

# CEVO: Comprehensive Event Ontology Enhancing Cognitive Annotation on Relations

Saeedeh Shekarpour\*, Faisal Alshargi†, Krishnaprasad Thirunarayan‡, Valerie L. Shalin‡, Amit Sheth‡

\* University of Dayton, USA, sshekarpour1@udayton.edu

† University of Leipzig, Germany, alshargi@informatik.uni-leipzig.de

‡ Kno.e.sis Center, {tkprasad, valerie, amit}@knoesis.org

**Abstract**—While the general analysis of named entities has received substantial research attention on unstructured as well as structured data, the analysis of relations among named entities has received limited focus. In fact, a review of the literature revealed a deficiency in research on the abstract conceptualization required to organize relations. We believe that such an abstract conceptualization can benefit various communities and applications such as natural language processing, information extraction, machine learning, and ontology engineering. In this paper, we present Comprehensive Event Ontology (CEVO), built on Levin’s conceptual hierarchy of English verbs that categorizes verbs with shared meaning, and syntactic behavior. We present the fundamental concepts and requirements for this ontology. Furthermore, we present three use cases employing the CEVO ontology on annotation tasks: (i) annotating relations in plain text, (ii) annotating ontological properties, and (iii) linking textual relations to ontological properties. These use-cases demonstrate the benefits of using CEVO for annotation: (i) annotating English verbs from an abstract conceptualization, (ii) playing the role of an upper ontology for organizing ontological properties, and (iii) facilitating the annotation of text relations using any underlying vocabulary. This resource is available at <https://shekarpour.github.io/cevo.io/> using <https://w3id.org/cevo> namespace.

**Index Terms**—cevo, relation annotation, event ontology, abstract ontology, cognitive annotation.

## I. INTRODUCTION

While the size of data on the Web is growing dramatically in both structured and unstructured, still a high proportion of the Web remains unstructured (i.e., textual data) such as social network feeds, blogs, news, logs. However, the size of structured data is also significant; e.g., so far more than 130 billion triples are published on Linked Open Data from over 9,960 datasets<sup>1</sup>. These large datasets subscribe to ad hoc, independently created and diverse lexicons, nomenclatures, vocabularies, taxonomies, and ontologies. Although this diversity provides flexibility, it complicates both reuse of ontologies and interlinking of datasets.

So far, named entity recognition and entity linking to the background knowledge base have received substantial attention in research studies. However, the analysis of relations over named entities has not. Although cognitive scientists e.g., Dumas and Humel, 2005 [4], [5] suggest that relational content is key to reasoning, a review of the literature on

unstructured as well as structured data revealed a deficiency in research on the abstract conceptualization required to organize relations. Observing deficiencies in (i) *relation extraction* from text, (ii) *contextual equivalencing of relations*, and (iii) dealing with the *diversity of ontologies* motivated us to investigate an abstract conceptualization of relations.

We found an answer to our struggle in the organized lexicon and knowledge base assembled by the Stanford linguist Beth Levin in [9]. While Levin relies on Schank’s conceptual dependency theory [12], [13] to organize this knowledge base, the key here is the psychologically principled inventory of English verbs aligned with the knowledge base. Classes in this knowledge base identify sets of semantically coherent verbs with corresponding syntactic properties. For example, the *communication* class refers to verbs causing a communication activity. A subcategory of this class is the category of verbs transferring a message/idea (i.e. shared meaning), such as *announce*, *say*, and *mention*. In addition, they share the same syntactic behaviour such as NP1 VP NP2 or NP1 VP NP2 to NP3 (NP is noun phrase and VP is verb phrase). For example, Jack says greetings to Sarah is following the syntactic pattern NP1 VP NP2 to NP3. The semantic behavior is *communication* activity between two entities (i.e., Jack and Sarah) with the transferred message “greetings”. This conceptualization provides more than 230 classes for over 3,000 English verbs. Using Levin’s work, we build an event ontology and lexicon called Comprehensive Event Ontology (CEVO)<sup>2</sup>. CEVO is designed to recognize and equate relations from both textual data sources as well as knowledge bases. Such an abstract conceptualization benefits many applications, such as natural language processing, information extraction, ontology engineering, and machine learning.

