

EventKG+Click: A Dataset of Language-specific Event-centric User Interaction Traces

Sara Abdollahi, Simon Gottschalk, and Elena Demidova

L3S Research Center, Leibniz Universität Hannover, Germany
{abdollahi,gottschalk,demidova}@L3S.de

Abstract. An increasing need to analyse event-centric cross-lingual information calls for innovative user interaction models that assist users in crossing the language barrier. However, datasets that reflect user interaction traces in cross-lingual settings required to train and evaluate the user interaction models are mostly missing. In this paper, we present the EventKG+Click dataset that aims to facilitate the creation and evaluation of such interaction models. EventKG+Click builds upon the event-centric EventKG knowledge graph and language-specific information on user interactions with events, entities, and their relations derived from the Wikipedia clickstream.

1 Introduction

With a rapidly growing number of events with significant international impact, cross-lingual analytics gains increased importance for researchers and professionals in many disciplines, including digital humanities, media studies, and journalism. The most prominent recent examples of such events include the COVID-19 outbreak, the migration crisis in Europe, and Brexit. From the information science perspective, research on event-centric information spread across languages and communities, as well as cross-cultural and cross-lingual differences in reporting, are of particular interest. However, very often, the language barrier hinders such research.

The development of novel methods for user interaction with event-centric cross-lingual information can help to overcome the language barrier in this context. Such methods can facilitate researchers with limited knowledge of target languages to narrow down the search space and to obtain an overview of the cross-lingual differences effectively and efficiently. However, currently, user interaction in multilingual settings is not sufficiently studied. The benchmarks and datasets suitable for the evaluation of new methods for user interaction with cross-lingual information are mostly missing.

With the recent development of knowledge graphs that provide cross-lingual information, such as Wikidata, DBpedia, and the event-centric EventKG knowledge graph [6], the availability of semantic event-centric cross-lingual informa-

tion has significantly increased. These knowledge graphs contain semantic information regarding events and their relations while providing labels in different languages along with the properties extracted from language-specific sources. For example, EventKG, in its version 2.1 released in February 2020, includes information on more than 1,200,000 events in nine languages. We believe that knowledge graphs containing event-centric cross-lingual data can build a backbone for the development of user interaction methods that can assist users in crossing the language barrier.

In this paper, we present a novel cross-lingual dataset that reflects the language-specific relevance of events and their relations. This dataset aims to provide a reference source to train and evaluate novel models for event-centric cross-lingual user interaction, with a particular focus on the models supported by knowledge graphs. Our dataset EventKG+Click is based on two data sources: 1) the Wikipedia clickstream² that reflects real-world user interactions with events and their relations within language-specific Wikipedia editions; and 2) the EventKG knowledge graph that contains semantic information regarding events and their relations that partially originates from Wikipedia. EventKG+Click is available online³ to enable further analyses and applications.

Without loss of generality, we adopt a language-specific event ranking as an envisioned user interaction paradigm to illustrate our discussion. For example, Table 1 reveals the different language-specific focus when ranking events. In each of the three languages contained in EventKG+Click, the list of most language-specific related events is clearly representing language-specific views (e.g., “2016 Berlin truck attack” for German) that can be used for further exploration of events from language-specific viewpoints. In the case of English, we see that the Southeast Asian Games are of high language-specific relevance, which can be explained by the large percentage of Asian users of the English Wikipedia⁴.

Table 1. Events with highest language-specific relevance per language in EventKG+Click.

Rank	English	German	Russian
1	Southeast Asian Games	2016 Berlin truck attack	2009 Russian Premier League
2	2017 Southeast Asian Games	German student movement	1993 Russian Top League
3	2014 United States Senate elections	2006 Austrian legislative election	2012–13 Russian Premier League

In EventKG+Click, we enrich the information obtained from the Wikipedia clickstream with event and entity references from EventKG. Furthermore, we

² https://meta.wikimedia.org/wiki/Research:Wikipedia_clickstream

³ <https://github.com/saraabdollahi/EventKG-Click>

⁴ <https://stats.wikimedia.org/wikimedia/squids/SquidReportPageViewsPerCountryBreakdown.htm>

create a cross-lingual view on the clickstream by combining information obtained from three Wikipedia language editions, namely English, German, and Russian. Moreover, we compute scores that reflect the language-specific relevance of events and their relations, as indicated by the user interactions in the clickstream. Finally, to support further development of the event-centric user interaction methods in the cross-lingual settings, we analyse the correlations of the proposed scoring function and selected influence factors.

