

The Rich Event Ontology

Susan Windisch Brown,¹ Claire Bonial,² Leo Obrst,³ and Martha Palmer¹

¹ University of Colorado, Boulder, Colorado

² U.S. Army Research Lab, Adelphi, Maryland

³ MITRE, McLean, Virginia

Susan.Brown@colorado.edu, Claire.N.Bonial.civ@mail.mil,
lobrst@mitre.org, Martha.Palmer@colorado.edu

Abstract

In this paper we describe a new lexical semantic resource, The Rich Event Ontology, which provides an independent conceptual backbone to unify existing semantic role labeling (SRL) schemas and augment them with event-to-event causal and temporal relations. By unifying the FrameNet, VerbNet, Automatic Content Extraction, and Rich Entities, Relations and Events resources, the ontology serves as a shared hub for the disparate annotation schemas and therefore enables the combination of SRL training data into a larger, more diverse corpus. By adding temporal and causal relational information not found in any of the independent resources, the ontology facilitates reasoning on and across documents, revealing relationships between events that come together in temporal and causal chains to build more complex scenarios. We envision the open resource serving as a valuable tool for both moving from the ontology to text to query for event types and scenarios of interest, and for moving from text to the ontology to access interpretations of events using the combined semantic information housed there.

1 Introduction

As NLP moves into tasks requiring deeper language understanding, inferencing, and reasoning, knowledge-based resources are being increasingly

called on to support and supplement probabilistic and other data-driven methods (Hogenboom et al., 2011). Ontologies have been recognized as useful for tasks such as information extraction (IE) (Maedche et al., 2003; Wimalasuriyu et al., 2010), metaphor analysis (Brown, 2014) and automatic question answering (Lopez et al., 2011). By providing a formal specification of the shared concepts in a domain, an ontology allows users to identify entities and relations between them despite the myriad ways these can be expressed in language.

Existing general-purpose ontologies, such as the Descriptive Ontology for Linguistic and Cognitive Engineering, DOLCE (Masolo et al., 2003), the Suggested Upper Merged Ontology, SUMO (Pease, 2002), Cyc (Lenat, 1995), and the Basic Formal Ontology, BFO (Smith & Grenon, 2002) have either focused on providing only a very under-specified upper level ontology to which domain-specific ontologies can attach or have created much more fully developed object hierarchies than event hierarchies. SUMO has links to the well-known WordNet lexicon (Fellbaum, 1998), which is also the foundation for the BabelNet ontology (Navigli and Ponzetto, 2012). WordNet has well-developed subsumption relations in its noun lexicon. It's verb lexicon, however, has hyponym/meronym relations only four to five nodes deep. This situation translates to an ontology rich with object concepts and relations but a rather impoverished event network. In addition, none of

these ontologies have incorporated information from lexical resources that focus on events.

Most applications using ontologies have made heaviest use of these ontologies' object hierarchies, drawing on their often extensive representations of physical objects, people, and locations. Events, being more difficult to delineate and define, often have a sparser and more shallow class hierarchy in an ontology. Classes representing events, however, can provide the nexus for relating objects and properties and prove useful for many language understanding tasks. By explicitly representing events, we can deal directly not only with relations between events and objects, but between multiple events as well. One of the more difficult language understanding tasks is identifying temporal and causal relations between events. The ontology we describe here is intended to provide a rich structure of event concepts that connects varying levels of event specificity, relates events to their key objects and participants, and encodes the temporal and causal relationships between events.

We found that existing ontologies were not suitable for bridging the gap between spatio-temporal ontological approaches to representing events and the representations stemming from SRL resources. Our ontology provides this bridge by drawing heavily from the upper-level distinctions of DOLCE, but also linking to the widely used lexical resources FrameNet (FN) (Fillmore et al., 2002) and VerbNet (VN) (Kipper et al., 2008). Not only do these provide wide-coverage lexicons having to do with events, they also contribute annotated corpora and additional semantic and syntactic information that can be crucial to identifying events and their participants (see section 2.5). In addition, the ontology provides links to the annotations, event typing, and role specifications of both the Automatic Content Extraction (ACE) (Doddington et al., 2004) and the Rich Entities, Relations, and Events (ERE) (Song et al., 2015) Projects. Both are DARPA-funded resources that have figured prominently in the TAC-KBP (Text Analysis Conference – Knowledge Base Population) evaluations. The ontology thus allows one to draw on multiple linguistic resources and combine their annotations. This firstly ensures a larger, more diverse training corpus with the potential to detect a wide variety of events. Secondly, this allows the resources to be integrated in terms of common temporal and causal relationships be-

tween annotated event types, making explicit higher-order relationships between events – information not found in any of the independent resources.

