

# Multilingual Entity, Relation, Event and Human Value Extraction

Manling Li<sup>1</sup>, Ying Lin<sup>1</sup>, Joseph Hoover<sup>3</sup>, Spencer Whitehead<sup>1</sup>,  
Clare R. Voss<sup>2</sup>, Morteza Dehghani<sup>3</sup>, Heng Ji<sup>1</sup>

<sup>1</sup> Rensselaer Polytechnic Institute, Troy, NY, USA

{lim22, liny9, whites5, jih}@rpi.edu

<sup>2</sup> US Army Research Laboratory, Adelphi, MD, USA

clare.r.voss.civ@mail.mil

<sup>3</sup> University of Southern California, Los Angeles, CA, USA

{jehoover, mdehghan}@usc.edu

## Abstract

This paper demonstrates a state-of-the-art end-to-end multilingual (English, Russian, and Ukrainian) knowledge extraction system that can perform entity discovery and linking, relation extraction, event extraction, and coreference. It extracts and aggregates knowledge elements across multiple languages and documents as well as provides visualizations of the results along three dimensions: temporal (as displayed in an event timeline), spatial (as displayed in an event heatmap), and relational (as displayed in entity-relation networks). For our system to further support users' analyses of causal sequences of events in complex situations, we also integrate a wide range of human moral value measures, independently derived from region-based survey, into the event heatmap. This system is publicly available as a docker container and a live demo,<sup>1,2</sup> with a video demonstrating the system<sup>3</sup>.

## 1 Introduction

Knowledge extraction aims to convert unstructured texts into structured entities, relations and events. Recently, we have developed a state-of-the-art multilingual knowledge extraction system for three languages including English, Russian, and Ukrainian (Zhang et al., 2018). However, individual extraction components lack the ability to aggregate knowledge from multiple languages and documents. For example, complementary salient information about the *Ukraine crisis* may be extracted from English, Ukrainian, and Russian news documents. We develop a novel framework, as illustrated in Figure 1, to aggregate knowledge elements from multiple documents in multiple languages and visualize these knowledge elements

<sup>1</sup>System: [http://nlp.cs.rpi.edu/demo/aida\\_pipeline-master.zip](http://nlp.cs.rpi.edu/demo/aida_pipeline-master.zip)

<sup>2</sup>Demo: <http://nlp.cs.rpi.edu/software/>

<sup>3</sup>Video: <https://youtu.be/cQPHaxGLn8k>

in three interfaces (temporal, spatial, and entity-relation networks) which support effective multi-dimensional search and filtering. The system is publicly available as a series of docker containers and it can be easily run via a single script. We also provide a live demo of the system that efficiently extracts knowledge elements from user input text.

The system improves the ease and speed with which users may discover inter-connections among knowledge elements from multiple languages and documents, so users can isolate subsets of activity that warrant further attention. The complementary dimensions of the three visualization interfaces provide distinct yet comprehensive views of the entities, relations, and events as well as, most notably, their implicit connections.

For example, in the *Ukraine crisis*, a *Transport-Person* event in an *airport* in *Kramatorsk* is part of the *Attack* event in *Sloviansk*. A causal relation between these two events may be discovered both in the event heat-map interface, where the former event in *Kramatorsk* is located near the latter event in *Sloviansk*, and in the event timeline interface, where these two events both occur in *April 2014*. Furthermore, the entity-relation network interface enables users to retrieve and relate entities of interest while reasoning about such events. The interface displays each retrieved entity with its one-hop relations to other entities, which then allows the user to retrieve one-hop relations for any of those entities, thereby traversing the network and discovering information. We see this in traversing the network following the *Leadership* relation from *Donbass People's militia* to *Pro-Russian separatists* and then the *Sponsorship* relation from *Pro-Russian separatists* to *Russia*, suggesting the *Donbass People's militia* is sponsored by *Russia*.