Figure 1 represents the evolution of ontologies as well as vocabularies regarding the *level of abstraction*. The early generation of vocabularies was primarily for annotating datasets (i.e., metadata) or describing a domain. The next generation of vocabularies have supported interoperability requirements and hence are created from a higher level of abstraction. The subsequent generation of ontologies has cognitive applicability and therefore has the highest level of abstraction. CEVO provides a high abstraction for annotating events.

This paper is organized as follows. Related work is presented in Section II. Section III presents the deficiencies

<sup>1</sup>observed on September 30th 2018 at <http://lodstats.aksw.org/>

<sup>2</sup>CEVO namespace: <https://w3id.org/cevo/>

that motivated us to develop the CEVO ontology. Section IV discusses fundamental requirements for designing the CEVO ontology. The principles and considerations that Beth Levin has taken into account for categorizing English verbs are presented in Section V. The main concepts of CEVO are introduced in Section VI. We discuss three use cases employing CEVO for annotation tasks in Section VII. We close with the conclusion and future work in Section VIII.

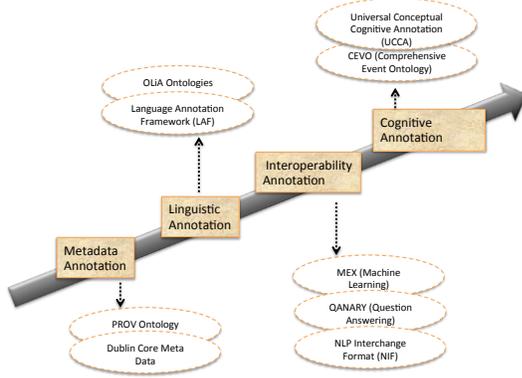


Fig. 1: Evolution of Abstraction Level of Vocabularies and Ontologies.

## II. RELATED WORK

We first review the definition of *event* which has been given previously in different ontologies and then present several existing ontologies that facilitate the annotation task and interoperability among various components.

*a) Event Classification: LODÉ:* An ontology for Linking Open Descriptions of Events<sup>3</sup> defines a single generic concept of an event as ‘Something that happened’, e.g., reported in a news article or with historical significance. This is a generic definition and does not specify the various types of events necessary for subsequent inference. **Schema.org**<sup>4</sup> introduces a similar generic concept of an event<sup>5</sup> that additionally considers temporal as well as location aspects and additionally provides a limited hierarchy. This hierarchy introduces types of events such as business events, sale events, and social events. Similarly, the **DBpedia ontology**<sup>6</sup> defines the generic concept of an event with a hierarchy which is broader, including lifecycle events (e.g. birth, death), natural events (e.g. earthquake, stormsurge), and societal events (e.g. concert, election). However, the state-of-the-art does not provide a detailed specification about various types of events. Having such specification is of high importance because in specific domains such as news domain it is required to differentiate between various types of events. *Event classification* task is necessary in case the employed background data model considers the specific type of events as part of an event annotation. In this case, event phrases have to be labeled by specific types of events using multi-class classifier trained

for distinguishing the specific type of a given event. For example, Table I shows samples of news headline annotated based on their major event (column 2) like the tweets no.2, no.5, no.8 of Table I having the specific type “*meet*”. To the best of our knowledge, CEVO is the first event ontology that provides a fine-grained abstract conceptualization of events. In fact, CEVO captures the whole of the hierarchy of Levin’s classification as the hierarchy representing specific events. Thus, it specifies 230 events.