We have completed the early stages of ontology development and are now working toward a formal evaluation. To that end, we are integrating the ontology into an end-to-end IE pipeline in order to evaluate the ontology's ability to 1) increase the number and types of events recognized and classified in text, and 2) allow users to refine, expand or alter queries about events by making use of ontological relations. We report results on two sample use cases related to these goals.

In the remainder of the paper, we provide a description of the upper level of the ontology, some of the major mid-level classes, and the linked lexical resources. We then explain the modular structure of the ontology and its advantages. In section 3, we describe our progress towards evaluation by discussing our two use cases. Finally, we conclude with a description of our future work.

2 Ontology Description

Intended as a resource for a wide range of tasks, the Rich Event Ontology (REO) has been designed to encompass both meta-level concepts in its upper level and many general domains in its mid level. REO has been implemented in OWL, which allows for easy extension with more detailed, domain-specific ontologies. The main reference ontology now encompasses 161 classes and 553 axioms. Including the lexical resource ontologies and the linking models (described in detail in sections 2.5 and 2.6) in these counts brings the totals to 3,065 classes and 60,531 axioms, as well as 16,005 individuals representing the vocabulary (unique lemmas) of event denotations.¹

This project's goal has been the development of a unified representation of events. To do this, however, we must be able to reference the participants of the events, necessitating a connection to a well-developed physical and abstract object ontology. Although this paper will include mentions of object classes, especially as they link to event classes as participants in those events, it will focus on the event portion of the ontology. In addition, we will focus on a description of the ontology's structure and content, rather than a description of

¹ For comparison: VN includes about 8,600 verb lemmas and FN includes about 13,000 lexical units.

our development methodology, which can be found in earlier work (Bonial et al., 2016).

2.1 Theoretical Framework & Approach

We attempt to describe those categories that underlie human language. DOLCE’s basic assumptions reflect our own: “We do not commit to a strictly referentialist metaphysics related to the intrinsic nature of the world: rather, the categories we introduce here are thought of as cognitive artifacts ultimately depending on human perception, cultural imprints and social conventions” (Masolo et al., 2003, p. 8). Our upper level ontological distinctions align with DOLCE’s largely spatio-temporal distinctions. However, given our practical NLP goals, our mid-level distinctions shift towards Davidsonian (Davidson, 1980) distinctions more aligned with SRL resources.

2.2 The Upper Ontology

The fundamental distinction at the top level of our ontology is between Endurant and Perdurant entities. Borrowing heavily from DOLCE, we define “Endurants” as *those entities that can be observed/perceived as a complete concept, no matter which given snapshot of time* and “Perdurants” as *those entities for which only a part exists if we look at them at any given snapshot in time. Various called events, processes phenomena, or activities and states, perdurants have temporal parts or spatial parts and participants*. We continue to follow DOLCE’s lead in dividing the PERDURANT class into the subclasses EVENTIVE PERDURANT and STATIVE PERDURANT. This dichotomy is based on the notions of homeomericity and cumulativity (Masolo et al., 2003). So, in this case, a stative would be distinguished from an eventive by way of possessing the property of cumulativity, i.e., a *sitting* occurrence type is a stative because the mereological sum of two sittings is still a sitting. This is somewhat similar to the “waterfall” analysis of Galton & Mizoguchi (2009), that more radically proposes a property of dissectivity for processes and matter, so that processes are similar to mass nouns in semantics. In the waterfall model, processes are dependent continuants, similar to objects, which are independent continuants. Unlike DOLCE and the waterfall model (see also Mizoguchi et al (2011), Galton (2012), Borgo & Mizoguchi (2014), Rovetto & Mizoguchi (2015), which more directly address notions of causality), however, we do not currently subdivide these cat-