Other types of implicit knowledge that are not readily discovered by traditional methods of knowledge extraction, such as human values, play

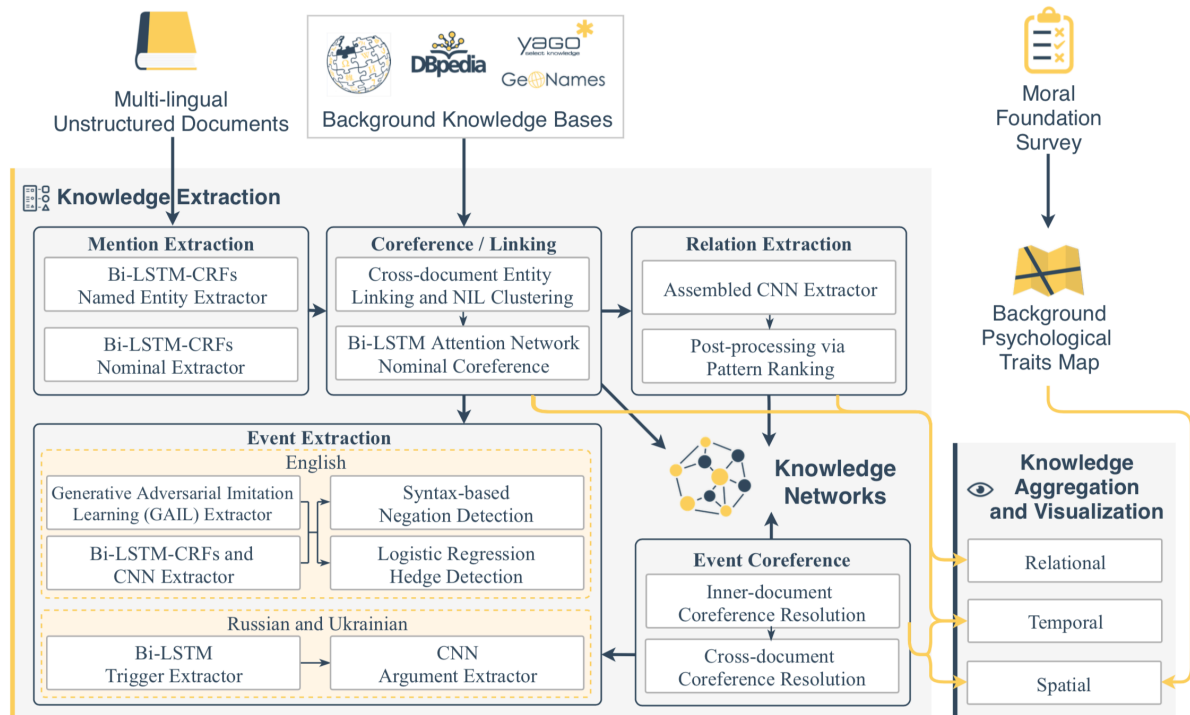


Figure 1: System overview.

a major role in social functioning and motivation (Rai and Fiske, 2011; Haidt, 2012; Graham et al., 2013; Schwartz, 2017). Numerous studies suggest that human values are often central motivating factors for protests, conflicts, and violence (Ginges and Atran, 2009; Fiske et al., 2014; Mooijman and Van Dijk, 2015; Skitka et al., 2017). Therefore, we integrate region-specific estimates of dominant psychological characteristics into the spatial event heat-map, which provides an additional layer of information that can be used to understand geo-spatial event patterns.

## 2 Multilingual Knowledge Extraction

The overall architecture of our multilingual knowledge extraction system is illustrated in Figure 1. The system performs entity discovery and linking (Pan et al., 2017; Lin et al., 2018), time expression extraction and normalization (Manning et al., 2014), relation extraction (Shi et al., 2018), event extraction (Zhang et al., 2017, 2019), and event coreference (Zhang et al., 2015). The system supports the extraction of 7 entity types, 23 relations, and 47 event types, as defined in the DARPA AIDA ontology.<sup>4</sup> Table 1 shows the main types.

For Russian and Ukrainian text input, we did

<sup>4</sup><https://www.darpa.mil/program/active-interpretation-of-disparate-alternatives>

<b>Entity</b>	Person, Organization, Geopolitical Entity, Facility, Location, Weapon, Vehicle
<b>Relation</b>	Physical, Part-Whole, Personal-Social, Measurement, Organization-Affiliation, General-Affiliation
<b>Event</b>	Life, Movement, Business, Conflict, Contact, Manufacture, Personnel, Justice, Transaction, Government, Inspection, Existence

Table 1: Main types of knowledge elements

not adopt the alternative approach of translating the source documents into English and then applying English knowledge extraction system due to the low-quality of state-of-the-art machine translation and word alignment for these two languages.

