regarding events and entities extracted from the reference sources including labels (`rdfs:label`), aliases (`dcterms:alternative`) and descriptions of events (`dcterms:description`).

In the context of this article, the term temporal relation refers to real-world relations between events and entities valid over a period of time. The set of temporal relations in EventKG includes event-entity, entity-event and entity-entity relations. Temporal relations between events and entities typically connect an event and its actors (as in SEM). A typical example of a temporal relation between two entities is a marriage. Temporal relations between entities can also indirectly capture information about events [7]. For example, the DBpedia property `http://dbpedia.org/property/` acquired can be used to represent an event of acquisition of one company by another. Temporal relations in SEM are limited to the situation where an actor plays a specific role in the context of an event. This yields two limitations: (i) there is no possibility to model temporal relations between events and entities where the entity acts as a subject. For example, it is not possible to directly model the fact that Barack Obama participated in the event “Second inauguration of Barack Obama”, as the entity “Barack Obama” plays the subject role in this relation; and (ii) a temporal relation between

two entities such as a marriage can not be modelled directly⁴.

To overcome these limitations, EventKG introduces the class `eventKG-s:Relation` representing relations between events and entities. This way of relation modelling facilitates flexible additional attributes describing a relation⁵. This class links two `sem:Core` instances (each representing an event or an entity). The resulting relation can be annotated with a validity time and a property `sem:RoleType` that characterises the relation using RDF predicates. Currently, the predicates are directly derived from the reference sources. In future work, we envision the normalisation of these predicates by mapping them to a dedicated ontology (e.g. the DBpedia ontology⁶). This way, arbitrary temporal relations between entity pairs or relations involving an entity and an event can be represented. This model provides flexibility to express heterogeneous temporal relations derived from the reference sources. Figure 3 visualises the example mentioned above using the EventKG data model.

Other event and entity relations: Relations between events (in particular sub-event, previous and next event relations) play an important role in the context of event series (e.g. Olympic Games), seasons containing a number of related events (e.g. in sports), or events related to a certain topic (e.g. operations in a military conflict). Sub-event relations are modelled using the `so:hasSubEvent` property. To interlink events within an event series such as the sequence of Olympic Games, the properties `dbo:previousEvent` and `dbo:nextEvent` are used. A location hierarchy is provided through the property `so:containedInPlace`.

Towards measuring relation strength and event popularity: Measuring relation strength between events and entities and event popularity enables answering question like “Who were the most important participants of the US Election 2016?” or “What are the most popular events related to the Summer Olympics 2016?”. Relation strength and event popularity are of importance for many practical applications. For example, relation strength can help when using the knowledge graph to jointly disambiguate entities and events

in text documents or in natural language questions in the context of question answering applications. Relation strength and event popularity can also support ranking-based applications including timeline generation and event-centric information retrieval.

Whereas the exact computation of relation strength and event popularity metrics can be application-dependent, we include two major factors required for such computations, namely *links* and *mentions* in the EventKG schema:

- 1. *Links:* This factor represents how often the description of one entity refers to another entity. Intuitively, this factor can be used to estimate the popularity of events and the strength of their relations. In EventKG the links factor is represented through the predicate `eventKG-s:links` in the domain of `eventKG-s:Relation`. `eventKG-s:links` denotes how often the Wikipedia article representing the relation subject links to the entity representing the object.
- 2. *Mentions:* `eventKG-s:mentions` represents the number of relation mentions in external sources. Intuitively, this factor can be used to estimate the relation strength. In EventKG, `eventKG-s:mentions` denotes the number of sentences in Wikipedia that mention both, the subject and the object of the relation.

Links and mentions factors provided by EventKG are computed using sources external to the knowledge graph, such as the entire Wikipedia corpus. Having this information included directly in the knowledge graph can help the relevant applications to obtain this information efficiently and to directly use it in their computations, including (but potentially not limited to) relation strength and event popularity metrics.

Provenance information: EventKG provides the following provenance information: (i) provenance of the individual resources; (ii) representation of the reference sources; and (iii) provenance of statements.

Provenance of the individual resources: EventKG resources typically directly correspond to the events and entities contained in the reference sources (e.g. an entity representing Barack Obama in EventKG corresponds to the DBpedia resource http://dbpedia.org/page/Barack_Obama). In this case, the `owl:sameAs` property is used to interlink both resources. EventKG resources can also be extracted from a resource collection. For example, philosophy events in 2007 can be extracted from the Wikipedia event list at https://en.wikipedia.org/wiki/2007_in_philosophy. In this

⁴Consider the difference between a wedding that is modelled as an event and a marriage between two people that can be modelled as a temporal relation.

⁵See W3C Working Group Note from 12 April 2006 on defining N-ary Relations on the Semantic Web: <https://www.w3.org/TR/swbp-n-aryRelations>.

⁶<https://wiki.dbpedia.org/services-resources/ontology>




Figure 3. Example of the event representing the participation of Barack Obama in his second inauguration as a US president in 2013 as modelled in EventKG. wdt:P793 is the Wikidata identifier for the “significant event” property.

case, the EventKG property `eventKG-s:extracted-From` is utilised to establish the link between the EventKG resource and the resource collection from which this resource was extracted. Through the provenance URIs, background knowledge contained in the reference sources can be accessed.

Representation of the reference sources: EventKG and each of the reference sources are represented through an instance of `void:Dataset`⁷. Such an instance in the namespace `eventKG-g` includes specific properties of the source (e.g. its creation date as in: `eventKG-g:dbpedia_pt dcterms:created "2016-10-01"^^xsd:date`).

Provenance information of statements: A statement in EventKG is represented as a quadruple, containing a triple and a URI of the named graph it belongs to. Through named graphs, EventKG offers an intuitive way to retrieve information extracted from the individual reference sources using SPARQL queries.

4.2. EventKG as a Temporal Knowledge Graph

A named graph such as `eventKG-g:event_kg` can be expressed as a temporal knowledge graph $TKG : \langle E_t, R_t \rangle$ as follows:

- *Entities and events:* Each instance of `sem:Core` is a temporal entity $e \in E_t$ and each instance of `sem:Event` is an event $v \in \mathcal{V}$, such that $E = E_t \setminus \mathcal{V}$ is the set representing real-world entities.
- *Time information for entities and events:* For each temporal entity $e = \langle e_{uri}, e_{time} \rangle, e \in E_t, e_{uri}$ is the URI of the corresponding EventKG entity. e_{start}

and e_{end} are set according to the `sem:hasBeginTimeStamp` and `sem:hasEndTimeStamp` values in the `eventKG-g:event_kg` named graph, correspondingly.

- *Temporal relations with known validity times:* Each instance of `eventKG-s:Relation` that has a start or an end time in the named graph is transformed into a temporal relation $r = \langle r_{uri}, r_{time}, e_i, e_j \rangle \in R_t$. Here, r_{uri} is the URI of the EventKG relation instance, e_i is the entity connected to the `eventKG-s:Relation` instance via `rdf:subject`, e_j is the entity connected via `rdf:object` and r_{time} includes the `sem:hasBeginTimeStamp` and `sem:hasEndTimeStamp` relations.
- *Indirect temporal relations:* Information regarding the temporal validity of a relation is not always explicitly provided in EventKG. However, this information can often be derived based on the existence times of the participating entities or the happening times of the events. For example, the validity of a “mother” relation can be determined using the birth date of the child entity. We refer to such relations as *indirect temporal relations*. Each instance of `eventKG-s:Relation` that represents such an indirect temporal relation is transformed into a temporal relation $r_t = \langle r_{uri}, r_{time}, e_i, e_j \rangle \in R_t, r_{time} = e_{j_{time}}$.

4.3. EventKG Generation Pipeline

The EventKG generation pipeline is shown in Figure 4.

Input & Pre-processing: First, the dumps of the reference sources in the corresponding languages are collected. Both Wikidata and YAGO provide multilin-

⁷The VoID vocabulary <https://www.w3.org/TR/void/>.




Figure 4. The EventKG generation pipeline.

qual information in a single data dump. DBpedia and Wikipedia provide language-specific dumps, so that we collect the dumps for the languages of interest, i.e. EN, FR, DE, RU and PT. The Wikipedia Current Events Portal is currently available in English only. The mapping from the Wikidata identifiers to the Wikipedia and DBpedia identifiers required for the integration is collected as part of the Wikidata dump.

As part of the pre-processing, the following information is created for each language⁸:

- *Terms*: Terms is a set of terms and regular expressions used throughout the extraction process. This includes the month names, weekday names, a black list of namespaces and prefixes of the Wikipedia articles to be ignored (e.g. the prefix “Chronological_list_of_” in English) as well as regular expressions to detect titles of the Wikipedia articles representing events.
- *Date expressions*: To extract dates from unstructured reference sources, a set of regular expressions is created. These expressions are sorted in the decreasing order of specificity, where time intervals are considered to be more specific than the individual dates or months. For example, a specific regular expression to extract a span of two dates in English is: `@regexMonthDay1@@hyphensOr@@regexMonthDay2@`, where `@regexMonthDay1@` denotes a month name followed by a date and `@hyphensOr@` is any kind of hyphen. This regular expression can match textual patterns such as “February 17 — April 23”. A less specific expression is `@regexDay1@` that only checks for day numbers such as “17”. Moreover, regular patterns to identify Wikipedia event lists such as “2007 in Science” are created, together with the rules to extract the temporal scope (the year 2007 in this example).

⁸To obtain a complete list of the manually defined terms, expressions and mappings adopted in this work, please see the readme file in the open source software release provided at: <https://github.com/sgottsch/eventkg>

- *Mapping of predicates representing event relations*: We define a mapping table to identify predicates that represent equivalent event relations in EventKG and its reference sources such as `so:hasSubEvent` and Wikidata’s “part of” property. Examples of such mappings are shown in Table 5. In this work we define the predicate mappings manually. In future work schema mapping techniques can be adopted to determine such links automatically.

Identification and Extraction of Events: Event instances are identified in the reference sources and extracted as follows:

- **Wikidata** [10]: We identify events as subclasses of Wikidata’s “event” (representing temporary and scheduled events like festivals or competitions) and “occurrence” (representing happenings like wars or ceremonies). Some of the identified subclasses are blacklisted manually. For example, the class “song” is blacklisted because of the subclass hierarchy `song > musical form > art form > format > arrangement > act > process > occurrence`.
- **DBpedia** [11]: For each language edition, we identify DBpedia events as instances of `dbo:Event` or its subclasses.
- **YAGO** [12]: We do not use the YAGO ontology for event identification due to the noisy event subcategories we observed (e.g. `event > act > activity > protection > self-defense > martial_art`).
- **Wikipedia**: We use Wikipedia category names that match a manually defined language-dependent regular expression (e.g. English category names that end with “events”) as an indication that a knowledge graph entry linked to such an article is an event.
- **Wikipedia Event Lists**: For each language, we identify Wikipedia event lists by adopting a set of regular expressions defined manually during pre-processing. This way, Wikipedia pages with titles such as “2007 in Science” and “August 11” are retrieved. Within these pages, textual descriptions of events are collected using methods similar to [25].

Using the ordered list of regular temporal expressions and Wikipedia link markup, representations of events including their descriptions, linked entities and dates are extracted.