egories into the aspectual classes of states, processes, achievements and accomplishments. Although these categories have a long history in linguistic and philosophical literature (Vendler, 1957; Moens and Steedman, 1988) and more recently in semantics, distinguishing kinds of states (Maienborn, 2011; Maienborn et al., 2011), these divisions are difficult to apply in a commonsense way to domains we consider coherent. For example, Vendlerian divisions would place a chatting eventuality in a fundamentally different section of the ontology from a telling eventuality. Instead, as we move into the middle level of the ontology, we shift to a neo-Davidsonian perspective, in which event participants become a greater focus. We expect that we will refine the underlying event formalization over time, as it becomes clearer how to reconcile our commonsense semantic application focus with more recent semantic and ontological analyses.

2.3 Mid-level Classes

The EVENTIVE PERDURANT class splits into many daughter classes, of which some of the most extensive are COGNITIVE EVENT, LIFE EVENT, INTENTIONALLY ACT, and MOTION. These are still very general concepts, and have no direct connections to the lexical resources and specific lexical items. For some of these classes, such as LIFE EVENT, the next level down introduces concepts with direct links to the lexical resources, such as the LIFE EVENT daughter class BIRTH linking to FN’s BEING_BORN frame and VN’s BIRTH class (among others).

For other classes, another sublevel with few direct lexical realizations seemed necessary. For example, INTENTIONALLY ACT includes the subclasses SOCIAL INTERACTION, INTENTIONALLY AFFECT, TRANSFER POSSESSION, and ORGANIZATIONAL EVENT. Each of these has multiple subclasses. To illustrate the level of class granularity, we present ORGANIZATIONAL EVENT in more detail (Figure 1).

Its daughter classes include START ORGANIZATION, END ORGANIZATION, MERGE ORGANIZATION, DECLARE BANKRUPTCY, START POSITION WITH AN ORGANIZATION, and END POSITION WITH AN ORGANIZATION. Most of these have no further subclasses, although START POSITION subdivides further into START LEADERSHIP POSITION, HIRING, and HIRING ON. END POSITION has similar subclasses. The decision to include the very specific classes