We structure the rest of the paper as follows: First, we review related work regarding cross-lingual analytics, knowledge graphs, and the Wikipedia clickstream in Section 2. Then, we introduce our EventKG+Click dataset in Section 3. In Section 4, we propose scores to represent the language-specific relevance of events and their relations. Given the EventKG+Click dataset and these scores, we analyse how selected factors influence the language-specific relevance in Section 5. Finally, we provide a conclusion in Section 6.

2 Related Work

In this section, we briefly summarise related work in the areas of cross-lingual analytics, along with the aspects related to knowledge graphs and the Wikipedia clickstream.

Cross-lingual analytics and interaction. With the rise of the Web, there came an uprise of user-generated content accessible over the whole world, leading to knowledge diversity across languages [7]. The identification and analysis of such knowledge diversity is an important method to understand language communities better. For example, Oeberst et al. identified different types of "collective biases" such as biased representations of intergroup conflicts that appear under collaborative circumstances [14]. Miz et al. identified how Wikipedia reflects cultural particularities [11]. Mocanu et al. have identified linguistic trends in Twitter usage in more than 100 countries [12].

In the context of cross-lingual analytics, events play a particularly important role: When an event breaks out, this event is usually reported by a large number of sources, whose coverage highly varies across language communities [3]. This phenomenon becomes visible when using EventRegistry, a tool that allows cross-lingual exploration of news articles which are assigned to event clusters [15]. Event-centric cross-lingual analytics are also viable across different Wikipedia language editions as illustrated by two case studies about the Brexit and the US withdrawal from the Paris Agreement, where researchers identified language-specific viewpoints [4].

With EventKG+Click, our goal is to promote further cross-lingual analytics and interaction, facilitated by a combination of semantic information given in knowledge graphs and user interaction traces obtained from a clickstream.

Knowledge graphs. An essential resource to facilitate interaction with cross-lingual information are knowledge graphs, in particular those containing language-specific labels and relations. Kaffee et al. [8] developed metrics that measure the multilingualism of knowledge graphs to identify those suitable for

usage in multilingual applications and to gain cross-lingual insights. For example, Marie et al. [10] discovered a "cultural prism" between the different DBpedia language editions when querying for entities related to facets of interest.

The importance of multilingualism in knowledge graphs becomes even more evident in the case of event-based applications. EventKG [6] is a knowledge graph that is tailored not only to the interaction with event-centric information but also contains information coming from several languages. An example application that makes use of this cross-lingual event knowledge is EventKG+TL [5] that relies on Wikipedia link counts present in EventKG to model the importance of events related to a given concept.

In our analysis, we observed that the closeness of event locations extracted from EventKG is an essential indicator to explain language-specific relevance. Thus, we confirm the importance of event-centric and multilingual knowledge graphs in the context of cross-lingual analytics.

Wikipedia clickstream. The Wikipedia clickstream has been used as a ground-truth to evaluate entity recommendation and relatedness in several examples, as it reveals the navigational behaviour of users and their preferences while exploring Wikipedia pages. Existing work, however, has not considered language-specific differences and mainly focused on the English Wikipedia clickstream: For example, Tran et al. used the English Wikipedia clickstream as ground truth for constructing entity-context queries [16] and Bhatia et al. constructed their query dataset based on the English Wikipedia clickstream [1]. Nguan et al. evaluated their relatedness ranking method by using the raw number of navigations in Wikipedia clickstream [13]. With the usage of the Wikipedia clickstream in different languages, EventKG+Click adds a new perspective onto EventKG, as it reflects real user behaviour across language communities, which goes beyond the consideration of knowledge graph relations and Wikipedia link counts.

3 EventKG+Click Dataset

The Wikipedia clickstream holds the interaction of real users with the articles representing events and entities in the specific Wikipedia language editions and their relations. In particular, the clickstream contains the counts of the (source, target) pairs extracted from Wikipedia's request logs. The clickstream contains all the requests to a Wikipedia page, including links from and to external web pages. As EventKG+Click and our analysis are based on Wikipedia click behaviour, we only consider those (source, target) click pairs in the clickstream where both the source and target are Wikipedia articles connected by a hyperlink.