Event Type	News Headlines Tweets
communication	no1. Michelle Obama tells #SXSWS crowd: I will not run for president
meet	no2. Instagram CEO meets with @Pontifex to discuss "the power of images to unite people"
murder (killing)	no3. Chemical accident in Bangkok bank kills eight people
communication	no4. State elections were "difficult day," German Chancellor Angela Merkel says
meet	no5. Pope Francis visits Cuba and Mexico
murder (killing)	no6. Storms kill at least three in Virginia
communication	no7. Obama and Justin Trudeau announce efforts to fight climate change
meet	no8. Pope to meet leader of Russian Orthodox Church for first time in nearly
murder (killing)	no9. 2 air force pilots from United Arab Emirates killed when warplane crashed over Yemen

TABLE I: Samples of news headlines from different publishers on Twitter.

*b) Interoperability Challenge:* the interoperability among heterogeneous datasets, schemas, various tools, and dependent components is an important issue. So, recently, a number of linguistic ontologies as well as annotation tools for related purposes have emerged. The Ontologies of Linguistic Annotation (OLiA) [3], [2] provides annotation tag sets using syntactical and morphological perspectives. OLiA covers over 110 OWL ontologies for over 34 tag sets in 69 different languages. Thus, all NLP tools can leverage OLiA’s tag sets for annotating the output. Among the ontologies introduced for promoting interoperability between tools, services, and components, we mention two recent ones:

- 1) NLP Interchange Format (NIF)<sup>7</sup> [8] provides a vocabulary for interoperability between Natural Language Tools (NLP). NIF allows tools to exchange annotations for any part of text using three layers: (i) *Structural layer*: This describes URI scheme for identifying any part of a document; thus, making each piece of text dereferenceable. (ii) *Conceptual layer*: The NIF Core Ontology<sup>8</sup> describes classes and properties for providing relations between tokens and documents. The core class is `nif:String`, which points to any word mentioned in Unicode characters. (iii) *Access layer*: NIF-aware applications publish their output using the NIF format via REST APIs.
- 2) qa vocabulary<sup>9</sup> [14], [1] is intended to facilitate interoperability among components of question answering systems. Currently, it defines a two-layered ontology: a *abstract layer* describing the generic functionalities of each component and a *binding layer* enabling binding output of each element to the abstract layer.

CEVO connects the conceptualization of events to a large lexicon of English verbs, i.e., 3000 verbs. Thus, it can provide an abstract annotation on relations. Please note that this knowledge base is different from lexicons such as *lemon*<sup>10</sup>

<sup>3</sup><http://linkedevents.org/ontology/>

<sup>4</sup><http://schema.org>

<sup>5</sup><http://schema.org/Event>

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

<sup>7</sup><http://persistence.uni-leipzig.org/nlp2rdf/>

<sup>8</sup><http://persistence.uni-leipzig.org/nlp2rdf/ontologies/nif-core/nif-core.html#>

<sup>9</sup><https://github.com/WDAqua/QAontology>

<sup>10</sup><http://lemon-model.net/learn/5mins.html>

[10], [15] which shares purely terminological and lexicological resources on the Semantic Web. Potentially, CEVO conceptualization will play a significant role in overcoming the existing challenges of (i) tagging relations, (ii) linking relations, and (iii) ontology alignment.

### III. PROBLEM STATEMENT

The CEVO ontology compensates for pervasive deficiencies in the abstract conceptualization of relations. In the following, we mention the three well-known deficiencies in annotating relations or ontological properties.