Once within-document knowledge elements for each language are extracted, the system performs cross-lingual entity linking to Wikipedia, cross-document entity clustering for unlinkable mentions, and cross-document event coreference resolution for cross-lingual information fusion. Further details of each component are described in (Zhang et al., 2018). Currently, each main component in the system outperforms the best reported results in the literature, as shown in Table 2.

Components	Ours	State-of-the-art
Name Tagging	91.8%	91.4% (Liu et al., 2018)
Relation Extraction	66.4%	65.2% (Fu et al., 2017)
Event Trigger Labeling	72.9%	69.6% (Sha et al., 2018)
Event Argument Labeling	59.0%	57.2% (Sha et al., 2018)

Table 2: F1 score comparisons of our approach vs. state-of-the-art for English knowledge extraction.

### 3 Knowledge Aggregation and Visualization

To demonstrate the capabilities of our aforementioned system, we process 10,984 documents about the Ukraine-Russia conflict scenario from the DARPA AIDA program, including 7,415 in English, 2,307 in Russian, and 929 in Ukrainian.

We organize the extracted events in our interfaces, as described below, along the temporal and spatial dimensions in order to assist users both in gaining a comprehensive view of the evolving situations in this scenario and in detecting shared patterns of occurrence and possible connections among events of interest over time and space.

#### 3.1 Event Timeline

We extract and normalize time arguments to construct an event timeline in Figure 2 using TimelineJS for visualization.<sup>5</sup> There are three zones in the web-enabled timeline interface. By clicking on an event in the timeline (*i.e.*, the gray area at the bottom of the screen), the pertinent context sentence for that event is displayed in the middle of the screen with the trigger and arguments highlighted in color, along with a link to the sentence’s source document (Figure 3). Clicking on the source document link retrieves the document with full inline annotations and its publication date, to support inference of the absolute date(s) from relative time expressions in the text (*e.g.*, “two days ago”). Additionally, at the top of the interface, users may search and filter with multiple criteria (*entity name, event type, event subtype, argument role, and time period*) to narrow down the results to a particular query of interest.

<sup>5</sup><https://timeline.knightlab.com/>

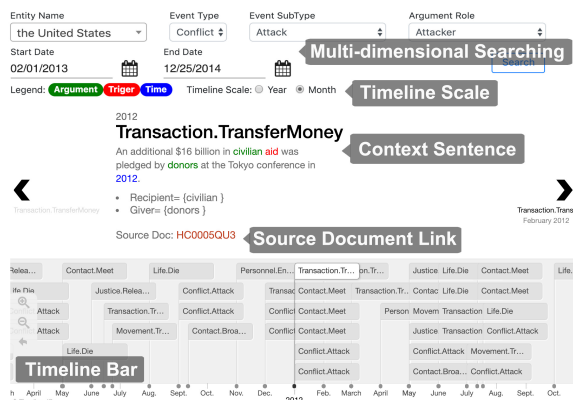


Figure 2: Example of the event timeline interface.

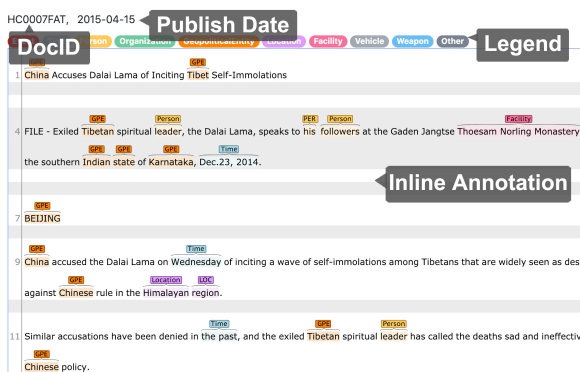


Figure 3: Example source document with inline annotation retrieved from the link in the event timeline interface.