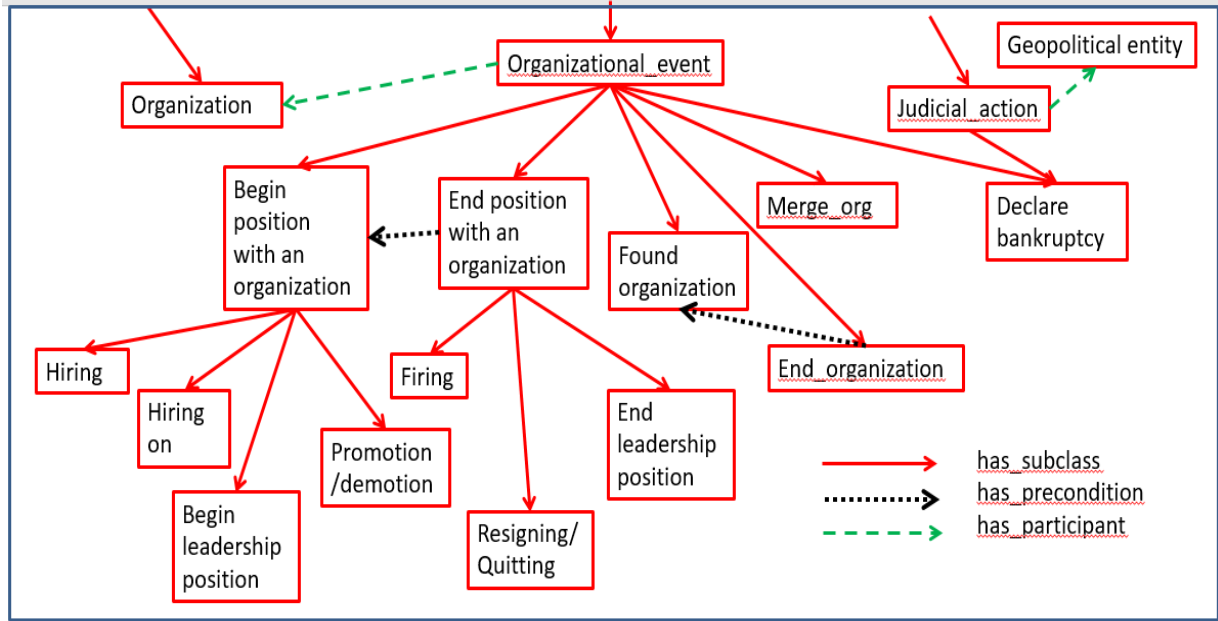


Figure 1: ORGANIZATIONAL EVENT section of the ontology.

concerning leadership positions resulted from the many lexical items, across languages, for events like ‘crown’, ‘ordain’, ‘oust’ and ‘depose’, and the frequency with which starting and ending leadership positions are discussed in print and oral corpora.

The decision to create the closely related classes HIRING and HIRING ON stems from a similar desire to take common human distinctions into account and to allow for the shift in role relations that usually accompany such shifts in perspective. The agent of a hiring event is the employer and the employee is a theme. However, the agent of a hiring-on event is the employee. Although *Company hiring Person* is arguably the same event as *Person hiring on with Company*, the shift in perspective is commonly lexicalized and therefore represented in the ontology. Such perspective-shifting classes are rare in the ontology and always share a common parent class, which ignores the perspective shift. They are important, however, in the TRANSFER POSSESSION domain, with such divisions as GIVE and GET. We highlight the perspective shift by having two relations between a class like TRANSFER and a class like GIVE: both TRANSFER *hasSubclass* GIVE and TRANSFER *hasPerspective* GIVE. For applications that need a more perspective-neutral classification, one can generalize to the parent class.

2.4 Relations between Classes

The main relation between classes (i.e., concepts) in the ontology is the subclass relation, which specifies that every subclass is a more specific

type of the superclass. This entails that a subclass inherits all the domain and range restrictions of the parent class as well as other types of relations the parent class holds, such as *hasResult*.

The subclass relation, however, barely taps into the rich, complex relations between events or between events and objects. To capture some of that, we have included temporal and causal relations extended from the Richer Event Description (RED) project (Ikuta et al., 2014; O’Gorman et al., 2016). The RED project aims to annotate text with mentions of eventualities and entities, with the goal of representing the temporal and causal relationships between those eventualities in such a way that an accurate timeline of events could be automatically constructed. We have adapted and expanded their relations to our *hasPrecondition*, *hasCause*, *hasResult*, and *hasSubevent* relations.² Examples of these relations include:

1. END ORG *hasPrecondition* BEGIN ORG
2. KILLING *hasResult* DYING
3. TRIAL *hasSubevent* VERDICT

The *hasSubevent* relation is intended to capture events that are temporally contained within another event and considered a proper part of that event. For example, Verdict is not a type of Trial, so the *Subclass* relation is inappropriate. The *has-*

² In some cases the relations encode opposite perspectives on the same relation between classes (e.g., DEAD *hasCause* DYING and DYING *hasResult* DEAD), but those relations do not always coincide (e.g., (2) does not entail that DYING *hasCause* KILLING).

Subevent relation, however, indicates that a verdict happens within the greater context of a trial.

We have currently defined ten such cross-event relations. As part of the process of selecting and defining these relations, we created 49 instances of event-to-event relations in a small portion of the existing ontology. Future work will involve applying these relations to the rest of the reference ontology.

Other relations connect events with object classes (physical or abstract), such as the *hasLocation*, *hasAgent*, and *hasPatient* relations. As mentioned earlier, these relations are inherited by descendent classes. For example, DECLARE BANKRUPTCY is a subclass of both ORGANIZATIONAL PROCESS and JUDICIAL ACTION. ORGANIZATIONAL PROCESS *hasParticipant* some ORGANIZATION, and JUDICIAL ACTION *hasParticipant* some GOVERNMENTAL AUTHORITY. DECLARE BANKRUPTCY would thus inherit both ORGANIZATION and GOVERNMENTAL AUTHORITY as participants in the event.