In this work, we adopt the Wikipedia clickstream that covers the period from December 1, 2019, to December 31, 2019, and contains nearly 19,521,580 click pairs for the English, 2,902,878 click pairs for the German, and 2,752,340 click pairs for the Russian Wikipedia.

EventKG is an event-centric knowledge graph that contains more than 1.2 million events and more than 4 million temporal relations in nine languages in its release from February 2020. Knowledge graphs such as EventKG, DBpedia, and Wikidata include information extracted from the multilingual Wikipedia as the basis. This way, data regarding user interaction with Wikipedia articles and links, available from the Wikipedia clickstream dataset, can be directly mapped to the events, entities and their relations in these knowledge graphs.

When creating the proposed EventKG+Click dataset, we assume that: 1) the events of global importance are reflected in Wikipedia clickstreams of several languages, and 2) a clickstream in a specific language reflects the importance of events and their relations as perceived by the users of the specific Wikipedia language edition. Based on these assumptions, we employ the intersection of language-specific clickstreams to build a dataset for training and evaluation of cross-lingual user interaction. In particular, we map the events and entities included in the Wikipedia clickstream to EventKG and extract relations for these events from all language-specific clickstreams. Furthermore, we compute scores that represent the language-specific relevance of events and their relations. These scores are presented in Section 4. To enable further cross-lingual analysis, we enrich EventKG+Click with several influence factors extracted from EventKG and Wikipedia, which are presented in Section 5.

In EventKG+Click, we only consider entities that are clicked at least 10 times per language, so that we capture those entities that are of global importance and do not consider entities solely present in single Wikipedia language versions. We also only consider pairs which exist in the clickstreams of all considered languages and in which the target page is an event.

The resulting EventKG+Click dataset is available online⁵ and contains relevance scores for more than 4 thousand events, and nearly 10 thousand event-centric click-through pairs.

4 Scores to Assess Language-Specific Relevance

To allow cross-lingual analytics with EventKG+Click, we need to capture the language-specific relevance of events and their relations. Based on the Wikipedia clickstream, we propose two scores that rule out language-independent relevance.

To describe our scores, we first define the concepts used for the computation:

- L is the set of languages under consideration. The current release of EventKG+Click comes in English, German, and Russian: $L = \{EN, DE, RU\}$.
- E is the set of entities contained in EventKG+Click, that are all represented by specific Wikipedia pages and EventKG resources. Formally, named events considered in this work are a specific type of entity and thus included in E .
- $clicks(e_s, e_t, l)$ represents the number of clicks from the source entity $e_s \in E$ to the target event $e_t \in E$ in the clickstream of the given language $l \in L$.

⁵ <https://github.com/saraabdollahi/EventKG-Click>

We distinguish between two scores defined in the following: language-specific event relevance and language-specific relation relevance.

4.1 Language-specific Event Relevance

Wikipedia language versions differ a lot concerning the number of their active users, edits, and articles. For example, the English Wikipedia has 7.2 times as many active users as the German Wikipedia⁶. The clickstream also reflects this imbalance: There are 7 times more clicks in the English clickstream than in the German one. To observe language-specific behaviour, we first need to level the effects that originate from the popularity of the specific Wikipedia language versions. To do so, we normalise the number of clicks with respect to the total number of clicks in the respective language, which leads to normalised scores in the range $[0, 1]$. In order to create balanced click counts, we then multiply the normalised score by the total number of clicks in the clickstreams, as follows:

$$balanced_clicks(e_s, e_t, l) = clicks(e_s, e_t, l) \cdot \frac{\sum_{l' \in L} \sum_{e'_s \in E} \sum_{e'_t \in E} clicks(e'_s, e'_t, l')}{\sum_{e'_s \in E} \sum_{e'_t \in E} clicks(e'_s, e'_t, l)}$$

The popularity of an event can be inferred by the number of user interactions with its Wikipedia page. That way, we can identify the most popular events in a given language $l \in L$ by summing up all clicks from and to an event $e \in E$:

$$balanced_clicks(e, l) = \sum_{e_t \in E} balanced_clicks(e, e_t, l) + \sum_{e_s \in E} balanced_clicks(e_s, e, l)$$

As we focus on the language-specific relevance in EventKG+Click, we need to rule out the events that are highly ranked across all languages under consideration. Therefore, we normalise the language-specific click count by the overall number of clicks in all languages:

$$event_relevance(e, l) = \frac{balanced_clicks(e, l)}{\sum_{l' \in L} balanced_clicks(e, l')} \in [0, 1]$$

With this relevance score, events that are clicked often in a given language $l \in L$, but rarely clicked in the other languages are assigned a relevance score close to 1.