#### 3.2 Event and Human Value Heatmap

We link event locations to the GeoNames database (Vatant and Wick, 2012) via the entity linking component and visualize involved events on a world map using Mapbox for visualization, as Figure 4 illustrates.<sup>6</sup> Each event is displayed as a dot or, when zooming in, an icon on the map. The color of a dot indicates the language of the source sentence, while the icon denotes the event type. Users can apply filters to the map to view the events of a certain type or language.

In addition to events, we also integrate regional estimates of human values into the heatmap. Specifically, the system encodes the geographic variations of 10 distinct dimensions of the human values in Table 3. These values are proposed in the Schwartz Basic Theory of Human Values (Schwartz, 2012) as a culturally universal taxonomy of human values.

The human values estimates are derived from the European Social Survey (ESS) (Round, 5, 6,

<sup>6</sup><https://www.mapbox.com/>

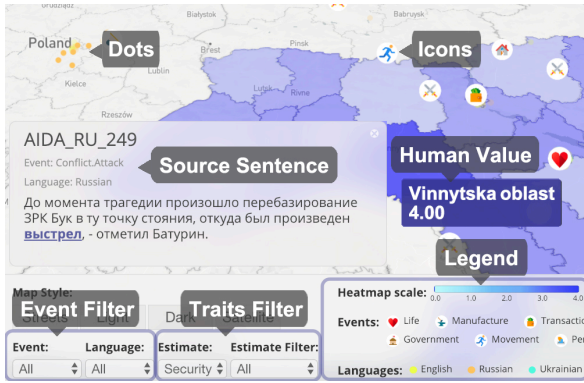


Figure 4: Example event heatmap with events and human values by region.

<b>Human Values</b>	Achievement, Benevolence, Conformity, Hedonism, Power, Security, Self-direction, Simulation, Tradition, Universalism
<b>Age Filter</b>	15-29, 30-44, 45-59, 60+
<b>Gender Filter</b>	Female, Male

Table 3: Human values. In the heatmap, the estimates for these values are displayed by region.

7), a nationally representative survey administered throughout the European Union. While the ESS data is sufficient for directly estimating *national* human values, it cannot be used to directly derive Oblast-level estimates because it is not representative at the Oblast-level.<sup>7</sup> To resolve this issue, we employ a state-of-the-art approach to survey adjustment and small-area estimation called Multi-level Regression and Synthetic Post-stratification with Spatial Smoothing (MrsP-SM) (Park et al., 2004; Selb and Munzert, 2011; Leemann and Wasserfallen, 2017; Hoover and Dehghani, 2018). This involves a model-based approach to post-stratification in which a hierarchical regression model is used to model person-level responses to a survey item as a function of demographic characteristics, region-level factors, and geographic indicators. Then, the model is used to generate predictions for each combination of demographic variables and geographic region. Finally, the predictions are weighted by the demographic population proportions within each region, yielding a set of regularized regional estimates that are adjusted for representativeness. To obtain regional human values estimates in the event heatmap, we estimate

<sup>7</sup>Our regional unit of analysis is the Oblast, of which there are 24 in Ukraine.

MrsP-SM models for each of the 10 Schwartz Human Values domains.

Human values have close ties to the intentions underlying events. A *Demonstration* event may result in violence, property destruction and involvement of extremist groups. The values of *Benevolence*, *Hedonism*, and *Conformity* among authority figures may impact their response to a protest. Additionally, people in areas where *Conflict* events are common may have higher values for *Security* and lower values for *Achievement*.

### 3.3 Entity-relation Networks

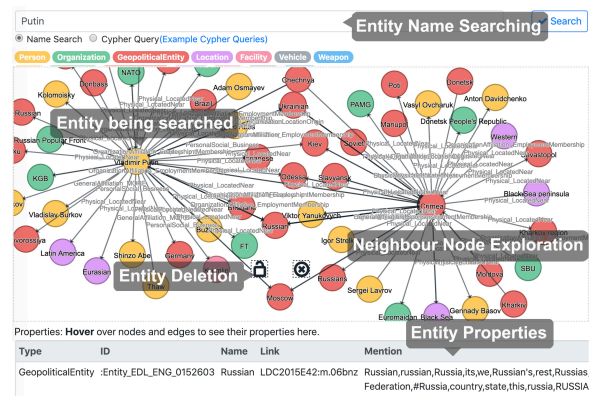


Figure 5: Entity-relation network.