- **WCEP**: In the Wikipedia Current Events Portal, events are represented through rather brief textual descriptions and refer to daily happenings. We extract WCEP events using the WikiTimes tool [13].

Extraction of Event and Entity Relations: We extract the following types of relations: 1) *Relations with temporal validity* are identified based on the availability of temporal information. Temporal relations are extracted from YAGO and Wikidata. DBpedia does not provide such information. 2) *Relations with indirect temporal information*: we extract all relations involving events as well as relations of entities with known existence time. 3) *Other event and entity relations*: we use the manually defined mapping table shown in Table 5 to identify predicates that represent event relations in EventKG such as `so:hasSubEvent` (e.g. we map Wikidata’s “part of” property (P361) to `so:hasSubEvent` in cases where the property is used to connect events), `dbo:previousEvent` and `dbo:nextEvent` as well as `so:containedInPlace` to extract location hierarchies. 4) *Relation strength and event popularity information*: For each event-entity relation we extract language-specific interlinking information from Wikipedia. In particular we extract the number of links and the number of mentions for each relation involving events. Link and mentions are extracted from each Wikipedia language edition by parsing all of its pages.

Integration: The statements extracted from the reference sources are included in the named graphs, such that each named graph corresponds to a reference source. In addition, we create a named graph `eventKG-g:event_kg` containing information resulting from integration and fusion. Each `sem:Event` and `sem:Core` instance in the `eventKG-g:event_kg` graph integrates event-centric and entity-centric information from the reference sources related to equivalent real-world instances.

The integration of entities and events obtained from knowledge graphs and Wikipedia articles is conducted using existing `owl:sameAs` links, as provided by the Wikidata dataset. In particular, the entities and events covered by YAGO and different language versions of DBpedia and Wikipedia are also present in Wikidata. We use `owl:sameAs` links to the Wikidata identifiers to represent each resource that is linked as equivalent

in multiple reference sources as one resource in EventKG. That way, information regarding this resource in different reference sources, e.g. labels in different languages, is integrated. In the current version of EventKG, we do not apply any entity resolution techniques to identify missing `owl:sameAs` links in these reference sources. This can be addressed in future work to further increase the degree of integration.

The events in the Wikipedia event lists and WCEP do not possess unique identifiers. Such events are integrated using a rule-based approach to identify equivalent events. Two events e_1 and e_2 extracted from such sources are represented as one EventKG event if the times of these events are identical ($e_1.time = e_2.time$) and the set of entities they link to overlaps. A special case is given if an event e_1 without an identifier links to an exactly one event e_n with a known identifier and their times are equal. In that case, the text of e_1 is added as a description to e_n .

Fusion: In the fusion step, we aggregate temporal, spatial and type information of `eventKG-g:event_kg` events using a rule-based approach.

- *Time Fusion:* For each entity, event or relation with a known existence or a validity time stamp, time fusion is conducted using the following rules: (i) ignore the dates at the beginning or end of a time unit (e.g. January, 1st), if alternative dates are available; (ii) apply majority voting among the reference sources; (iii) take the time stamp from the more trusted source (in order: Wikidata, DBpedia, Wikipedia, WCEP, YAGO).
- *Location Fusion:* For each event in `eventKG-g:event_kg`, we take the union of its locations from the different reference sources and exploit the `so:containedInPlace` relations to reduce this set to the minimum (e.g. the set {Paris, France, Lyon} is reduced to {Paris, Lyon}, while France can still be induced as a location using `so:containedInPlace` transitively).
- *Type Fusion:* We provide `rdf:type` information according to the DBpedia ontology (`dbo`), using types and `owl:sameAs` links in the reference sources.

Output: Finally, extracted instances and relations are represented in RDF according to the EventKG data model (see Section 4.1). As described above, information extracted from each reference source and the results of the fusion step are provided in separate named graphs.

Table 5
Example property mapping between EventKG and its reference sources.

| EventKG | Wikidata | DBpedia | YAGO |
|-----------------------|---|---|---|
| sem:hasPlace | wd:P276 (location) wd:P30 (continent) ... | dbo:place | yago:isLocatedIn yago:happenedIn |
| sem:hasBeginTimeStamp | wd:P580 (start time) wd:P585 (point in time) wd:P1619 (date of official opening) ... | — | yago:startedOnDate yago:happenedOnDate |
| sem:hasEndTimeStamp | wd:P582 (end time) wd:P585 (point in time) ... | — | yago:endedOnDate yago:happenedOnDate |
| so:hasSubEvent | wd:P361 (part of) | dbo:isPartOf dbo:isPartOfMilitaryConflict ... | — |
| so:previousEvent | wd:P155 (follows) | dbo:previousEvent dbo:previousWork | — |
| so:nextEvent | wd:P156 (followed by) | dbo:followingEvent dbo:subsequentWork | — |
| so:containedInPlace | wd:P36 (capital) wd:P706 (located on terrain feature) ... | — | — |

4.4. Running Example: Barack Obama

In the context of our running example, we now provide an exemplary overview of the EventKG generation pipeline and illustrate how exemplar relations are expressed in the EventKG model and in the TKG. We refer to individual heterogeneous instances in the input data that are not yet expressed in the EventKG schema as *data items*. Table 6 provides exemplary data items involving Barack Obama obtained from Wikidata, YAGO and different language editions of Wikipedia and DBpedia.

Identification and Extraction of Events. The first data item is extracted from the English Wikipedia event list in the article “2018 in the United States”. The entities “first inauguration of Barack Obama”, “United States presidential election, 2012” and “Death of Osama bin Laden” from the data items #2, #3 and #5 are identified as events using the class hierarchies in the reference sources. In this example, Obama’s first inauguration is identified as an event, because it is an instance of “United States presidential inauguration”, which can be tracked back to `inauguration > key event > occurrence` in Wikidata. Thus, the text event from data item #1 and the event “first inauguration of Barack Obama” are stored as event instances with ad-

ditional values such as a textual description for the former and a title for the latter event.

Extraction of Event and Entity Relations. Given the set of events, we can now detect relations between them and other entities. For example, the statement that Barack Obama was involved in his own inauguration as US president is extracted from Wikidata. This statement represents an indirect temporal relation, as it alone does not provide the required temporal validity information, which needs to be extracted from a related fact about the event. Similarly, we can extract the information that Barack Obama was a candidate of the US elections in 2012 from the French DBpedia.

With the help of Wikipedia links, we connect Barack Obama to the death of Osama bin Laden (data item #5). Given the relation `?rel` that links to Barack Obama as the subject and to the event “Death of Osama bin Laden” as the object, the link information is modelled as follows, using a named graph (where `eventKG-r:entity_11973762` represents Barack Obama and `eventKG-r:event_527087` represents the event “Death of Osama bin Laden”):

Table 6
Example data items about Barack Obama extracted from different reference sources.

| # | Reference Source | Data Item | Related Data Items |
|---|-------------------------|--|---|
| 1 | Wikipedia ^{EN} | 8 May 2018: President Trump announces his intention to withdraw the United States from the Iranian nuclear agreement. In a statement, former U.S. President Barack Obama calls the move "a serious mistake". | — |
| 2 | Wikidata | Barack Obama, significant event, first inauguration of Barack Obama | Wikidata: first inauguration of Barack Obama, point in time, 20 January 2009 YAGO: first inauguration of Barack Obama, was created on, 17 July 1981 Wikidata: first inauguration of Barack Obama, instance of, United States presidential inauguration Wikidata: United States presidential inauguration, subclass of*, occurrence |
| 3 | Wikidata | Barack Obama, spouse, Michelle Obama start time: 3 October 1992 | — |
| 4 | DBpedia ^{FR} | Barack Obama, prop-fr:candidat, Élection présidentielle américaine de 2012 | DBpedia ^{FR} : Élection présidentielle américaine de 2012 owl:sameAs United States presidential election, 2012 Wikidata: United States presidential election, 2012, point in time, 6 November 2012 |
| 5 | Wikipedia ^{PT} | [The Portuguese Wikipedia page of Barack Obama links to the page "Death of Osama bin Laden" once.] | Wikidata: Death of Osama bin Laden, point in time, 2 May 2011 |

```
?rel rdfs:type
    eventKG-s:Relation .
?rel rdf:subject
    eventKG-r:entity_11973762 .
?rel rdf:object
    eventKG-r:event_527087 .

eventKG-g:wikipedia_pt {
    ?rel eventKG-s:links 1 .
} .
```

For the relation `?rel`, link information can be added using specific named graphs. For example, such information can model the co-mentions of Barack Obama and the death of Osama bin Laden in the Portuguese Wikipedia.

Another type of information is coming from the temporal relations between two temporal entities: Here, the *spouse* relation between Barack and Michelle Obama is directly assigned a temporal validity time by Wikidata.

Integration. The entities “Élection présidentielle américaine de 2012” and “United States presidential election, 2012,” are modeled as the same event resource in EventKG, using DBpedia’s owl:sameAs link.

Fusion. There are two different dates provided for the first inauguration of Barack Obama (data item #2). While both dates are stored in EventKG together with their provenance information (i.e. as named graphs for Wikidata and YAGO), a single happening time for that event is created with our rule-based fusion approach (see Section 4.3). As the majority voting is not sufficient here, we take the date from the higher trusted source. In this case, Wikidata’s date (20 January 2009) is selected for EventKG’s named graph.

With that time information, the indirect temporal relation about Obama’s participation in his own inauguration can be transformed into the following temporal relation in the *TKG* generated from the named graph `eventKG-g:event_kg`:

```
Barack Obama,
significant event:
first inauguration of Barack Obama
[2009-01-20, 2009-01-20]
```

5. EventKG Characteristics & Evaluation

To demonstrate the quality of the data extraction, integration and fusion steps, we first show characteristics of EventKG and provide several comparisons to

its reference sources in Section 5.1. Then, we provide evaluation results based on user annotations in Section 5.2.

5.1. Characteristics

In EventKG V1.1, we extracted event representations and relations in five languages – English (EN), German (DE), French (FR), Russian (RU) and Portuguese (PT) – from the latest available versions of each reference source as of 12/2017. EventKG uses open standards and is publicly available under a persistent URI⁹ under the CC BY 4.0 license¹⁰. Our extraction pipeline is available as open source software on GitHub¹¹ under the MIT License¹². A description of EventKG and example SPARQL queries are online¹³. Two example SPARQL queries are also presented in Appendix A.

Table 7 summarises selected statistics from the EventKG V1.1, released in 03/2018. Overall, this version provides information for over 690 thousand events and over 2.3 million temporal relations. Nearly half of the events (46.75%) originate from the existing knowledge graphs; the other half (53.25%) is extracted from semi-structured sources. The data quality of the individual named graphs directly corresponds to the quality of the reference sources. In eventKG-g:event_kg, the majority of the events (76.21%) possess a known start or end time. Locations are provided for 12.21% of the events. The coverage of locations can be further increased in future work, e.g. using NLP techniques to extract locations from event descriptions. Along with over 2.3 million temporal relations, EventKG V1.1 includes relations between events and entities for which the time is not available. This results in overall over 88 million relations. Approximately half of these relations possess interlinking information.