4.2 Language-specific Relation Relevance

To identify events relevant to a given source entity, we define the language-specific relation relevance score. This score assigns a relevance score to the relation between a source entity e_s and a target event e_t in a given language.

⁶ https://en.wikipedia.org/wiki/List_of_Wikipedias

Similarly to the language-specific event relevance, the language-specific relation relevance is computed as the fraction of clicks in the given language compared to all languages:

$$relation_relevance(e_s, e_t, l) = \frac{balanced_clicks(e_s, e_t, l)}{\sum_{l' \in L} balanced_clicks(e_s, e_t, l')} \in [0, 1]$$

Note that this score rules out the effects resulting from the relevance of the source entity: Events that are highly related to an entity e can obtain relevance scores close to 1 independent of e 's click count.

4.3 Examples of Scores

In Table 1 in Section 1, we have given an example of the language-specific event relevance, i.e., that table provides the top-ranked events per language, according to our language-specific event relevance score. As discussed before, we can clearly observe events which are intuitively important for the respective language community.

Table 2. The three events most relevant to the 2012 Summer Olympics for English, German and Russian.

Rank	English	German	Russian
1	2012 Summer Olympics opening ceremony	Olympic Games	2020 Summer Olympics opening ceremony
2	Swimming at the 2012 Summer Olympics	Modern pentathlon at the 2012 Summer Olympics	Olympic Games
3	Badminton at the 2012 Summer Olympics	Equestrian at the 2012 Summer Olympics	Weightlifting at the 2012 Summer Olympics

Table 2 presents the language-specific relation relevance by showing the concrete example of events relevant to the Summer Olympics in 2012. According to our score, the opening ceremony that happened in London is the most relevant event for the 2012 Summer Olympics from the English perspective. Apart from that, we can observe that sports particularly popular in a specific language community are ranked higher (e.g., swimming for English, equestrian sports for German, and weightlifting for Russian).

Our examples illustrate that user click behaviour is not only based on globally relevant entities but takes the language-specific relevance into account. Both relevance scores can be used in language-specific contexts, e.g. for event retrieval or recommendation.

5 Influence Factors for Language-Specific Relevance

Given the EventKG+Click dataset with the relevance scores defined in the previous section, we now discuss several *influence factors* that can potentially im-

pact the language-specific relevance of events and analyse their correlations with the proposed relevance scores. As influence factors we consider language community relevance, event location closeness and event recency, as defined in the following. In future work, we plan to investigate the role of further influence factors, as for example the event type that has been shown to influence the click-behaviour [2].

5.1 Language Community Relevance

The language community relevance factor reflects the importance of an event for the community that speaks this language. We assume that events relevant for the language community should be mentioned and referred to more often in a language-specific corpus.

Based on this assumption, we measure the language community relevance by counting the links to the event article and mentions of the event within the specific Wikipedia language edition⁷. Dependent on the context (i.e., event or relation relevance), we make use of two influence factors:

- Links pointing to the event: The number of links in the whole Wikipedia language edition that link to the event article.
- Co-mentions of a relation: The number of sentences in the whole Wikipedia language edition that jointly mentions the (source, target) pair participating in the relation.

5.2 Event Location Closeness

The event location closeness factor expresses the intuition that users are likely to be interested in the exploration of local events, i.e., events located in spatial proximity of the user. To reflect this intuition, we introduce a binary influence factor that indicates whether an event happened in a location where the respective language $l \in L$ is an official language. For example, the Battle of Stalingrad may be particularly important from the Russian perspective in the context of the Second World War. To compute this factor, we first identify event location(s) using the `sem:hasPlace`⁸ property of EventKG and then derive the official languages of the location’s country.