A critical task for users gaining an understanding of complex scenarios is to explore implicit entity relations beyond the scope of traditional in-line document annotation. Our interface provides interactive knowledge graph exploration, using Neo4j<sup>8</sup> (Figure 5), where entities can be searched by name and a sub-graph for each entity with its one-hop neighbors and their relations is returned, with entity properties displayed at the bottom of the interface. Users may either explore each retrieved neighbour by double clicking on it for its subgraph, or reduce their search graph by deleting entities no longer of interest. Thus, users can construct a multi-hop entity-relation graph, discovering variable length paths between entities. Each entity is labeled with its canonical name mention, while the entities without name mentions are removed from the network.

## 4 Conclusions and Future Work

In this paper, we demonstrate a comprehensive multi-lingual knowledge extraction, aggregation

<sup>8</sup><https://neo4j.com/>

and visualization system which can effectively discover and synthesize knowledge elements from multiple data sources, and present them to users in multiple dimensions. In the future, we plan to conduct utility experiments with users to compare and evaluate the quality and speed of generating summary reports with and without using our interfaces.

## Acknowledgments

This work was supported by the U.S. ARL NS-CTA No. W911NF-09-2-0053 and DARPA AIDA Program No. FA8750-18-2-0014. The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the U.S. Government. The U.S. Government is authorized to reproduce and distribute reprints for Government purposes notwithstanding any copyright notation here on.

## References

- A P Fiske, T S Rai, and S Pinker. 2014. *Virtuous Violence: Hurting and Killing to Create, Sustain, End, and Honor Social Relationships*. Cambridge University Press.
- Lisheng Fu, Thien Huu Nguyen, Bonan Min, and Ralph Grishman. 2017. Domain adaptation for relation extraction with domain adversarial neural network. In *Proceedings of the Eighth International Joint Conference on Natural Language Processing (Volume 2: Short Papers)*, volume 2, pages 425–429.
- Jeremy Ginges and Scott Atran. 2009. What motivates participation in violent political action. *Annals of the New York Academy of Sciences*, 1167(1):115–123.
- Jesse Graham, Jonathan Haidt, Sena Koleva, Matt Motyl, Ravi Iyer, Sean P Wojcik, and Peter H Ditto. 2013. Moral foundations theory: The pragmatic validity of moral pluralism. In *Advances in experimental social psychology*, volume 47, pages 55–130. Elsevier.
- Jonathan Haidt. 2012. *The righteous mind: Why good people are divided by politics and religion*. Vintage.
- J Hoover and M Dehghani. 2018. The big, the bad, and the ugly: Geographic estimation with flawed psychological data. *PsyArXiv*. October.
- Lucas Leemann and Fabio Wasserfallen. 2017. Extending the use and prediction precision of subnational public opinion estimation: EXTENDING USE AND PRECISION OF MrP. *American journal of political science*, 61(4):1003–1022.
- Ying Lin, Shengqi Yang, Veselin Stoyanov, and Heng Ji. 2018. A multi-lingual multi-task architecture for low-resource sequence labeling. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, volume 1, pages 799–809.
- Liyuan Liu, Jingbo Shang, Xiang Ren, Frank Fangzheng Xu, Huan Gui, Jian Peng, and Jiawei Han. 2018. Empower sequence labeling with task-aware neural language model. In *Thirty-Second AAAI Conference on Artificial Intelligence*.
- Christopher Manning, Mihai Surdeanu, John Bauer, Jenny Finkel, Steven Bethard, and David McClosky. 2014. The stanford corenlp natural language processing toolkit. In *Proceedings of 52nd annual meeting of the association for computational linguistics: system demonstrations*, pages 55–60.
- Marlon Mooijman and Wilco W Van Dijk. 2015. The self in moral judgement: How self-affirmation affects the moral condemnation of harmless sexual taboo violations. *Cognition and Emotion*, 29(7):1326–1334.
- Xiaoman Pan, Boliang Zhang, Jonathan May, Joel Nothman, Kevin Knight, and Heng Ji. 2017. Cross-lingual name tagging and linking for 282 languages. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, volume 1, pages 1946–1958.
- David K Park, Andrew Gelman, and Joseph Bafumi. 2004. Bayesian multilevel estimation with post-stratification: State-Level estimates from national polls. *Political analysis: an annual publication of the Methodology Section of the American Political Science Association*, 12(4):375–385.
- Tage Shakti Rai and Alan Page Fiske. 2011. Moral psychology is relationship regulation: moral motives for unity, hierarchy, equality, and proportionality. *Psychological review*, 118(1):57–75.
- ESS Round. 5. 5: European social survey round 4 data (2010). *Data file edition*, 5.
- ESS Round. 6. 6: European social survey round 6 data (2012). *Data file edition*, 6.
- ESS Round. 7. 7: European social survey round 7 data (2014). *Data file edition*, 7.
- Shalom H Schwartz. 2012. An overview of the schwartz theory of basic values. *Online Readings in Psychology and Culture*, 2(1):11.
- Shalom H Schwartz. 2017. The refined theory of basic values. In Sonia Roccas and Lilach Sagiv, editors, *Values and Behavior: Taking a Cross Cultural Perspective*, pages 51–72. Springer International Publishing, Cham.