5.1.1. Comparison of EventKG to its Reference Sources

We compare EventKG to its reference sources in terms of the number of identified events and completeness of their representations. The results of the event identification and extraction step in Section 4.3 are shown in Table 8. EventKG with 690,247 events contains a significantly higher number of events than any

of its reference sources. This is especially due to the integration of knowledge graphs and semi-structured sources.

Table 9 presents a comparison of the event representations in EventKG and its reference knowledge graphs (Wikidata, YAGO, DBpedia). As we can observe, through the integration of event-centric information, EventKG: 1) enables better event identification (e.g. we can map 322,669 events from EventKG to Wikidata, whereas only 266,198 were identified as events in Wikidata initially - see Table 8) and 2) provides more complete event representations (i.e. EventKG provides a higher percentage of events with specified temporal and spatial information compared to Wikidata, that is the most complete reference source). The most frequent event types are source-dependent (see Table 10).

5.1.2. Relation & Fusion Statistics

Over 2.3 million temporal relations are an essential part of EventKG. The majority of the frequent predicates in EventKG such as “member of sports team” (882,398 relations), “heritage designation” (221,472), “award received” (128,125) and “position held” (105,333) originate from Wikidata. The biggest fraction of YAGO’s temporal relations have the predicate “plays for” (492,263), referring to football players. Other YAGO predicates such as “has won prize” are less frequent. Overall, about 93.62% of the temporal relations have a start time from 1900 to 2020. 81.75% of events extracted from knowledge graphs are covered by multiple sources. At the fusion step, we observed that 93.79% of the events that have a known start time agree on the start times across the different sources.

5.1.3. Textual Descriptions

EventKG V1.1 contains information in five languages. Overall, 87.65% of the events extracted from knowledge graphs provide an English label whereas only a small fraction (4.49%) provide labels in all languages. Among the 367,578 events extracted from the semi-structured sources, just 115 provide a description in all five languages, e.g. the first launch of a Space Shuttle in 1981. This indicates potential for further enrichment of multilingual event descriptions in future work.

5.2. Evaluation of EventKG

The aim of the evaluation is to assess the effectiveness of the event identification, time fusion and location fusion steps of the pipeline.

⁹<https://doi.org/10.5281/zenodo.1112283>

¹⁰<https://creativecommons.org/licenses/by/4.0/>

¹¹<https://github.com/sgottsch/eventkg>

¹²<https://opensource.org/licenses/MIT>

¹³<http://eventkg.l3s.uni-hannover.de/>

Table 7

Number of events and relations in eventKG-g:event_kg.

| | #Events | Known time | Known location |
|-------------------------------------|------------|------------|----------------|
| Events from KGs | 322,669 | 163,977 | 84,304 |
| Events from semi-structured sources | 367,578 | 362,064 | not extracted |
| Relations | 88,473,111 | 2,331,370 | not extracted |

Table 8

Number of events identified in extracted from the reference sources.

| Wikidata | DBpedia | | | | | Wikipedia event lists | | | | | WCEP |
|----------|---------|--------|-------|-------|--------|-----------------------|---------|--------|--------|--------|--------|
| | EN | FR | DE | RU | PT | EN | FR | DE | RU | PT | |
| 266,198 | 60,307 | 43,495 | 9,383 | 5,730 | 14,641 | 131,774 | 110,879 | 21,191 | 44,025 | 18,792 | 61,382 |

Table 9

Comparison of the event representation completeness in the source-specific named graphs (after integration).

| | EventKG | Wikidata | YAGO | DBpedia | | | | |
|--------------|---------|----------|---------|---------|--------|--------|--------|--------|
| | | | | EN | FR | DE | RU | PT |
| #Events with | 322,669 | 322,669 | 222,325 | 214,556 | 78,527 | 62,971 | 47,304 | 35,682 |
| Location (L) | 26.13% | 11.70% | 26.61% | 6.21% | 8.32% | 4.03% | 10.60% | 6.15% |
| Time (T) | 50.82% | 33.00% | 39.02% | 7.00% | 17.21% | 2.00% | 1.35% | 0.08% |
| L&T | 21.97% | 8.83% | 19.02% | 4.29% | 0.00% | 4.84% | 1.18% | 0.08% |

Table 10

The most frequent event types extracted from the references sources and the percentage of the events in that source with the respective type.

| | Wikidata | DBpedia | | | | |
|-----------|----------|----------------------|-----------------|----------------------|----------------------|----------------------|
| | | EN | FR | DE | RU | PT |
| dbo:type | season | Military Conflict | Sports Event | Tennis Tournament | Military Conflict | Soccer Tournament |
| Events, % | 11.37% | 6.31% | 21.86% | 33.00% | 11.87% | 16.17% |

5.2.1. Event Identification

We manually evaluated a random sample of the events identified in the event identification step of EventKG (Section 4.3). For each reference source, we randomly sampled 100 events and manually annotated whether they represent real-world events or not. The results are shown in Table 11.

For DBpedia and Wikidata, where we rely on the event types and type hierarchies, we achieve a precision of 98% on average. On a random sample of 100 events extracted from the category names in the English and the Russian Wikipedia, we achieve 94% and 88% precision, correspondingly. One example for an entity wrongly identified as an event is the canceled project “San Francisco Municipal Wireless”, which was part of the “Cancelled projects and events” category in Wikipedia.

5.2.2. Time Fusion

To evaluate the quality of the proposed rule-based time fusion approach, we randomly sampled 100 events from EventKG, where each event has at least two reference sources that differ in the event happening time (i.e. start and/or end time). Three users have annotated this sample by providing a start and end time for at least 20 events each. Additionally, we asked the users to denote which source they used to research the actual event dates. For our evaluation, we then checked how many of the user-given start and end dates are available in the reference sources and the joint EventKG named graph, and we computed how many of these dates are correct with respect to the user annotations.

Table 12 provides the result overview: As the time fusion does always adopt accessible time information from any reference source, all events in our random sample possess time information. Wikidata and YAGO

Table 11

User-evaluated precision for the identification of events with selected reference sources.

| | Wikidata | DBpedia _{DE} | DBpedia _{RU} | DBpedia _{PT} | Wikipedia _{EN} | Wikipedia _{RU} |
|------------------|----------|-----------------------|-----------------------|-----------------------|-------------------------|-------------------------|
| Precision | 96% | 100% | 100% | 98% | 94% | 88% |

Table 12

Evaluation of EventKG’s time information. For EventKG and the reference sources, the percentage of correct, wrong and missing event dates with respect to the user annotations in our sample is shown. These are based on the random sample of events where the reference sources show disagreement between time information provided.

| Source | Start Dates | | | End Dates | | | Start and End Dates | | | Precision |
|-----------------------|-------------|-------|---------|-----------|-------|---------|---------------------|-------|---------|-------------|
| | Correct | Wrong | Missing | Correct | Wrong | Missing | Correct | Wrong | Missing | |
| EventKG | 71 | 25 | 0 | 73 | 23 | 0 | 144 | 48 | 0 | 0.75 |
| Wikidata | 40 | 33 | 23 | 33 | 29 | 34 | 73 | 62 | 57 | 0.54 |
| YAGO | 21 | 60 | 15 | 20 | 57 | 19 | 41 | 117 | 34 | 0.26 |
| DBpedia _{EN} | 12 | 5 | 79 | 13 | 4 | 79 | 25 | 9 | 158 | 0.74 |
| DBpedia _{DE} | 0 | 2 | 94 | 2 | 0 | 94 | 2 | 2 | 188 | 0.5 |
| DBpedia _{FR} | 6 | 17 | 73 | 15 | 8 | 73 | 21 | 25 | 146 | 0.46 |
| DBpedia _{RU} | 0 | 2 | 94 | 0 | 2 | 94 | 0 | 4 | 188 | 0 |

Table 13

Time Fusion Evaluation: The most frequent sources used by the users to lookup event start and end dates.

| Source | #Uses | Percentage |
|-----------------------------|-------|------------|
| en.wikipedia.org | 117 | 58.5% |
| www.google.com | 37 | 18.5% |
| de.wikipedia.org | 14 | 7.0% |
| no source used | 7 | 3.5% |
| fr.wikipedia.org | 6 | 3.0% |
| www.singapore-elections.com | 2 | 1.0% |
| www.un.org | 2 | 1.0% |
| ... | | |

provide the next highest coverage of time information. In terms of precision, EventKG outperforms these two reference sources by 21% (Wikidata) and 49% (YAGO). This result confirms the quality of the proposed rule-based time fusion approach. The results of a McNemar’s test [26] has shown a two-tailed p-value of less than 0.0001, which confirms the statistical significance of this result.

Table 13 provides an overview of the sources most often used for finding the event dates by the users participating in the evaluation. In 69% of the cases, the users adopted Wikipedia articles in different languages as their source. When the users did not use Wikipedia, either the information presented on the search engine’s result page (18.5% of the cases) or domain-specific web sites such as www.singapore-elections.com or www.un.org were used.

5.2.3. Location Fusion

To evaluate the correctness of the extracted locations, we selected a random sample of 100 events with at least one location. In case of locations, multiple correct values are possible, for example South America, the United States of Colombia and the Colombia-Ecuador border are valid locations for the Ecuadorian-Colombian War. We presented all locations from each reference source to the users and for each location asked the users to verify whether that location is correct or not. Four users have annotated that sample.

Table 14 provides the result for our evaluation of the location fusion. We distinguish between the locations directly provided by EventKG and those which could be inferred using sub-location information via `so:containedInPlace`. We refer to this extended knowledge graph as EventKG* throughout this evaluation. EventKG and EventKG* have by far the highest coverage of locations (EventKG* finds 78.13% more event locations than YAGO and 159.10% more than in Wikidata), while keeping the number of wrong locations low (approx. 7%), although it also inherits wrong locations as provided by the reference sources due to the adopted location fusion mechanism. The results of a McNemar’s test [26] has shown a two-tailed p-value of 0.0005, which confirms statistical the significance of this result.

Table 15 lists the sources used by the users in this task. Similarly to the evaluation of the time fusion, Wikipedia and Google were the most frequently used sources, followed by domain-dependent ones such as

Table 14

Evaluation of EventKG’s location information. For each event in the sample, users judged for each location in EventKG and the reference sources whether it is correct.

| Source | Correct | Wrong | Precision |
|-----------------------|---------|-------|-----------|
| EventKG* | 116 | 7 | 94.31% |
| EventKG | 87 | 4 | 95.60% |
| YAGO | 64 | 2 | 96.97% |
| Wikidata | 44 | 2 | 95.65% |
| DBpedia _{EN} | 15 | 1 | 93.75% |
| DBpedia _{FR} | 7 | 0 | 100.0% |
| DBpedia _{DE} | 1 | 0 | 100.0% |
| DBpedia _{RU} | 4 | 1 | 80.0% |
| DBpedia _{PT} | 3 | 1 | 75.0% |

Table 15

Location Fusion Evaluation: The most frequent sources used by the users to lookup event locations.

| Source | #Uses | Percentage |
|------------------|-------|------------|
| en.wikipedia.org | 58 | 43.94% |
| no source used | 35 | 26.51% |
| de.wikipedia.org | 7 | 5.3% |
| www.google.com | 5 | 3.79% |
| everipedia.org | 3 | 2.0 % |
| fr.wikipedia.org | 3 | 2.0 % |
| www.kicker.de | 2 | 1.51% |
| ... | | |

kicker.de for locating football matches. However, in 26.51% of the cases in this task, the users did not use a source at all, mainly because many event locations are self-explanatory or contained in the event names. For example, no source was needed to verify the locations Monaco and Circuit de Monaco for the 1956 Monaco Grand Prix.