- (I) **Relation Extraction:** Decades of research in the field of information extraction has resulted in tools that successfully recognize and annotate Named Entities (NE) and link them to entities available in a background knowledge base (e.g., [7], [6], [11]). In contrast, recognizing, tagging, and linking relationships among entities has received limited attention.
- (II) **Contextual Equivalence of Relations:** Relations embedded in plain text can be expressed in various ways either explicitly or implicitly. Explicit relations often appear as a verb phrase and implicit ones are usually hidden or embedded in other phrases, e.g., an adjective phrase such as "European Countries" (countries located in Europe). Moreover, a single explicit relation can be expressed using several distinct verbs. E.g., consider the two sentences 'Jack visits Sara', and 'Jack consults Sara'. In both of these cases, the *abstract event of meeting* is taking place using two different verbs, 'visit' and 'consult'. Nevertheless, these two verbs do not have a simple synonymous relationship recoverable by inference using lexicons such as WordNet. These verbs convey the same event in a specific context.
- (III) **Diversity in conceptualization:** Each ontology is created based upon a specific interpretation of a domain. This strategy yields various ontologies used by different users or communities. While this heterogeneity provides flexibility, it complicates reusability and introduces integration challenges. Linking the mention of either entity or relation from plain text to the corresponding background knowledge base is ontology-specific.

### IV. REQUIREMENTS FOR CEVO

On the basis of these observations, we derive a series of requirements for a cognitive ontology for annotating relations on both structured as well as unstructured data. The coarse-grained requirements are listed below.

#### *Requirement 1 (Relation Tagging on Textual Data):*

Similar to tagging Named Entities in plain text, each mention of a relation must be recognized, normalized, and tagged. Thus, a tags set is required for distinguishing relations.

*Requirement 2 (Relation Linking):* Beyond recognizing and tagging mentions of relations in plain text, it is necessary to link textual relations to ontological properties. To do that, having an upper ontology that can be used for annotating both ontological properties and textual relations is required.

#### *Requirement 3 (Integration and Alignment of Properties):*

The variety of ways to conceptualize a domain results in different ontologies. Certainly, overlaps between ontologies require alignment or integration. Thus, annotating ontologies based on an upper ontology which has a higher abstraction can help in integrating and aligning ontologies.

*Requirement 4 (Reusability):* One of the main obstacles for the reuse of ontologies is the additional effort required for interpreting the conceptualization represented. Providing annotations based on a cognitive conceptualization which is from a higher abstraction level indeed boosts reusability because it is not involved in domain-specific considerations.

*Requirement 5 (Simplicity):* Since we desire CEVO to be widely adopted and reused, the captured cognitive conceptualization has to be as simple as possible to minimize integration and adoption efforts.

To the best of our knowledge, there is no existing ontology on Linked Data that fulfills the above requirements. To this end, we present CEVO (a cognitive event ontology), built upon [9]. We contend that CEVO fulfills an abstract conceptualization of relations. In the following section, we discuss principles behind this conceptual hierarchy.

### V. LEVIN CONCEPTUAL HIERARCHY

The entries of Levin's lexical knowledge base are verb classes. For example, Figure 3 illustrates two distinct English verb classes (1) *transformation and creation* and (2) *change of the state* both of which subsume several verbs. The members of each class (i.e. English verbs) have two characteristics: (i) semantic coherence and (ii) shared syntactic behavior. These characteristics are described as follows:

(i) *Semantic Coherence:* Each class of verbs has a unique set of properties that shape the meaning of the member verbs (i.e. shared meaning). In fact, the conjunction of properties provides a distinctive meaning for each class. A single meaning property might be attributed to several classes depending upon context. Moreover, an individual verb can belong to multiple classes, creating a graph instead of a tree. For example, the class of *Creation and Transformation* (shown in Figure 3) refers to a class of verbs causing alternation. This class contains both transitive verbs (an agent creates an entity) and intransitive verbs (describing transformation of an entity). A subclass of this class is the *Build* class with the shared properties: (1) material/product alternation, (2) total transformation alternation, (3) unspecified object alternation, (4) benefactive alternation, (5) causative alternation, (6) raw material subject alternation, and (7) sum of money subject alternation; whereas another subclass namely *Grow* class shares only three properties: (1) material/product alternation, (2) total transformation alternation, and (3) causative alternation.