The relations described in this section are being applied to the main, “reference” REO ontology. For an explanation of how the main ontology links to the lexical resources, see section 2.6.

2.5 Lexical Resource Ontologies and Their Linking Models

One of the primary goals of the ontology is to provide a means of combining the information in multiple lexical resources, despite differences in their categorization of lexical items. With our focus on event modeling, we have chosen to link to resources with rich event representations and broad coverage of English verbs and eventive nouns. We have represented the categorizations, lexical items, and participant roles included in each of these resources as separate OWL ontologies.

FrameNet: This resource, based on Fillmore’s frame semantics (Fillmore, 1976; Fillmore & Baker, 2001), groups verbs, nouns and adjectives into “frames” based on words or “frame elements” that evoke the same semantic frame: a description of a type of event, relation, or entity and the participants in it. For example, the Apply_heat frame includes the frame elements Cook, Food, Heating_instrument, Temperature_setting, etc. The “net” of frames makes up a rather complex network, including simple *isA* inheritance relations as well as more complex relations such as *Precedes* and *PerspectiveOn*.

These relations highlight important aspects of many frames, for example, the Apply_heat frame is *UsedBy* the Cooking_creation frame, but often the frames involved are not anchored to the main *isA* hierarchy. In addition, the automatic reasoning capabilities of ontologies implemented in OWL are restricted to strictly logical relationships between classes. The complexity of FN precludes complete representation in OWL, as others have found (e.g., Scheffczyk et al., 2006). Therefore, we flattened the FN hierarchy, connecting every frame to a single parent node, FrameNetFrame, and relying on our main ontology to provide *isA* and event-event relations. This decision reduces the relational information from FN that is directly represented in our ontology, but users can of course trace the frames back to FN proper and access FN’s full relational structure there.

VerbNet: This resource, based on Levin (1993), groups verbs into “classes” using their compatibility with certain syntactic alternations (e.g., *She rolled the ball down the hill* vs. *The ball rolled down the hill*). Although the groupings are primarily syntactic, the classes do share semantic features as well, since, as Levin posited, the syntactic behavior of a verb is largely determined by its meaning. Each class specifies its member verbs and their typical participants (i.e., semantic roles), lists the syntactic patterns they are all compatible with, and connects those patterns to semantic representations (Kipper et al., 2008).

By linking to VN, the ontology gains valuable syntactic information about how events are expressed in English. Generally, a VN class is linked in a one-to-one relation to one of the main ontology classes. A class’s syntactic alternations, however, sometimes cut across semantic distinctions made by the main ontology. For example, events expressible with causative-inchoative alternations are grouped in the same VN class, but are divided in the main ontology (since the main ontology makes distinctions based on the number and types of event participants). For these VN classes, we link an ontology class to specific frames in a class, using VN thematic roles to distinguish the appropriate frames. These cases coincide with places where VN’s semantic representation also differs for a particular frame, indicating that the reference ontology is consistent with VN semantic distinctions.

ERE/ACE: ERE is based on the ACE project’s semantic role annotation schema. The goal of the

ERE/ACE projects is to mark up the events and the entities involved in them, and to mark coreference between these. This provides a somewhat shallow representation of the meaning of the text. The ERE/ACE schema can also serve as a lexicon imported into the ontology, with its event type designations serving as links to the lexical items marked up with that designation. ERE annotated eventualities are limited to certain types of special interest within the defense community, with top-level types referred to as *Life, Movement, Transaction, Business, Conflict, Manufacture, Contact, Personnel* and *Justice* events.

Both the FN and VN resource ontologies model lexical units and class members, respectively, as individuals that represent lemmas, which may be used as references for particular event concepts in REO. Because ERE and ACE are resources developed specifically for annotating data to be used as training data, they do not include pre-specified individuals or “triggers,” of certain event types. Instead, these are always marked up in context. Thus, these resources provide a data-driven, ground-up perspective on event semantics that is very distinct from the other resources. The ACE and ERE models include as individuals English lemmas that have been annotated either in the freely available ACE 2005 Multilingual Training Corpus (Walker et al., 2005), or the as-of-yet unreleased ERE corpus, respectively.