6. Generation of Biographical Timelines

In this section, we show how EventKG can be applied as a temporal knowledge graph for the task of biographical timelines generation.

First, we present our approach based on distant supervision in Section 6.1. The features used in the relevance model are introduced in Section 6.2. Subsequently, we describe the benchmarks involved in our process to generate biographical timelines in Section 6.3 and discuss how the model is used to generate them in Section 6.4. Finally, we illustrate these steps on our running example of Barack Obama’s timeline in Section 6.5.

6.1. Approach

Given a timeline entity e for which we need to generate a biographical timeline, the number of *candidate timeline entries* (i.e. temporal relations involving e) is potentially very high, especially for popular entities and a large-scale temporal knowledge graph. In fact, for our set of popular persons described later in Section 7.1, EventKG contains 272.75 temporal relations per person entity on average. In order to determine the relevance of a temporal relation to the timeline entity we propose a classification approach using distant supervision. The key idea of our approach is to learn a relevance model for temporal relations using occurrences of these relations extracted from *biographical sources*. Examples of such biographical sources include collections of biographical or encyclopedic articles. We adopt a distant supervision approach, where we assume that a particular temporal relation r is relevant for the entity’s biography if this relation occurs in a known biographical source. An overview of the training phase and the timeline generation is depicted in Figure 5, which illustrates the role of the TKG, the biographical and reference sources and the benchmark. Initially, we use the temporal knowledge graph and a biographical source to create a benchmark that provides relevance judgements for candidate timeline entries. We train the prediction model with features extracted for each candidate timeline entry. This includes entity type and interlinking information included in the named graphs corresponding to the reference sources of EventKG. To generate a timeline for a timeline entity e , we collect its candidate timeline entries R_e from TKG and identify the relevant entries using the trained model.

6.2. Relevance Model

In our approach we train a classification model that identifies the relevance of a candidate timeline entry towards a biography of the timeline entity e . The candidate timeline entry is a temporal relation involving e and obtained from a knowledge graph. To train such classification models, we adopt a range of features in several categories reflecting the characteristics of the timeline entity, the entity connected to it via a temporal relation, the temporal relation and time information. In total, we consider 4 language-independent numerical features, 6 language-dependent features, as well as a number of binary features representing frequent entity types and properties in EventKG.

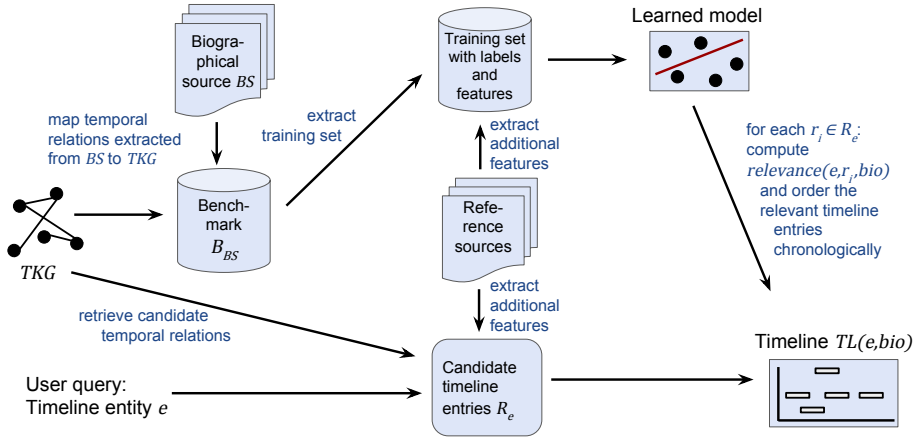


Figure 5. Creating a timeline for a timeline entity e , after training a model from a biographical source to predict the relevance of temporal relations in the TKG for biographical timelines.

We illustrate the features described in the following at the example of the candidate timeline entry representing Barack Obama’s participation in his second inauguration (see Figure 3) in Table 16.

Timeline Entity Features

The timeline entity features (TEF) reflect specific characteristics of the timeline entity e . These features address the intuition that the relevance of the particular temporal relation r for a given timeline entity e depends on the specific characteristics of e . For example, winning an award may be more important for athletes or actors than for politicians. Based on this intuition, we introduce the timeline entity features:

TEF-C Timeline entity characteristics: A set of binary features denoting if the entity is an instance of the specific type (e.g. a politician or an actor).

Connected Entity Features

The connected entity features (CEF) take into account characteristics of the connected entity e' . In particular, we consider indications of the importance and popularity of e' in the context of the reference collections by using mention counts, similar to Thalhammer et al. [27]. In particular, we consider different representations of the mention counts of e' .

CEF-M Connected entity mentions: The set of features, each reflecting the absolute number of mentions of the connected entity e' in a reference collection.

CEF-MR Connected entity mentions rank: For each reference collection, we rank the entities connected to the timeline entity e by the number of their men-

tions. This feature represents the rank of the specific connected entity, where the rank of 1 is assigned to the entity with the highest number of mentions.

CEF-MRR Connected entity mentions relative rank: We normalise the CEF-MR rank by the maximal rank.

CEF-E Connected entity represents a real-world event: A binary feature denoting whether the connected entity is an event (i.e. $e' \in \mathcal{V}$).

Features of Temporal Relations

The features of temporal relations (TRF) reflect semantics of the temporal relation between the timeline entity and the connected entity. Furthermore, we consider features related to the importance and popularity of entity relations.

TRF-PI Property identifier: Temporal relations possess property identifiers r_{uri} that express semantics of the relation (e.g. *dbo:spouse*). Each property identifier is modelled as a binary feature.

TRF-M Relation mentions: The number of co-mentions of both entities involved in the temporal relation in a reference collection (independent of relation semantics).

TRF-MR Relation mentions rank: We rank the connected entities according to the number of their co-mentions with the timeline entity in a reference collection. This feature represents the rank of the specific connected entity involved in the relation.

TRF-MRR Relation mentions relative rank: We normalise the TRF-MR rank by its maximal rank.

Temporal Features

The temporal features (TF) reflect the relevance of the temporal relations based on the time information. This includes the temporal differences in the existence time of the entities or happening times of the events involved in the relation. For example, Barack Obama gave a speech related to World War II - a historical event finished before Obama's birth date in 1961. Here, the temporal difference in the existence times of both entities can be an indication of the low relevance of this speech for Obama's biography. Therefore, we attempt to learn to discard the temporal relations involving events that happened too early for the entity timeline. This had been also observed by Althoff et al. [2] who implemented a rule to discard such relations. In addition to that, our temporal features could help to learn whether some events may be more relevant at specific stages of the entity's life or existence. Furthermore, temporal features include the provenance of the temporal information by denoting whether a relation was induced from an indirect temporal relation or not.

To capture this intuition, we introduce the following temporal features:

- TF-TDS Temporal distance (start): The temporal distance between the beginning of the existence time of the timeline entity and the start of the relation validity time $e_{start} - r_{start}$.
- TF-TDE Temporal distance (end): The same feature as TF-TDS, but using the entity existence end time $e_{end} - r_{start}$.
- TF-TP Time provenance: This categorical feature specifies the provenance of the relation validity time. If the relation has initially been a temporal relation, the feature value is set to 3. If the temporal validity was induced from an event happening time ($e_j \in \mathcal{V}$), then the feature value is set to 2; 1 otherwise ($e_j \in \mathcal{E}'$).

6.3. Benchmarks for Distant Supervision

To facilitate supervised model training, we require a benchmark that provides relevance judgements for temporal relations. These judgements can be obtained from the specific biographical source.

Definition 5. A benchmark B is a mapping of the form: $relevance(e_i, r_j, bio) \mapsto J, J \in \{0, 1\}$, where e_i is a temporal entity, r_j is a temporal relation involving e_i and J is a relevance judgement.

Given the large number of entities and temporal relations in the existing knowledge graphs, manual relevance judgements appear unfeasible. Therefore, we adopt an automatic approach to benchmark generation. We extract entities and temporal relations contained in the biographical sources and map them to the temporal relations in TKG using an automatic procedure involving source-specific heuristics (described later in Section 7.1). Temporal relations extracted from the biographical sources are considered relevant.

Although the resulting benchmarks can potentially contain noisy relevance judgements due to the automatic extraction and mapping methods applied, our experimental results demonstrate that these benchmarks, used as a training set in a distant supervision method, facilitate generation of high quality timelines.

The benchmarks created in this work are publicly available online¹⁴.

6.4. Model Training and Timeline Generation

We address the relevance estimation for a timeline relation r with respect to the timeline entity e as a classification problem. For each biographical source BS , we build a classification model using the features presented in Section 6.2 and a binary classifier.

Note that a classification model is chosen over a ranking-based approach because of two reasons: First, the timeline entries are ordered chronologically and not by their importance. Therefore, for the purpose of timeline generation we can assume that each timeline entry is equally relevant. Second, if a ranked list of timeline entries would be provided, a cut-off threshold value would still be required to decide which of the entries are to be shown.

To facilitate efficient training we limit the number of instances of the TEF-C and TRF-PI features considered. In particular, the 50% most frequent types in the training set are added as a TEF-C feature. Furthermore, only properties that occur in at least 25% of the relations in the training set are added as a TRF-PI feature.

Our benchmark is equally divided into a training and a test set of person entities, so that the relevance judgements are obtained from the training set. We adopt a binary notion of relevance. The datasets used as biographical sources to build the classification models are presented in Section 7.1.

¹⁴<http://eventkg.l3s.uni-hannover.de/timelines.html>

Table 16

Selected feature values for the candidate timeline entry “Barack Obama, significant event, Second inauguration of Barack Obama” for the timeline entity “Barack Obama”.

| Feature | Feature Instance | Value | Note |
|---------|-----------------------|-------|--|
| TEF-C | Politician | 1 | Barack Obama is an instance of dbo:Politician. |
| | President | 1 | Barack Obama is an instance of dbo:President. |
| | Scientist | 0 | Barack Obama is not an instance of dbo:Scientist. |
| CEF-M | CEF-M _{EN} | 84 | The inauguration is linked 84 times in the English Wikipedia. |
| CEF-MR | CEF-MR _{EN} | 361 | Among all entities connected to Obama in the English Wikipedia, the inauguration is linked the 361st most times. |
| CEF-MRR | CEF-MR _{EN} | 0.817 | Among all entities connected to Obama in the English Wikipedia, there are 442 different CEF-MR _{EN} scores, such that inauguration’s relative rank is $\frac{361}{442} \approx 0.817$. |
| CEF-E | CEF-E | 1 | The inauguration is an instance of sem:Event. |
| TRF-PI | wd:significantEvent | 1 | Obama is connected to the inauguration through Wikidata’s “significant event” property. |
| | wd:spouse | 0 | Barack Obama is not connected to the inauguration through Wikidata’s “spouse” property. |
| TRF-M | TRF-M _{PT} | 4 | In the Portuguese Wikipedia, there are 4 sentences mentioning both Barack Obama and the inauguration. |
| TRF-MR | TRF-MR _{PT} | 18 | Among all co-mentions of Barack Obama and an event, the co-mention with the inauguration is the 18th most frequent one the Portuguese Wikipedia. |
| TRF-M | TRF-M _{ALL} | 36 | In all the five involved Wikipedia language editions together, there are 36 sentences mentioning both Obama and the inauguration. |
| TRF-MR | TRF-MR _{ALL} | 39 | Among all co-mentions of Barack Obama and an event, the co-mention with the inauguration is the 39th most frequent one in all the five involved Wikipedias together. |
| TF-TDS | TF-TDS | 18798 | The inauguration started 18798 days (51 years) after Barack Obama’s birth. |
| TF-TDE | TF-TDE | 18798 | The inauguration ended 18798 days (51 years) after Barack Obama’s birth. |
| TF-TP | TF-TP | 2 | The validity time assigned to this temporal relation is induced from the happening time of an event instance. |

We use the resulting classification model to build a timeline $TL(e, bio)$. Each candidate timeline entry (i.e. a temporal relation involving the timeline entity e in TKG) is classified using the classification models learned from a biographical source. The classification function $relevance(e, r, bio)$ uses this model to classify the temporal relations of the timeline entity e as either 0 (non-relevant) or 1 (relevant). As illustrated in Figure 5, the timeline is generated by ordering the timeline entries classified as relevant by their start time.