An English verb might belong to two distinct classes. For example, the verbs 'cook' and 'boil' belong to two distinct classes: (a) creation and transformation and (b) change of state. These two classes are represented in Figure 3. Thus, depending on the *context*, the appropriate class is distinguished.

(ii) *Shared syntactic behavior:* Meaning influences the syntactic behavior of a verb in terms of its expression and the



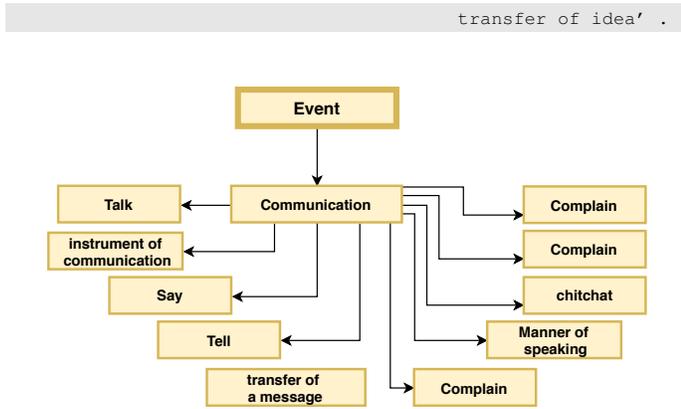


Fig. 4: Hierarchy of `cevo:Communication` event.

The next main class is `cevo:Verb` refers to words with the part-of-speech *verb*. This class is equivalent to the class *main verb* of *OLiA* ontology [3], [2]<sup>11 12</sup> which is an annotation model based on morphology.

```
cevo:Verb          a          owl:Class      ;
                  owl:equivalentClass  OLiA:MainVerb .
```

### B. CEVO Instance Level (Verb Individuals)

So far, we described the schema level classes containing two major classes `cevo:Event` and `cevo:Verb`; the next important step is to map each individual English verb to the corresponding event class. Because any **appearance of a verb** shows the likelihood of occurrence one or multiple events. In other words, verbs expose occurrences of events. Thus, each individual verb has a dual sort of `rdf:type` (i) the class `cevo:Verb` meaning this individual linguistically is a verb and (ii) `cevo:Event` meaning this verb manifests occurrence of an event. We use the prefix `cevov: <https://w3id.org/cevo/verb/>` for referring verb individuals. Thus, we instantiate each English verb at the instance level and type that primarily as `cevo:Verb` and then map it to the associated event(s). In the following, the two English verbs *say* and *cook* are instantiated. They are primarily typed as `cevo:Verb`, and furthermore, they are typed to their corresponding event classes, respectively `cevo:Communication` and `cevo:Creation and Transformation` events. In fact, by specifying the type of a verb as `cevo:Verb`, explicitly, the syntactic role of that verb in the English language is determined and by associating the relevant events to a verb, domain-specific semantic roles of that verb are determined. Please note that, each individual verb can be associated with several event classes. For instance, the verb `cevov:cook`, in addition to the event `cevo:Creation and Transformation`, is also associated with the event `cevo:Change of the State` (shown in Figure 5).

<sup>11</sup>OLiA Ontology is available at <http://nachhalt.sfb632.uni-potsdam.de/owl/>  
<sup>12</sup>Documentation of OLiA is available at <http://nachhalt.sfb632.uni-potsdam.de/owl/olia.owl>.

```
(a) cevov:say      rdf:type   cevo:MainVerb      .
(a) cevov:say      rdf:type   cevo:Communication .
(b) cevov:cook     rdf:type   cevo:MainVerb      .
(b) cevov:cook     rdf:type   cevo:Creation_Transformation .
(b) cevov:cook     rdf:type   cevo:Change_of_the_state .
(b) cevov:cook     rdf:type   cevo:Cooking      .
(b) cevov:cook     rdf:type   cevo:Build       .
```