2.6 Modular Architecture

The structure of the ontology is modeled after the architecture of the Ontologies of Linguistic Annotation (OLiA) (Chiarcos et al., 2016). OLiA serves as a reference hub for annotation terminology for largely (morpho-)syntactic information across a variety of languages. Similarly, REO can act as a bridge between semantic annotation resources. In this modular architecture (Figure 2), one reference ontology houses the schema-independent, primary event concepts and relations of REO. Each of the lexical resources currently included, FrameNet, VerbNet, ERE and ACE, are modeled as independent OWL ontologies, as described above. For each annotation resource model, a linking model defines the relationships between the concepts and

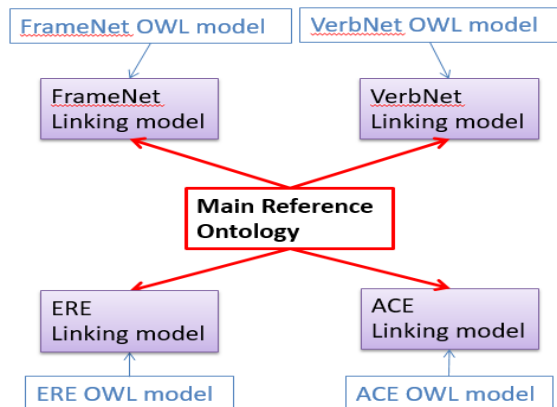


Figure 2: Modular architecture of the ontology

properties in the resource model and those of the reference model. Specifically, each linking model imports both the respective resource model and the reference ontology, and concepts in the reference ontology are linked to those of the resource model via the *hasReferenceGroup* relation. For example, the LEGAL ACTION event subclass DISCHARGE has the reference group Release-Parole from ERE and Releasing from FN (see Figure 3). Thus, all of the lexical units that are members of the Releasing Frame and all of the triggers annotated as Release-Parole form the group of references for a DISCHARGE event: *free, parole, release, let go, set free*, etc. Each of the linking models can be imported into a single ontology to query across all resources simultaneously. However, as Chiarcos et al. (2016) point out, maintaining independent ontologies in this modular structure allows one to integrate, or remove, terminology from different resources in a lossless and reversible way. Additionally, given the ongoing development of resources like FN, this structure also allows for independent lexical resource models to be updated without impacting the ontology as a whole. Finally, the modularity offers a certain level of customization for users. For example, if a user is looking for somewhat synonymous references to events, then it may be desirable to leave FN out of the final model, since FN frames include Frame Elements that may not be references to the event (e.g., *cop* in the Arrest frame).