6.5. Running Example: Barack Obama

As discussed in Section 4.4, EventKG contains many relations involving Barack Obama. In order to create a timeline of his life, we collect all relations with Obama as a subject or an object, together with their temporal validity. One example is the temporal relation about Obama’s first inauguration shown at the end of Section 4.4.

Due to the more than 2,500 candidate timeline entries for Obama, we now need to apply the previously trained model to determine the timeline entries rele-

vant for a biography. To this end, we train the classifier that predicts whether a candidate timeline entry is relevant given a biographical source, i.e. whether it is probable to be part of entity biography in such source. All candidate timeline entries that are classified as relevant by this model are inserted into the timeline in chronological order.

Figure 1 provides a visual representation of Obama’s timeline obtained using a model trained on a Wikipedia abstracts dataset (BS-ENC) described later in Section 7.

7. Setup and Evaluation of the Biographical Timeline Generation

In this section we first describe the biographical sources and the set of timeline entities used to create our biographical timeline benchmark used to train the classification models (Section 7.1) and to run our experiments described in Section 7.2. Then, we evaluate our approach against a baseline (Sections 7.3 and 7.4).

Table 17

Example data extracted from the biographical sources for Barack Obama.

| | BS-BIO | BS-ENC |
|---------------------|--|---|
| Source | biography.com, thefamouspeople.com | Wikipedia _{EN} abstracts |
| Example Data | 1961-8-4, {Honolulu} 1979, {Punahou School, Basketball} 2000, {Democratic Party, Bobby Rush} 2010-8, {War in Afghanistan, Iraq} | 1961, {Honolulu} 2013, {US presidential election 2012, Mitt Romney, Second inauguration of Barack Obama} 2009, {Nobel Peace Prize} |

7.1. Benchmark: Entities and Biographical Sources

We collect a dataset \mathcal{P} that contains 2,760 timeline entities of the type *Person*, including its subtypes like politicians, actors, musicians and athletes. This set of 2,760 entities contains all persons that are included in EventKG and described in each biographical source described below. Consequently, the training and the test set consist of 1,380 person entities each, after random division.

To train the relevance models for the biographical timeline generation, we consider the following biographical sources:

- BS-BIO: Biographical articles;
- BS-ENC: Encyclopedic articles.

Biographical articles (BS-BIO):

Biographies of important entities (e.g. famous people) are available in form of textual descriptions from dedicated Web sources. We collect data from two publicly accessible biographical web sources (Thefamouspeople.com¹⁵ and Biography.com¹⁶). After collecting the biographical texts from both websites, they are pre-processed as follows: 1) The texts are split into sentences using the Stanford Tokenizer [28]. 2) Time expressions are collected from each sentence using HeidelTime [29]. 3) Entity mentions are identified using DBpedia Spotlight [30]. Table 17 illustrates example annotations in the *BS-BIO* and *BS-ENC* datasets extracted for the entity Barack Obama, including his birth, education and political activities. In order to map the extracted information to the temporal relations in the *TKG*, we use the following rule-based approach: An annotated sentence in the biographical article is mapped to the temporal relation in *TKG* if they both happened on exactly the same date, or if they share both entities and time. A special case is given if one of the linked entities is an event in \mathcal{V} . In that case, temporal overlap is not required, as events are typi-

¹⁵www.thefamouspeople.com

¹⁶www.biography.com

Table 18

Statistics of the dataset \mathcal{P} involving 2,760 entities of type person.

| | thefamous- people.com | biogra- phy.com | Wikipedia Abstracts |
|-------------------------|--------------------------|--------------------|------------------------|
| Time expressions | 50,919 | 41,318 | 18,099 |
| Entity links | 107,126 | 92,149 | 32,516 |

cally inherently connected to a validity time span. The mapped temporal relations from the *TKG* are added to the B_{BS-BIO} benchmark.

Encyclopedic articles (BS-ENC):

Wikipedia is a rich source of encyclopedic information. Wikipedia articles usually provide an abstract – a brief overview of the specific entity (e.g. person’s life) that typically contains important biographical sentences [31, 32]. From these abstracts, we extract all the event mentions, i.e. links to the event articles, as these represent significant events in the entity’s life. For example, Table 17 shows selected events for the entity Barack Obama based on BS-ENC. In contrast to the annotations in B_{BS-BIO} , these events are more focused on the political happenings with major public impact. The benchmark B_{BS-ENC} includes all relations of the specific entity to the events linked from the abstract of the Wikipedia article representing this entity.

Statistics of the entity-related information for the entities contained in the dataset \mathcal{P} in the biographical sources, including in particular the number of relevant entity links and time expressions are provided in Table 18. Figure 6 illustrates the distribution of the number of relevant relations per person in the BS-BIO benchmark. Except for very few popular entities such as David Bowie and Barack Obama, the number of relevant relations is typically below 100, with an average of 13.64.

We generate a benchmark B_{BS} for each biographical source *BS* considered in this work. The statistics regarding these benchmarks are presented in Table 19.

Table 20 provides the percentage of person types in the benchmarks. Actors and musical artists are the




Figure 6. The number of person entities with the given number of relevant relations in the BS-BIO benchmark. The top-3 entities with the highest number of relevant relations are marked.

Table 19

Benchmark statistics: the number of entities and relevant temporal relations (temp. rel.).

| | #Per- sons | #Relevant Tem- poral Relations | Avg. # Temp. Rel. per Entity |
|--------------|---------------|-----------------------------------|---------------------------------|
| B_{BS-BIO} | 2,760 | 37,638 | 13.64 |
| B_{BS-ENC} | 2,760 | 33,106 | 12.00 |

Table 20

Percentage of top-5 entity types in the training and test set.

| | Training | Test |
|-----------------------|----------|--------|
| Actor | 27.73% | 28.57% |
| Musical Artist | 13.32% | 16.17% |
| Athlete | 10.50% | 6.16% |
| Politician | 10.35% | 10.44% |
| Writer | 6.95% | 11.31% |

most frequent person types in both the training and test set.

7.2. Classifier Setup and Timeline Statistics

As our binary classifier we adopted a Support Vector Machine (SVM) due to its good generalisation ability, in particular when applied to smaller datasets. We trained this classifier on the training dataset containing 1,380 person entities, with input data normalisation, an increased weight of 3.0 for predicting relevant instances and a linear kernel, using Weka's LibSVM implementation [33]. From the training data, a balanced set of relevant and irrelevant instances is given to the SVM.

As described in Section 6.4, the timelines are generated by ordering the timeline entries classified as relevant chronologically by their start time. On average,

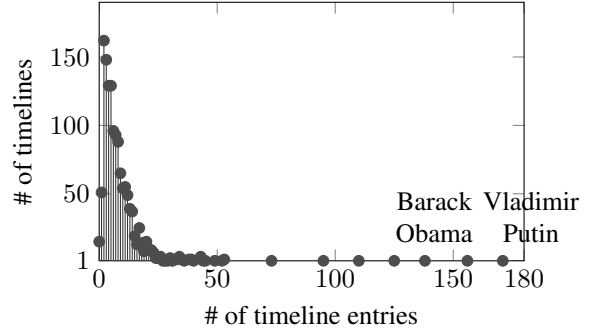


Figure 7. The number of timelines with the specific number of entries generated for the BS – BIO test set.

each biographical timeline of the person entities in the test set contains 8.54 entries after training the classifier on B_{BS-BIO} (B_{BS-ENC} : 7.81). Figure 7 illustrates the number of timelines generated for the BS – BIO with the specific number of entries.

7.3. The TM Baseline Algorithm

We compare our proposed approach with the state-of-the-art Time Machine (TM) approach for timeline generation proposed by Althoff et al. [2]. The TM approach creates events from the entity-entity relations in a knowledge graph, where one entity possesses a property with a time value. Resulting events are filtered using frequency and existence time heuristics; then a greedy algorithm selects the events that maximise a relevance score. To facilitate a fair comparison, we perform the following adjustments to implement the TM baseline:

- The TM approach in [2] was initially proposed for entity-centric knowledge graphs such as Freebase. Therefore, events in the TM terminology mean link structures in an entity-centric knowledge graph that vary with respect to their complexity. In EventKG, the events are connected to the entities directly via temporal relations. To facilitate the comparison, we adopt the TM baseline such that so-called "simple events" in the TM-terminology are generated. Such "simple events" in TM directly correspond to the temporal relations in EventKG.
- In the original TM approach, the maximal number of temporal relations on the timeline is restricted due to the visualisation constraints; i.e. these relations are ranked by their relevance and retrieved until the visualisation constraint is met. Our goal is to provide all relevant relations, such that we do

not enforce any visualisation-based constraints on the number of relations. To facilitate comparison, we retrieve an equal number of relations from the baseline and our approach.

- TM was initially evaluated on the Freebase dataset, and the relevance scores were computed using a search engine query log and a textual corpus. We apply all methods on the EventKG data; we use the same reference sources (i.e. Wikipedia articles) to estimate the parameters related to the global importance of entities, their occurrences and temporal relations for all baselines and approaches evaluated in this article.

7.4. Evaluation of the Timeline Generation

The goals of the evaluation of the timeline generation are to assess the effectiveness of the proposed method for timeline generation and the role of the reference and biographical sources.

In particular, we assess:

- G1 Quality of the generated timelines in comparison to the baseline (in a user evaluation).
- G2 Impact of the individual features on the timeline generation (using correlation measures).
- G3 Relevance of the timeline entries with respect to the biographical source (by measuring performance of the classification model).
- G4 Coverage of the timeline entries with respect to the reference sources (by measuring the mean coverage of the temporal relations in the reference sources).

7.4.1. Timeline Quality Evaluation

In order to evaluate the timeline quality we performed a user evaluation. We generated timelines for 60 popular entities of the types actors, athletes, musical artists, politicians and writers for both biographical sources BS-BIO and BS-ENC. These entities were selected from the persons in the test set described in Section 7.1 based on their popularity (measured as the link count of the corresponding Wikipedia article).