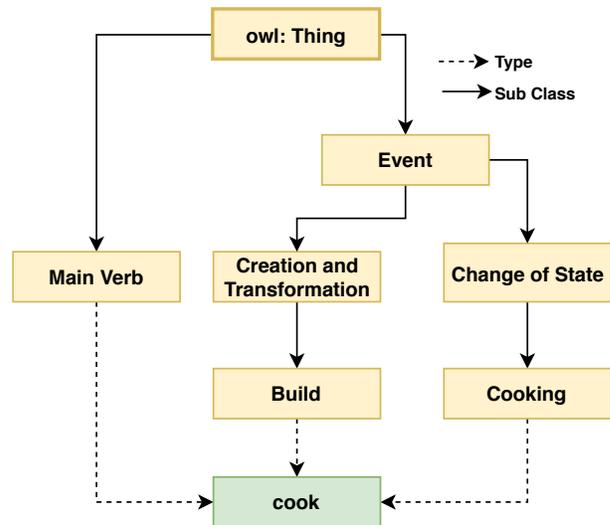


Fig. 5: Classes associated with the verb `cevov:cook`.

On the other hand, a group of verbs reveals the occurrence of an event. For example, Figure 6 represents all the verbs which may cause a *Transferring a message* event (a subclass of `cevo:Communication`).

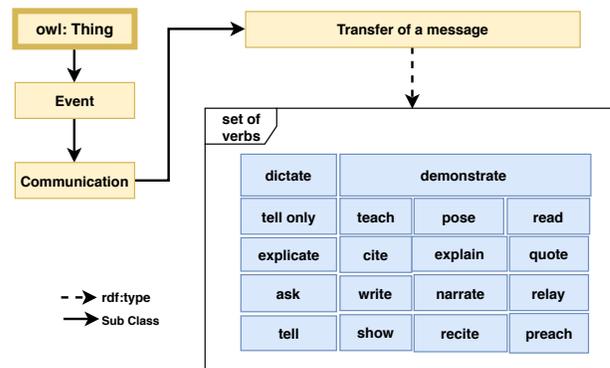


Fig. 6: Verbs associated with the event `cevo:Transfer of a message`.

### C. Querying CEVO

A typical concern is how we can obtain the list of events related to a particular verb. Let us assume that the input verb `cook` is given. The following SPARQL query retrieves the event classes associated with the input verb `cook`.

```
SELECT distinct ?v ?event
WHERE {
  ?v rdf:type cevo:Verb .
  ?v rdf:type ?event .
  ?v rdfs:label ?vlabel .
  Filter (regex(?vlabel,'cook')) .
  Filter (?event != cevo:Verb) .
  Filter (?event != owl:NamedIndividual) .}
```

## VII. USE CASES

We present three use cases employing CEVO for annotating textual relations and ontological relations.

### A. Use Case 1: Annotating Relations in Text using CEVO

Here, we show the applicability of CEVO for annotating relations in plain text. Figure 7 shows two headlines from news on Twitter. The first tweet was published by the BBC and the second one was published by the New York Times. Tweet #1 is headed by the verb *announce* and the tweet #2 is headed by the verb *says*. Both of these tweets have a similar meaning in the sense that a *message is transferred*. Annotating these two tweets via CEVO uses the same tag *communication* for both of these verbs, whereas they do not hold any lexical relation such as synonymy.

Tweet #1: BBC on 10/3/2016:

Obama and Justin Trudeau **announce** efforts to fight climate change.  
cevo:Communication

Tweet #2: NYT on 14/3/2016:

State elections were "difficult day," German Chancellor Angela Merkel **says**.  
cevo:Communication

Fig. 7: Annotating relations expressed in headlines from news published on Twitter using CEVO. Both of these tweets obtained similar tags because, in both, an event of ‘transferring a message’ appears.