5.3 Event Recency

Wikipedia is heavily influenced by recent events: Users tend to edit and read articles about events that are happening right now [9]. To observe the impact of recency on the language-specific user click behaviour, we introduce a recency score, which is computed as the number of days between the event start date and the start date of the clickstream dataset (the dates of the specific entries in the dataset are not available). To identify the event start dates, we use `sem:hasBeginTimeStamp` values in EventKG.

⁷ We derive these counts from EventKG that contains link and mention counts [6].

⁸ <http://semanticweb.cs.vu.nl/2009/11/sem/hasPlace>

5.4 Correlations with Influence Factors

Given EventKG+Click and the influence factors, we now investigate the correlations between such influence factors and the language-specific relevance scores. To this end, we compute the Pearson correlation coefficients in several configurations.

First, we compute the correlations of influence factors with language-specific event relevance scores of the events covered in the Wikipedia clickstream of all considered languages (i.e., *event_relevance*, as defined in Section 4). As influence factors we select the event location closeness (*Location*), the number of links pointing to the respective event (*Links*), and the event recency (*Recency*). Results are shown in Table 3.

Table 3. Correlations of influence factors with event relevance scores in EventKG+Click.

		Influence Factors						Recency
		Location			Links			
		EN	DE	RU	EN	DE	RU	
Language-specific relevance	EN	0.4	-0.13	-0.25	0	-0.02	0.01	-0.19
	DE	-0.18	0.20	0.02	-0.01	0	0.01	0.12
	RU	-0.19	-0.08	0.26	-0.01	-0.02	0.03	0.07

The *Location* influence factor for events indicates the largest positive correlation, which confirms the existence of different language viewpoints. This effect can be most notably observed in the case of English, which has a correlation of 0.4 between the event relevance score and the *Location* closeness influence factor. The other two influence factors, namely *Links* and *Recency*, do not show any notable correlation. We assume that this is because the users are interested in both, recent and historical events, whereas recent events might not be well interlinked in Wikipedia yet.

Until now, we have considered the language-specific event relevance scores, i.e., scores assigned to each event in isolation. Now, we investigate the user click behaviour from the perspective of the event relations (i.e., *relation_relevance*, as defined in Section 4). In particular, we focus on the properties of the target event, as the language-specific relation relevance score is independent of the source entity’s relevance.

The following influence factors are used in this correlation analysis:

- i *Location*: The location closeness of the target event.
- ii *Links*: The number of links to the target event in Wikipedia.
- iii *Recency*: The recency of the target event.
- iv *Co-Mentions*: The number of co-mentions of the relation source and target in Wikipedia.

The correlation results are shown in Table 4. The correlation coefficient for the language-specific relation relevance confirms our observations concerning the

Table 4. The Pearson correlation coefficient of the relation relevance score with selected influence factors in EventKG+Click.

		Influence Factors									
		Location			Links			Recency	Co-mentions		
		EN	DE	RU	EN	DE	RU		EN	DE	RU
Language-specific relevance	EN	0.41	-0.09	-0.27	0	0	-0.01	-0.16	0.11	-0.06	-0.04
	DE	-0.13	0.10	0.01	0	0	-0.01	0.1	-0.01	0	0.01
	RU	-0.15	-0.05	0.18	0	0	0.02	0.06	-0.07	-0.07	0.13

language-specific event relevance. The closeness of the target event location has the largest influence on language-specific relevance. The links, recency and co-mentions do not correlate with the relevance scores in any of the three languages. That means, if the user reads a particular Wikipedia article, there is a higher chance that the next click leads to a spatially close event than to an event that is mentioned many times together with the source entity.

6 Conclusion and Outlook

In this paper, we presented the EventKG+Click dataset and suggested scores for capturing language-specific relevance scores for events and their relations. EventKG+Click builds upon the EventKG knowledge graph and language-specific traces of user interaction with events derived from the Wikipedia clickstream. The resulting EventKG+Click dataset contains click counts and relevance scores for more than 4 thousand events and more than 10 thousand (source, target) pairs in English, German, and Russian. Furthermore, we analysed several influence factors of language-specific relevance. We believe that the EventKG+Click dataset is a valuable resource to evaluate event relevance in language-specific contexts. In future work, we plan to develop novel user interaction models supporting cross-lingual event-centric analytics, where we will adopt the EventKG+Click dataset for training and evaluation.