In each task, the user was presented with: (i) a task description, (ii) a timeline entity including its label and a Wikipedia link, and (iii) a pair of timelines. One timeline in the pair was generated by the specific configuration of our approach, the other timeline was generated by the TM baseline described in Section 7.3. Both timelines were visualised as illustrated in Figure 1. Each timeline contained all entries generated by the corresponding generation method. The user could

scroll and zoom within each individual timeline. In the user interface, both timelines were presented simultaneously, one above the other, in a random order. We asked the users to vote for their preferred timeline in the pair. We provided four options: two options to vote for one of the timelines, a neutral option indicating no preference for a specific timeline, and a "don't know" option. We encouraged the users to research the timeline entity (e.g. using Wikipedia) before evaluating the timeline pair, if necessary.

Each pair of timelines was rated by three or four users each. Then, majority voting was applied. In total 11 users (graduate Computer Science students) participated in the user evaluation. A user evaluated 42 timeline pairs on average. On average, the users took 69 seconds to decide between two timelines.

We compute the rater preference $RPref$ score adopted from [2] as the fraction of votes for the particular method, based on the annotation that is most frequent among the three users per timeline entity. The results of the user evaluation are presented in Table 21. The timelines generated by our approach with both biographical sources (BS-BIO and BS-ENC) were preferred over the baseline by the users most of the time, for all entity types. For example, all of the 16 timelines for politicians generated by our approach with BS-ENC were preferred over the TM timelines. In total the timelines from BS-BIO were preferred in 67.21% of the cases and the BS-ENC timelines were preferred in 69.35% of the cases.

For BS-BIO, the mean number of ratings favouring our timeline is 1.50 (BS-ENC: 1.58) with a standard deviation of 0.72 (BS-ENC: 0.97), for the TM baseline the mean is 0.40 (BS-ENC: 0.59) with a standard deviation of 0.67 (BS-ENC: 0.74). The results of the paired t-test confirm statistical significance of this result for the confidence level of 99%.

7.4.2. Feature Impact

In total, 411 features are utilised by the model during the timeline generation. In order to better understand the impact of the individual features on the classification task, we compute the correlation between the features and the benchmark judgements using the Pearson Correlation Coefficient ($PCC \in [-1, 1]$, with $PCC = 0$ corresponding to no linear relationship), shown in Table 22.

For both biographical sources, the highest PCC is achieved for the property "born" ($PCC = 0.39$ for BS-ENC, $PCC = 0.25$ for BS-BIO). The "died" property and the time provenance feature TRF-TP are of sim-

Table 21

RPref scores from user ratings for different timeline configurations and entity types. As users could also give a neutral rating or skip a rating, the RPref scores do not necessarily sum up to 100%.

| Biographical Source | BS-BIO | | BS-ENC | |
|-----------------------|---------------|---------------|---------------|---------------|
| | BS-BIO | TM baseline | BS-ENC | TM baseline |
| Actor | 81.82% | 9.09% | 72.73% | 9.09% |
| Athlete | 75.00% | 8.33% | 58.33% | 25.00% |
| Musical Artist | 70.00% | 0.00% | 50.00% | 30.00% |
| Politician | 53.33% | 13.33% | 100.00% | 0.00% |
| Writer | 61.54% | 30.77% | 53.85% | 25% |
| Total | 67.21% | 13.11% | 69.35% | 14.52% |

Table 22

PCC correlation coefficient between top-5 features and the benchmark judgments, sorted by the absolute PCC values.

| Rank | BS-BIO | | BS-ENC | |
|------|-------------------------------------|-------|-------------------------------------|-------|
| | Feature | PCC | Feature | PCC |
| 1 | TRF-PI: <i>born</i> | 0.25 | TRF-PI: <i>born</i> | 0.39 |
| 2 | TF-TP: Time provenance | 0.21 | TRF-PI: <i>died</i> | 0.27 |
| 3 | TRF-PI: <i>died</i> | 0.19 | TF-TP: Time provenance | 0.23 |
| 4 | TRF-MR: Relation mentions rank, EN | -0.19 | TRF-MR: Relation mentions rank, EN | -0.19 |
| 5 | TRF-MR: Relation mentions rank, all | -0.18 | TRF-MR: Relation mentions rank, all | -0.18 |
| ... | | | | |
| 10 | TRF-PI: <i>spouse</i> | 0.13 | TRF-MR: Relation mentions rank, RU | -0.14 |
| ... | | | | |
| 65 | TRF-PI: <i>director</i> | 0.03 | TRF-PI: <i>spouse</i> | 0.03 |
| ... | | | | |
| 410 | TRF-PI: <i>cover artist</i> | 0.00 | TRF-PI: <i>military rank</i> | 0.00 |
| 411 | TRF-PI: <i>illustrator</i> | 0.00 | TRF-PI: <i>draft team</i> | 0.00 |

ilar relevance in both biographical sources, followed by the features related to relation mentions. In contrast, properties like “cover artist” and “draft team” do not correlate with the relation importance. One interesting difference between the biographical sources is the property “spouse” that is highly relevant in the biographical source BS-BIO, but is ranked lower in BS-ENC. Such personal happenings are often not included in Wikipedia’s encyclopedic abstracts.

7.4.3. Relevance of the Timeline Entries

We evaluated the performance of the classification models for predicting the relevance of the individual temporal relations with respect to the benchmarks presented in Section 7.1. The results of this automated evaluation using a 10-fold cross validation are presented in Table 23. In general, our models learned from the training set are generalisable to the test set, reaching F-measure values of 0.827 in the case of BS-ENC and 0.738 for BS-BIO. Across the biographical sources, the usage of all features combined leads to the

best precision and recall scores. The removal of features leads to a decrease in performance: leaving out property labels or the features based on mentions leads to the biggest performance decrease.

7.4.4. Coverage of the Reference Sources

To demonstrate the gain of integrating data from multiple reference sources into EventKG, we assess the coverage of temporal relations in the biographical sources. That means, for each person in our benchmark, we compute the percentage of benchmark relations that are found in the temporal relations of a reference source. Table 24 shows the results, measured by mean coverage per person entity. For example, 27.45% of the relations extracted from BS-ENC can be mapped to a temporal relation in Wikidata. Additionally, we compute the coverage for *extended* reference sources, i.e. we still only consider relations from the specific source, but use the fused information about temporal entities (i.e. existence and happening times) from EventKG.

Table 23

Weighted precision and recall scores for both classes (relevant and irrelevant) for predicting the benchmark labels of the temporal relations using a 10-fold cross validation. Additionally, the F-measure as harmonic mean of precision and recall is reported. † All language-dependent features except for EN are omitted.

| Features | Omitted Features | BS-BIO | | | BS-ENC | | |
|-----------------------------|-----------------------------|-----------|--------|--------------|-----------|--------|--------------|
| | | Precision | Recall | F-Measure | Precision | Recall | F-Measure |
| all features | / | 0.796 | 0.749 | 0.738 | 0.848 | 0.829 | 0.827 |
| no property labels | TRF-PI | 0.753 | 0.691 | 0.671 | 0.822 | 0.802 | 0.799 |
| no mentions | TRF-RM | 0.769 | 0.700 | 0.679 | 0.802 | 0.734 | 0.719 |
| no temporal features | TF-TP, TF-TDS, TF-TDE | 0.795 | 0.747 | 0.736 | 0.847 | 0.829 | 0.827 |
| English only | † | 0.791 | 0.737 | 0.724 | 0.843 | 0.821 | 0.819 |

Table 24

Mean coverage of the temporal relations in the benchmarks per reference source and biographical source.

| | BS-BIO | | BS-ENC | |
|-------------------------------|-------------------|------------------------------|-------------------|------------------------------|
| | Mean coverage (%) | Mean Coverage (%) (extended) | Mean Coverage (%) | Mean Coverage (%) (extended) |
| Wikidata | 14.39 | 16.09 | 36.15 | 38.64 |
| YAGO | 11.96 | 12.34 | 37.90 | 38.40 |
| Wikipedia_{EN} | 0.51 | 14.56 | 0.80 | 23.65 |
| Wikipedia_{FR} | 0.34 | 11.04 | 0.61 | 18.96 |
| Wikipedia_{DE} | 0.16 | 0.86 | 0.40 | 16.66 |
| Wikipedia_{PT} | 0.00 | 8.61 | 0.16 | 15.73 |
| Wikipedia_{RU} | 0.22 | 8.68 | 0.43 | 15.41 |
| Wikipedia | 0.86 | 15.08 | 1.37 | 23.74 |
| DBpedia_{EN} | 5.05 | 9.27 | 27.94 | 34.97 |
| DBpedia_{FR} | 4.10 | 7.27 | 22.01 | 28.40 |
| DBpedia_{DE} | 4.48 | 6.41 | 25.69 | 28.90 |
| DBpedia_{PT} | 0.0 | 2.60 | 0.0 | 4.75 |
| DBpedia_{RU} | 0.0 | 1.48 | 0.0 | 2.64 |
| DBpedia | 5.73 | 14.53 | 30.02 | 45.10 |
| EventKG | 23.29 | — | 55.09 | — |

The results show that there is a higher coverage for BS-ENC than for BS-BIO across all reference sources. This can be explained by the fact that the texts from BS-BIO are longer and less event links are provided: not only does the BS-BIO benchmark rely on named entity recognition, as this source does not contain any links, but events are also harder to recognise as they can be described in several ways (e.g. “first inauguration of Barack Obama” and “Barack Obama was sworn in as the president on January 20, 2009”). In general, YAGO and Wikidata clearly outperform Wikipedia and DBpedia (as DBpedia does not contain statements with validity times). Through the integra-

tion and fusion in EventKG, the coverage increases to more than 50% in BS-ENC.

8. Related Work

In this section, we discuss related work in the areas of event knowledge graphs and the task of biographical timeline generation.

8.1. Event Knowledge Graphs

To the best of our knowledge, currently there are no dedicated knowledge graphs aggregating event-

centric information and temporal relations for historical and contemporary events directly comparable to EventKG. The heterogeneity of data models and vocabularies for event-centric and temporal information (e.g. [7, 14, 15, 23, 34, 35]), the large scale of the existing knowledge graphs, in which events play only an insignificant role, and the lack of clear identification of event-centric information, makes it particularly challenging to identify, extract, fuse and efficiently analyse event-centric and temporal information and make it accessible to real-world applications in an intuitive and unified way. Through the light-weight integration and fusion of event-centric and temporal information from different sources, EventKG enables to increase coverage and completeness of this information. Furthermore, existing sources lack structured information to judge event popularity and relation strength as provided by EventKG – the characteristic that gains the key relevance given the rapidly increasing amount of event-centric and temporal data on the Web and the resulting information overload.

Data models and vocabularies for events: Several data models and the corresponding vocabularies (e.g. [7, 23, 34–36]) provide means to model events. For example, the ECKG model proposed by Rospocher et al. [7] enables fine-grained textual annotations to model events extracted from news collections. CAMEO [36] is a framework to model events extracted from news, in particular in the political domain. The Simple Event Model (SEM) [23], schema.org [35] and the Linking Open Descriptions of Events (LODE) ontology [34] provide means to describe events and interlink them with actors, times and places. In EventKG, we build upon SEM and extend this model to represent a wider range of temporal relations and to provide additional information regarding events.