For specificity, we first annotate the two strings *announce* and *says* using NIF vocabulary.

```
:tweetID#char=26,33      a          NIF:String      ;
                          NIF:beginIndex 26          ;
                          NIF:endIndex    33          ;
                          NIF:anchorOf   "announce"   ;
                          NIF:oliaCategory Olia:MainVerb .

:tweetID#char=71,74      a          NIF:String      ;
                          NIF:beginIndex 71          ;
                          NIF:endIndex    74          ;
                          NIF:anchorOf   "says"       ;
                          NIF:oliaCategory Olia:MainVerb .
```

We enhance the annotations with the relevant event type using CEVO as follows:

```
tweetID#char=26,33      a          cevo:Communication .
tweetID#char=71,74      a          cevo:Communication .
```

### B. Use Case 2: Annotating Ontological Properties

CEVO can play the role of upper ontology, thus it can be utilized for annotating properties of any ontology. One way of providing such an annotation is using the Web Annotation Data Model<sup>13</sup> (WADM), which is a framework for expressing annotations. A WADM annotation has two elements (i) a target which indicates the resource being annotated, and (ii) a body which indicates the description. Annotating properties of various ontologies with CEVO addresses integration and

alignment problems. Assume that we have the property `dbp:spouse`<sup>14</sup> from DBpedia ontology, which represents the relation of *marrying* that is semantically related to the class `cevo:Amalgamate` (please note that the verb `cevov:marry` belongs to the the class `cevo:Amalgamate`). The annotation of this property is presented in Turtle syntax using the WADM framework as follows:

```
example:annotation1    a          oa:Annotation      ;
                          oa:hasTarget  dbp:spouse       ;
                          oa:hasBody    cevo:Amalgamate .
```

### C. Use case 3: Linking Relations

CEVO facilitates linking occurrences of relations in plain text to ontological properties. For example, on the March 4th 2016, the BBC published this headline: Rupert Murdoch and Jerry Hall marry. The embedded relation in this part of text is *marry*. This relation is annotated as `cevo:Amalgamate` employing CEVO ontology. We show this annotation using the NIF vocabulary below.

```
:headline#char=31,35    a          nif:String      ;
                          nif:beginIndex 31          ;
                          nif:endIndex    35          ;
                          nif:anchorOf   'marry'     ;
                          nif:oliaCategory Olia:MainVerb .

                          a          cevo:Amalgamate .
```

`:headline#char=31,35` is the assigned URI for the verb *marry* in the headline mentioned. By taking into account the previous annotation for `dbp:spouse` property, we are now empowered to link the *marry* relation in the headline directly to the property `dbp:spouse` due to the similar tag of the CEVO event. This annotation is represented using WADM as follows:

```
:annotation3           a          oa:Annotation      ;
                          oa:hasTarget  :headline#char=31,35 ;
                          oa:hasBody    dbo:spouse       .
```

Thus, using SPARQL query, we can link a textual relation to the appropriate ontological relation based on their common CEVO annotations. For example, the verb *marry* is linked to the property `dbp:spouse` as follows:

```
:headline#char=31,35  itsrdf:taIdentRef  dbp:spouse .
```

## VIII. CONCLUSIONS AND FUTURE DIRECTIONS

In this paper, we introduced an event ontology called *CEVO* (available at <https://shekarpour.github.io/cevo.io/>). This ontology relies on an abstract conceptualization of English verbs provided by Beth Levin in [9]. Such an abstract conceptualization largely obviates deficiencies in (i) relation extraction from text, (ii) contextual equivalencing of relations, and (iii) diversity of ontologies. This ontology presents more than 230 event classes for over 3,000 English verbs as individuals. We plan to extend CEVO in the direction of integrating

<sup>13</sup>W3C Recommendation since 23 February 2017, <http://www.w3.org/TR/annotation-model>

<sup>14</sup>We use the prefix `dbo:` <<https://dbpedia.org/ontology/>> and the prefix `dbp:` <<https://dbpedia.org/property/>>.

and interlinking to other existing ontologies, especially those that contain a conceptualization of events. Currently, we are applying CEVO to the domain of news for extracting events from news headlines.