Acknowledgements This work was partially funded by H2020-MSCA-ITN-2018-812997 under “Cleopatra”.

References

1. Bhatia, S., Vishwakarma, H.: Know thy Neighbors, and More! Studying the Role of Context in Entity Recommendation. In: Proceedings of the 29th on Hypertext and Social Media. pp. 87–95 (2018)
2. Dimitrov, D., Lemmerich, F., Flöck, F., Strohmaier, M.: Query for Architecture, Click through Military: Comparing the Roles of Search and Navigation on Wikipedia. In: Proceedings of the 10th ACM Conference on Web Science. pp. 371–380 (2018)

3. El Ali, A., Stratmann, T.C., Park, S., Schöning, J., Heuten, W., Boll, S.C.: Measuring, Understanding, and Classifying News Media Sympathy on Twitter after Crisis Events. In: Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems. pp. 1–13 (2018)
4. Gottschalk, S., Bernacchi, V., Rogers, R., Demidova, E.: Towards Better Understanding Researcher Strategies in Cross-lingual Event Analytics. In: Proceedings of the International Conference on Theory and Practice of Digital Libraries. pp. 139–151. Springer (2018)
5. Gottschalk, S., Demidova, E.: EventKG+TL: Creating Cross-lingual Timelines from an Event-centric Knowledge Graph. In: Proceedings of the Extended Semantic Web Conference. pp. 164–169. Springer (2018)
6. Gottschalk, S., Demidova, E.: EventKG—the Hub of Event Knowledge on the Web—and Biographical Timeline Generation. *Semantic Web* **10**(6), 1039–1070 (2019)
7. Hecht, B., Gergle, D.: The Tower of Babel meets Web 2.0: User-generated Content and its Applications in a Multilingual Context. In: Proceedings of the SIGCHI conference on human factors in computing systems. pp. 291–300 (2010)
8. Kaffee, L.A., Endris, K.M., Simperl, E., Vidal, M.E.: Ranking Knowledge Graphs By Capturing Knowledge about Languages and Labels. In: Proceedings of the 10th International Conference on Knowledge Capture. pp. 21–28 (2019)
9. Kaltenbrunner, A., Laniado, D.: There is no Deadline: Time Evolution of Wikipedia Discussions. In: Proceedings of the Eighth Annual International Symposium on Wikis and Open Collaboration. pp. 1–10 (2012)
10. Marie, N., Gandon, F., Giboin, A., Palagi, É.: Exploratory Search on Topics through Different Perspectives with DBpedia. In: Proceedings of the 10th International Conference on Semantic Systems. pp. 45–52 (2014)
11. Miz, V., Hanna, J., Aspert, N., Ricaud, B., Vandergheynst, P.: What is Trending on Wikipedia? Capturing Trends and Language Biases Across Wikipedia Editions. In: Companion Proceedings of the Web Conference 2020. ACM (2020)
12. Mocanu, D., Baronchelli, A., Perra, N., Gonçalves, B., Zhang, Q., Vespignani, A.: The Twitter of Babel: Mapping World Languages through Microblogging Platforms. *PloS one* **8**(4) (2013)
13. Nguyen, T., Tran, T., Nejd, W.: A Trio Neural Model for Dynamic Entity Relatedness Ranking. In: Proceedings of the 22nd Conference on Computational Natural Language Learning. pp. 31–41 (2018)
14. Oeberst, A., Cress, U., Back, M., Nestler, S.: Individual Versus Collaborative Information Processing: The Case of Biases in Wikipedia. In: Mass collaboration and education, pp. 165–185. Springer (2016)
15. Rupnik, J., Muhic, A., Leban, G., Skraba, P., Fortuna, B., Grobelnik, M.: News across Languages-cross-lingual Document Similarity and Event Tracking. *Journal of Artificial Intelligence Research* **55**, 283–316 (2016)
16. Tran, N.K., Tran, T., Niederée, C.: Beyond time: Dynamic Context-aware Entity Recommendation. In: European Semantic Web Conference. pp. 353–368. Springer (2017)