Extracting event-centric and temporal information: Most approaches for automatic knowledge graph construction and integration focus on entities and related facts rather than events. Examples include DBpedia [11], Freebase [37], YAGO [12] and YAGO+F [38]. In contrast, EventKG is focused on events and temporal relations. In [13], the authors extract event information from WCEP. EventKG builds upon this work to include WCEP events. For the extraction of temporal information, there are several approaches to annotate both textual data [39] and relations [40, 41] with temporal scopes inferred from external sources. In EventKG, we rely on the temporal information already contained in the reference sources, which gives highly precise values as shown in Section 5.2. In-

creasing the coverage for temporal annotations in case of missing values by using external resources is a potential extension for future work.

The question of how to model temporal data is an important question as it comes to considering time expressions of different levels of granularity or with uncertainty. Examples to tackle such issues include the use of multiple potential start and end times as in the temporal slot filling task [42] or adding uncertainty scores to temporal relations [43]. The representation of this information is facilitated through existing relational models [44], the Extended Date-Time Format (EDTF) [24] or with the Time Ontology in OWL [45]. The Simple Event Model adopted in this work supports a simple notion of temporal time spans, which is sufficient to represent temporal information provided by the reference sources of EventKG and is compatible with the time representation in these sources. Nevertheless, we see more advanced time models as a potential future extension, in particular in the context of a possible enrichment of EventKG with additional, and in particular automatically inferred, temporal information.

Extraction of events and facts from news: Recently, the problem of building knowledge graphs and datasets directly from plain text news articles [7, 8, 16, 17], and extraction of named events from news [14, 46] have been addressed. These approaches apply Open Information Extraction methods and develop them further to address specific challenges in the event extraction in the news domain. State-of-the-art approaches that automatically extract events from news potentially obtain noisy and unreliable results (e.g. the state-of-the-art extraction approach in [7] reports an accuracy of only 0.551). Furthermore, such systems provide billions of events at a very high granularity level, as typically represented in news articles. Compared to the established knowledge repositories such as DBpedia or Wikidata, such events indicate significant differences in the representation accuracy and event granularity. In contrast, contemporary events included in EventKG originate from high quality community curated sources such as WCEP and Wikipedia event lists and represent significant societal happenings at a different granularity and abstraction level, compared to news sources.

8.2. Biographical Timeline Generation

Existing work on timeline generation from knowledge graphs has mainly focused on the selection of relevant events or relations. The works of Althoff et al.

[2] and Tuan et al. [47] come closest to our task definition. In [2], the authors create timelines for politicians, actors and athletes from the Freebase knowledge graph, adding visual and diversity constraints on the generated timelines. In [47], person timelines are generated by ranking relations extracted from Wikipedia and YAGO knowledge graphs. Similarly, in [27] entity summarisation is created based on link counts, but without taking temporal data into account. In difference to our work, in both these approaches the feature weights are handcrafted and no machine learning is involved. [31] and [32] aim at generating biographies in a natural language, that means to generate textual summaries for people, by mapping facts from knowledge graphs to one-sentence biographies. Both works incorporate neural models to learn text, but the biographies are limited to few facts such as birth dates and entity types.

Other approaches generate timelines for different use cases, for example to get an overview over news articles over a large time span [48, 49] or for depicting singular events such as football matches in a very fine-grained manner [50]. For visualisation, there are approaches to transform relationship paths from knowledge graphs into sentences [2, 51] and different interaction models that let a user explore the timeline [2, 49, 52]. In this article, we focus on the generation of timelines containing relevant temporal relations and do not limit the approach by any visual constraints. This way, the models obtained by our methods can be used in a broader range of interfaces and application scenarios.

One important subtask of the timeline generation is to judge whether a temporal relation is relevant in a certain context. This task has been addressed by other works using classification and ranking approaches. For example, to rank news articles related to a query entity, Singh et al. [53] employ a diversified ranking model based both on the aspect and temporal dimension. Approaches such as the one proposed by Setty et al. [54] impose methods to rank the importance of events, but without taking into account the specific timeline entity. In comparison to these approaches, the task addressed in our work is more specific, as it considers the relevance of individual temporal relations to a timeline entity.

Further methods to access semantic information included in knowledge graphs in an intuitive way include question answering applications (e.g. [21],[4]) and interactive query construction interfaces proposed in our previous work (e.g. [55], [22]). Application of these

approaches to EventKG is an interesting direction for future research.

9. Conclusions

In this article we presented the concept of a temporal knowledge graph that interconnects real-world entities and events using temporal relations. Furthermore, we presented an instantiation of the temporal knowledge graph - EventKG. EventKG is a multilingual knowledge graph that integrates and harmonises event-centric and temporal information regarding historical and contemporary events. EventKG V1.1 includes over 690 thousand event resources and over 2.3 million temporal relations. Unique EventKG features include the light-weight integration and fusion of structured and semi-structured multilingual event representations and temporal relations in a single knowledge graph, as well as the provision of information to facilitate assessment of relation strength and event popularity, while providing provenance. The light-weight integration enables to significantly increase the coverage and completeness of the included event representations, in particular with respect to time and location information.

We analysed the characteristics of the resulting knowledge graph and observed a significant increase in coverage compared to the reference sources. For example, EventKG V1.1 contains 50K more events than identified in Wikidata and more than 262K events than identified in the English DBpedia. Additionally, 360K events are extracted from semi-structured sources. The quality of this resulting dataset was confirmed in a manual evaluation. This evaluation indicated high precision for the event identification step (with an average precision of 96%), the time fusion step (with precision of 75% for the events that had a disagreement regarding their time information in the reference sources) and the precision of the location fusion (94.31%).

Furthermore, in this article we addressed the problem of biographical timeline generation from a temporal knowledge graph. In order to generate biographical timelines from a large-scale temporal knowledge graph, we proposed a method based on distant supervision. This method uses features extracted from the temporal knowledge graph as well as a benchmark extracted from external biographical sources to train an effective relevance model. Our evaluation includes the results of a user study and an automatic evaluation and demonstrates the effectiveness of the pro-

posed method. Our method significantly outperforms the baseline in the biography generation. According to the rater preference score, our method achieves 68% on average, in contrast to the baseline that achieves only 14%.

We make the datasets described in this article publicly available to stimulate further research in this area.

In the future work, we plan to further extend EventKG to include additional sources. We would also like to explore the development of further methods and applications using EventKG.

Acknowledgements

This work was partially funded by the ERC ("ALEXANDRIA", 339233), the BMBF ("Data4UrbanMobility", 02K15A040) and H2020-MSCA-ITN-2018-812997 "Cleopatra".

References

- [1] J. Allan, R. Papka and V. Lavrenko, On-Line New Event Detection and Tracking, *SIGIR Forum* **51**(2) (2017), 185–193.
- [2] T. Althoff, X.L. Dong, K. Murphy, S. Alai, V. Dang and W. Zhang, TimeMachine: Timeline Generation for Knowledge-base Entities, in: *Proceedings of the 21th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 2015, pp. 19–28.
- [3] S. Gottschalk and E. Demidova, EventKG+TL: Creating Cross-Lingual Timelines from an Event-Centric Knowledge Graph, in: *Proceedings of the European Semantic Web Conference*, 2018.
- [4] K. Höffner and et al., Survey on challenges of Question Answering in the Semantic Web, *Semantic Web* **8**(6) (2017), 895–920.
- [5] S. Gottschalk, E. Demidova, V. Bernacchi and R. Rogers, Ongoing Events in Wikipedia: A Cross-lingual Case Study, in: *Proceedings of the 2017 ACM on Web Science Conference*, 2017, pp. 387–388.
- [6] R. Rogers, *Digital Methods*, MIT Press, 2013.
- [7] M. Rospocher and et al., Building Event-centric Knowledge Graphs from News, *Web Semantics* **37** (2016), 132–151.
- [8] M. Al-Badrashiny, J. Bolton, A.T. Chaganty, K. Clark, C. Harman, L. Huang, M. Lamm, J. Lei, D. Lu, X. Pan et al., TinkerBell: Cross-lingual Cold-Start Knowledge Base Construction, in: *TAC*, 2017.
- [9] M. Färber, F. Bartscherer, C. Menne and A. Rettinger, Linked Data Quality of DBpedia, Freebase, OpenCyc, Wikidata, and YAGO, *Semantic Web* (2016), 1–53.
- [10] F. Erxleben, M. Günther, M. Krötzsch, J. Mendez and D. Vrandečić, Introducing Wikidata to the Linked Data Web, in: *Proceedings of the International Semantic Web Conference*, 2014, pp. 50–65.
- [11] J. Lehmann and et al., DBpedia - A Large-scale, Multilingual Knowledge Base Extracted from Wikipedia, *Semantic Web Journal* **6**(2) (2015), 167–195.
- [12] F. Mahdisoltani, J. Biega and F. Suchanek, YAGO3: A Knowledge Base from Multilingual Wikipedias, in: *Proceedings of the Conference on Innovative Data Systems Research*, 2014.
- [13] G.B. Tran and M. Alrifai, Indexing and Analyzing Wikipedia's Current Events Portal, the Daily News Summaries by the Crowd, in: *Proceedings of the 23rd International Conference on World Wide Web*, 2014, pp. 511–516.
- [14] Q. Yuan, X. Ren, W. He, C. Zhang, X. Geng, L. Huang, H. Ji, C.-Y. Lin and J. Han, Open-Schema Event Profiling for Massive News Corpora, in: *Proceedings of the 27th ACM International Conference on Information and Knowledge Management*, ACM, 2018, pp. 587–596.
- [15] R.E. Prasojo, M. Kacimi and W. Nutt, StuffIE: Semantic Tagging of Unlabeled Facets Using Fine-Grained Information Extraction, in: *Proceedings of the 27th ACM International Conference on Information and Knowledge Management*, ACM, 2018, pp. 467–476.
- [16] K. Leetaru and P.A. Schrodt, GDELT: Global Data on Events, Location and Tone, 1979–2012, in: *ISA Annual Convention*, Vol. 2, Citeseer, 2013, pp. 1–49.
- [17] E. Boschee, J. Lautenschlager, S. O'Brien, S. Shellman, J. Starz and M. Ward, ICEWS Coded Event Data, *Harvard Dataverse* **12** (2015).
- [18] S. Gottschalk and E. Demidova, EventKG: A Multilingual Event-Centric Temporal Knowledge Graph, in: *Proceedings of the European Semantic Web Conference*, 2018, pp. 272–287.
- [19] S. Gottschalk and E. Demidova, MultiWiki: Interlingual Text Passage Alignment in Wikipedia, *TWEB* **11**(1) (2017), 6–1630.
- [20] S. Gottschalk, V. Bernacchi, R. Rogers and E. Demidova, Towards better Understanding Researcher Strategies in Cross-lingual Event Analytics, in: *Proceedings of the 22nd International Conference on Theory and Practice of Digital Libraries*, 2018.
- [21] W. Zheng, H. Cheng, L. Zou, J.X. Yu and K. Zhao, Natural Language Question/Answering: Let Users Talk With The Knowledge Graph, in: *Proceedings of the 2017 ACM on Conference on Information and Knowledge Management*, CIKM '17, 2017, pp. 217–226.
- [22] E. Demidova, X. Zhou and W. Nejdl, Efficient Query Construction for Large Scale Data, in: *Proceedings of the 36th international ACM SIGIR Conference on Research and Development in Information Retrieval*, 2013, pp. 573–582.
- [23] W.R. Van Hage, V. Malaisé, R. Segers, L. Hollink and G. Schreiber, Design and Use of the Simple Event Model (SEM), *Web Semantics* **9**(2) (2011), 128–136.
- [24] Library of Congress, Extended Date Time Format (EDTF) Specification, Accessed: 2018-11-19.
- [25] D. Hienert, D. Wegener and H. Paulheim, Automatic Classification and Relationship Extraction for Multi-Lingual and Multi-Granular Events from Wikipedia, in: *Proceedings of the Workshop on Detection, Representation, and Exploitation of Events in the Semantic Web*, 2012.
- [26] Q. McNemar, Note on the Sampling Error of the Difference between Correlated Proportions or Percentages, *Psychometrika* **12**(2) (1947), 153–157.
- [27] A. Thalhammer, N. Lasierra and A. Rettinger, LinkSUM: Using Link Analysis to Summarize Entity Data, in: *Proceedings of the International Conference on Web Engineering*, 2016, pp. 244–261.