## IX. ACKNOWLEDGMENTS

We acknowledge partial support from the National Science Foundation (NSF) award: EAR 1520870: Hazards SEES: Social and Physical Sensing Enabled Decision Support for Disaster Management and Response.

## REFERENCES

- [1] Andreas Both, Jens Lehmann, Sören Auer, and Martin Brümmer. Integrating NLP using linked data. In *The Semantic Web - ISWC 2013 - 11th International Extended Semantic Web Conference, Sydney, NSW, Australia, October 21-25, 2013, Proceedings, Part II*, pages 98–113, 2013.
- [2] Christian Chiarcos. Ontologies of linguistic annotation: Survey and perspectives. In *Proceedings of the Eight International Conference on Language Resources and Evaluation (LREC'12)*, Istanbul, Turkey, may 2012. European Language Resources Association (ELRA).
- [3] Christian Chiarcos and Maria Sukhareva. Olia - ontologies of linguistic annotation. *Semantic Web*, 6(4):379–386, 2015.
- [4] LA Doumas and John E Hummel. Approaches to modeling human mental representations: What works, what doesn't and why. *The Cambridge handbook of thinking and reasoning*, ed. KJ Holyoak & RG Morrison, pages 73–94, 2005.
- [5] Leonidas AA Doumas and John E Hummel. A symbolic-connectionist model of relation discovery. In *Proceedings of the Annual Meeting of the Cognitive Science Society*, volume 27, 2005.
- [6] Paolo Ferragina and Ugo Scaiella. Fast and accurate annotation of short texts with wikipedia pages. *IEEE Software*, 29(1):70–75, 2012.
- [7] Jenny Rose Finkel, Trond Grenager, and Christopher Manning. Incorporating non-local information into information extraction systems by gibbs sampling. In *Proceedings of the 43rd Annual Meeting on Association for Computational Linguistics, ACL '05*, 2005.
- [8] Sebastian Hellmann, Jens Lehmann, Sören Auer, and Martin Brümmer. Integrating NLP using linked data. In *The Semantic Web - ISWC 2013 - 12th International Semantic Web Conference, Sydney, NSW, Australia, October 21-25, 2013, Proceedings, Part II*, pages 98–113, 2013.
- [9] Beth Levin. *English Verb Classes and Alternations:~A Preliminary Investigation*. University of Chicago Press, 1993.
- [10] John P. McCrae, Dennis Spohr, and Philipp Cimiano. Linking lexical resources and ontologies on the semantic web with lemon. In *The Semantic Web: Research and Applications - 8th Extended Semantic Web Conference, ESWC 2011*.
- [11] Pablo N. Mendes, Max Jakob, Andrés García-Silva, and Christian Bizer. Dbpedia spotlight: shedding light on the web of documents. In *Proceedings the 7th International Conference on Semantic Systems, I-SEMANTICS 2011, Graz, Austria, September 7-9, 2011*, pages 1–8, 2011.
- [12] R. C. Schank. Conceptual dependency: A theory of natural language processing. 3, 1972.
- [13] R.C. Schank. *Conceptual Information Processing*. North-Holland, Amsterdam, 1975.
- [14] Kuldeep Singh, Andreas Both, Dennis Diefenbach, and Saeedeh Shekarpour. Towards a message-driven vocabulary for promoting the interoperability of question answering systems. In *Tenth IEEE International Conference on Semantic Computing, ICSC 2016, Laguna Hills, CA, USA, February 4-6, 2016*, pages 386–389, 2016.
- [15] Christina Unger, John P. McCrae, Sebastian Walter, Sara Winter, and Philipp Cimiano. A lemon lexicon for dbpedia. In *Proceedings of the NLP & DBpedia workshop co-located with the 12th International Semantic Web Conference (ISWC 2013), Sydney, Australia, October 22, 2013*.