- P Selb and S Munzert. 2011. Estimating constituency preferences from sparse survey data using auxiliary geographic information. *Political analysis: an annual publication of the Methodology Section of the American Political Science Association*.
- Lei Sha, Feng Qian, Baobao Chang, and Zhifang Sui. 2018. Jointly extracting event triggers and arguments by dependency-bridge rnn and tensor-based argument interaction. In *Thirty-Second AAAI Conference on Artificial Intelligence*.
- Ge Shi, Chong Feng, Lifu Huang, Boliang Zhang, Heng Ji, Lejian Liao, and Heyan Huang. 2018. Genre separation network with adversarial training for cross-genre relation extraction. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 1018–1023.
- Linda J Skitka, Brittany E Hanson, and Daniel C Wisneski. 2017. Utopian hopes or dystopian fears? exploring the motivational underpinnings of moralized political engagement. *Personality & social psychology bulletin*, 43(2):177–190.
- Bernard Vatant and Marc Wick. 2012. Geonames ontology. *Dostupné online*; [http://www.geonames.org/ontology/ontology\\_v3](http://www.geonames.org/ontology/ontology_v3), 1.
- Boliang Zhang, Di Lu, Xiaoman Pan, Ying Lin, Halidanmu Abudukelimu, Heng Ji, and Kevin Knight. 2017. Embracing non-traditional linguistic resources for low-resource language name tagging. In *Proceedings of the Eighth International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, volume 1, pages 362–372.
- Tongtao Zhang, Heng Ji, and Avirup Sil. 2019. Joint entity and event extraction with generative adversarial imitation learning. *Data Intelligence*.
- Tongtao Zhang, Hongzhi Li, Heng Ji, and Shih-Fu Chang. 2015. Cross-document event coreference resolution based on cross-media features. In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, pages 201–206.
- Tongtao Zhang, Ananya Subburathinam, Ge Shi, Lifu Huang, Di Lu, Xiaoman Pan, Manling Li, Boliang Zhang, Qingyun Wang, Spencer Whitehead, Heng Ji, Alireza Zareian, Hassan Akbari, Brian Chen, Ruiqi Zhong, Steven Shao, Emily Allaway, Shih-Fu Chang, Kathleen McKeown, Dongyu Li, Xin Huang, Kexuan Sun, Xujun Peng, Ryan Gabbard, Marjorie Freedman, Mayank Kejriwal, Ram Nevatia, Pedro Szekely, T.K. Satish Kumar, Ali Sadeghian, Giacomo Bergami, Sourav Dutta, Miguel Rodriguez, and Daisy Zhe Wang. 2018. [Gaia - a multi-media multi-lingual knowledge extraction and hypothesis generation system](#). In *Proceedings of TAC KBP 2018, the 25th International Conference on Computational Linguistics: Technical Papers*.