- [28] C.D. Manning, M. Surdeanu, J. Bauer, J.R. Finkel, S. Bethard and D. McClosky, The Stanford CoreNLP Natural Language Processing Toolkit, in: *Proceedings of 52nd Annual Meeting of the Association for Computational Linguistics: System Demonstrations*, 2014, pp. 55–60.
- [29] J. Strötgen and M. Gertz, HeidelTime: High Quality Rule-based Extraction and Normalization of Temporal Expressions, in: *Proceedings of the 5th International Workshop on Semantic Evaluation*, 2010, pp. 321–324.
- [30] P.N. Mendes, M. Jakob, A. García-Silva and C. Bizer, DBpedia Spotlight: Shedding Light on the Web of Documents, in: *Proceedings of the 7th International Conference on Semantic Systems*, 2011, pp. 1–8.
- [31] A. Chisholm, W. Radford and B. Hachey, Learning to Generate One-sentence Biographies from Wikidata, in: *Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics*, 2017.
- [32] R. Lebrecht, D. Grangier and M. Auli, Neural Text Generation from Structured Data with Application to the Biography Domain, in: *Proceedings of the Conference on Empirical Methods in Natural Language Processing*, 2016.
- [33] I.H. Witten, E. Frank, M.A. Hall and C.J. Pal, *Data Mining: Practical Machine Learning Tools and Techniques*, Morgan Kaufmann, 2016.
- [34] R. Shaw, R. Troncy and L. Hardman, LOD: Linking Open Descriptions of Events, in: *Proceedings of the Asian Semantic Web Conference*, 2009, pp. 153–167.
- [35] R. Guha, Introducing schema.org: Search Engines Come Together for a Richer Web, *Google Official Blog* (2011).
- [36] P.A. Schrodt, CAMEO: Conflict and Mediation Event Observations Event and Actor Codebook (2012), Pennsylvania State University.
- [37] K. Bollacker and et al., Freebase: A Collaboratively Created Graph Database for Structuring Human Knowledge, in: *Proceedings of the ACM SIGMOD International Conference on Management of Data*, 2008, pp. 1247–1250.
- [38] E. Demidova, I. Oelze and W. Nejdl, Aligning Freebase with the YAGO Ontology, in: *Proceedings of the 22nd ACM International Conference on Information and Knowledge Management*, 2013, pp. 579–588.
- [39] E. Kuzey, V. Setty, J. Strötgen and G. Weikum, As Time Goes By: Comprehensive Tagging of Textual Phrases with Temporal Scopes, in: *Proceedings of the 25th International conference on World Wide Web*, 2016, pp. 915–925.
- [40] A. Rula, M. Palmonari, A.-C.N. Ngomo, D. Gerber, J. Lehmann and L. Bühmann, Hybrid Acquisition of Temporal Scopes for RDF Data, in: *Proceedings of the European Semantic Web Conference*, 2014, pp. 488–503.
- [41] P.P. Talukdar, D. Wijaya and T. Mitchell, Coupled Temporal Scoping of Relational Facts, in: *Proceedings of the 6th ACM International Conference on Web Search and Data Mining*, 2012, pp. 73–82.
- [42] M. Surdeanu, Overview of the TAC2013 Knowledge Base Population Evaluation: English Slot Filling and Temporal Slot Filling., in: *TAC*, 2013.
- [43] M.W. Chekol, G. Pirrò, J. Schoenfish and H. Stuckenschmidt, Marrying Uncertainty and Time in Knowledge Graphs, in: *Proceedings of the 31st AAAI Conference on Artificial Intelligence*, 2017, pp. 88–94.
- [44] M.W. Chekol, Scaling Probabilistic Temporal Query Evaluation, in: *Proceedings of the 2017 ACM on Conference on Information and Knowledge Management*, ACM, 2017, pp. 697–706.
- [45] J.R. Hobbs and F. Pan, Time Ontology in OWL, *W3C working draft 27* (2006), 133.
- [46] E. Kuzey, J. Vreeken and G. Weikum, A Fresh Look on Knowledge Bases: Distilling Named Events from News, in: *Proceedings of the 23rd ACM International Conference on Conference on Information and Knowledge Management*, 2014, pp. 1689–1698.
- [47] T.A. Tuan, S. Elbassuoni, N. Preda and G. Weikum, CATE: Context-aware Timeline for Entity Illustration, in: *Proceedings of the 20th International Conference Companion on World Wide Web*, 2011, pp. 269–272.
- [48] G. Tran, M. Alrifai and E. Herder, Timeline Summarization from Relevant Headlines, in: *Proceedings of the European Conference on Information Retrieval*, 2015, pp. 245–256.
- [49] R. Swan and J. Allan, Automatic Generation of Overview Timelines, in: *Proceedings of the 23rd annual international ACM SIGIR conference on Research and development in information retrieval*, 2000, pp. 49–56.
- [50] O. Alonso and K. Shiells, Timelines as summaries of popular scheduled events, in: *Proceedings of the 22nd International Conference on World Wide Web*, 2013, pp. 1037–1044.
- [51] N. Voskarides, E. Meij and M. de Rijke, Generating Descriptions of Entity Relationships, in: *Proceedings of the European Conference on Information Retrieval*, 2017, pp. 317–330.
- [52] J. Zhao, S.M. Drucker, D. Fisher and D. Brinkman, TimeSlice: Interactive Faceted Browsing of Timeline Data, in: *Proceedings of the International Working Conference on Advanced Visual Interfaces*, 2012, pp. 433–436.
- [53] J. Singh, W. Nejdl and A. Anand, History by Diversity: Helping Historians Search News Archives, in: *Proceedings of the 2016 ACM on Conference on Human Information Interaction and Retrieval*, 2016, pp. 183–192.
- [54] V. Setty, A. Anand, A. Mishra and A. Anand, Modeling Event Importance for Ranking Daily News Events, in: *Proceedings of the Tenth ACM International Conference on Web Search and Data Mining*, 2017, pp. 231–240.
- [55] E. Demidova, X. Zhou and W. Nejdl, FreeQ: An Interactive Query Interface for Freebase, in: *Proceedings of the 21st World Wide Web Conference, WWW 2012, Lyon, France, April 16-20, 2012 (Companion Volume)*, 2012, pp. 325–328. doi:10.1145/2187980.2188040. <https://doi.org/10.1145/2187980.2188040>.

Appendix A. Example Queries

Here, we present example SPARQL queries to illustrate the retrieval of particular event and entity characteristics.

A.1. Query 1: Provenance and Event Locations

The SPARQL query in Listing 1 uses the named graph notation to find the locations of the event “Second inauguration of Barack Obama” in any source. This is done using the `sem:hasPlace` predicate introduced in Section 4.1. Table 25 lists the query results. While YAGO has the United States Capitol and Washington D.C. as location, Wikidata has Washington D.C. only. There are no locations for this event found in any of the DBpedia language editions. After fusion, the union of potential locations (United States Capitol, Washington, D.C.) is reduced to the United States Capitol only, which is located in Washington D.C.¹⁷. Fused locations are placed within EventKG’s named graph.

Table 25

Locations of the first inauguration of Barack Obama in EventKG.

| ?location | ?named_graph |
|---------------------------|--------------------|
| dbr:United_States_Capitol | eventKG-g:event_kg |
| dbr:Washington,_D.C. | eventKG-g:wikidata |
| dbr:United_States_Capitol | eventKG-g:yago |
| dbr:Washington,_D.C. | eventKG-g:yago |

A.2. Query 2: Important Events of an Entity

The second query shown in Listing 2 employs the relation strength information contained in EventKG. It returns a list of events connected to Barack Obama, sorted by the number of common mentions (`eventKG-s:mentions`) with Barack Obama in the English Wikipedia (`GRAPH eventKG-g:wikipedia_en`). Additionally, if there is an event start date available, this is returned as well, using the named EventKG graph to retrieve the fused date. The results in Table 26 reveal that the United States presidential election of 2008 is the event mentioned most often together with Barack Obama.

Table 26

Events that are most often mentioned together with Barack Obama.

| ?event | ?cnt | ?startDate |
|---|------|------------|
| dbr:United_States_presidential_election,_2008 | 719 | 2008-11-04 |
| dbr:United_States_presidential_election_in_New_Jersey,_2012 | 530 | 2012-11-06 |
| dbr:United_States_presidential_election_in_New_Jersey,_2008 | 522 | 2008-11-04 |
| ⋮ | | |
| dbr:First_inauguration_of_Barack_Obama | 68 | 2009-01-20 |

¹⁷This information could be inferred using `so:containedInPlace*`.

```

SELECT ?location ?named_graph

WHERE {
  ?event owl:sameAs dbr:First_inauguration_of_Barack_Obama .

  GRAPH ?named_graph {
    ?event sem:hasPlace ?loc
  } .

  GRAPH eventKG-g:dbpedia_en
    ?loc owl:sameAs ?location .
}
ORDER BY ?named_graph

```

Listing 1: SPARQL query for retrieving the locations of the first inauguration of Barack Obama using `sem:hasPlace`, together with their named graph for provenance information.

```

SELECT ?event ?cnt ?startDate

WHERE {
  ?obama owl:sameAs dbr:Barack_Obama .
  ?relation rdf:subject ?obama .
  ?relation rdf:object ?eventEKG .

  GRAPH eventKG-g:wikipedia_en {
    ?relation eventKG-s:mentions ?cnt .
  }

  ?eventEKG rdf:type sem:Event .

  GRAPH eventKG-g:dbpedia_en {
    ?eventEKG owl:sameAs ?event
  } .

  OPTIONAL {
    GRAPH eventKG-g:event_kg {
      ?eventEKG sem:hasBeginTimeStamp ?startDate
    }
  } .
}
ORDER BY DESC (?cnt)

```

Listing 2: SPARQL query for retrieving the events that are most often mentioned together with Barack Obama. Instances of `eventKG-s:Relation` are searched who are connected to Barack Obama as their subject and an instance of `sem:Event` as their object.