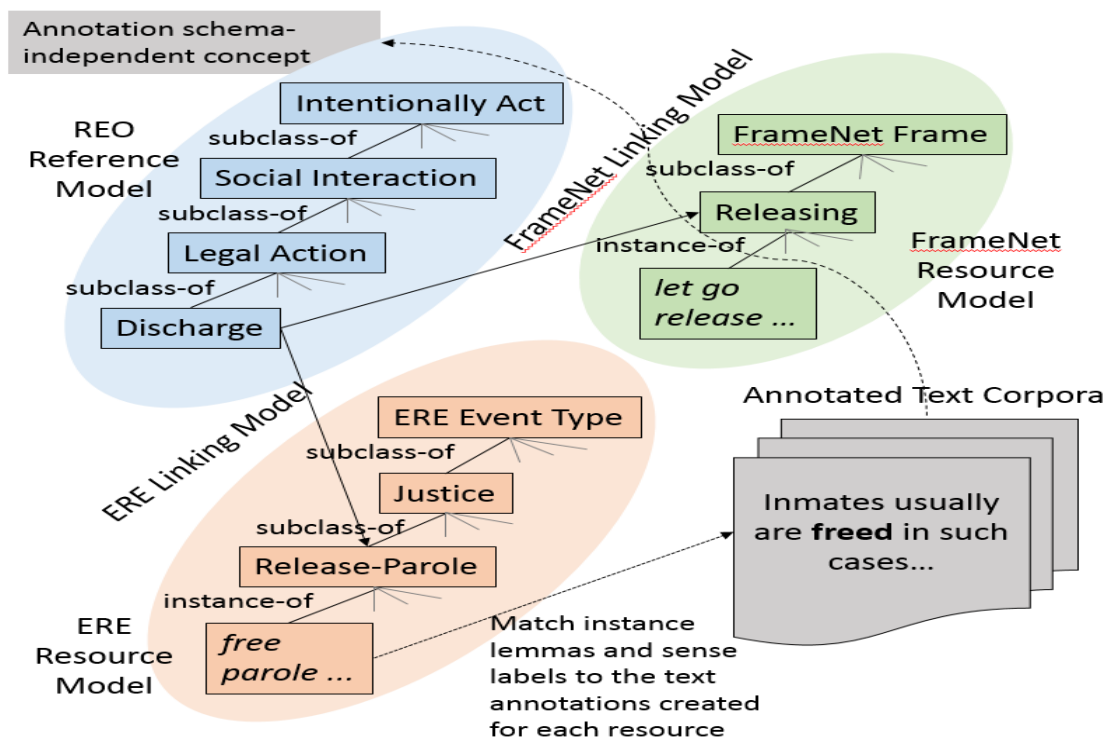


Figure 3: DISCHARGE, a daughter of LEGAL ACTION, *hasReferenceGroup* Releasing in FN and Release-Parole in ERE. The lexical items in these classes can be matched to sense-annotated lemmas found in the annotated corpora, and/or users can query the ontology for events to obtain a schema-neutral representation, including event-event relations, and access the combined semantic information from resources.

3 Use Cases of REO

We are working to integrate REO into an IE pipeline designed for intelligence analyst use. Within the pipeline we will evaluate the ontology’s impact on two main areas. 1) Increasing the number and types of events recognized and classified in text. We will be incrementally examining the precision, recall and F-score of trigger identification and classification in systems that are trained on just ACE data, then ACE+ERE, ACE+ERE+FN, and finally all data sources: ACE+ERE+FN+VN. 2) Allowing users to refine, expand or alter queries about events by making use of ontological relations. We will be completing user studies for this evaluation and comparing efficiency in decision-making using the IE pipeline with and without the event ontology component. In the interim, we report results below on two sample use cases related to these goals.

3.1 Expanding Lexical Triggers for IE

The ontology can be leveraged to support event detection in IE systems by expanding the number and variety of lexemes recognized as potentially referring to a given event type. The aforementioned ACE program, and its inclusion in TAC, has established the ACE annotated data as a benchmark dataset for IE systems. As a result, many existing IE systems are tailored to, and can be limited to, the detection of events recognized and marked up in the ACE annotated data. To avoid the need for additional manually annotated data, the ontology and associated lexical resources can be used in backoff techniques to augment the trigger words associated with certain types of events, thus expanding the domain of application.

To explore the potential efficacy of the ontology in this application, we examined the reference groups associated with the LEGAL ACTION portion of the ontology. LEGAL ACTION is a type of SOCIAL INTERACTION, and is the parent class of several subclasses, including ARREST, SUE, and DECLARE BANKRUPTCY (which also inherits from ORGANIZATIONAL PROCESS). We first established a baseline of what a typical system, trained on

ACE, might recognize as triggers associated with the event concepts in this portion of the ontology. To do this, we examined what ACE types and subtypes are linked to the subclasses of LEGAL ACTION via the *hasReferenceGroup* relation. We then extracted all of the individuals that have been tagged as triggers for the *hasReferenceGroup* linked event types and subtypes. In total, we found 102 lexemes associated with the LEGAL ACTION subtypes in ACE. Presumably, systems trained on ACE data have the potential to recognize these lexemes as triggers of the LEGAL ACTION events.

To determine how the ontology may help to move beyond this baseline, we examined what other triggers might be found by using the ontology to access lexemes in the reference groups associated with LEGAL ACTION in ERE, FN and VN. This allowed us to extract groups of 204, 69 and 14 lexemes from ERE, FN and VN, respectively. Thus, we were able to expand the vocabulary of what lexemes may denote subtypes of LEGAL ACTION from 102 words to 389 words. This is summarized in Table 1.

Source	ACE	ACE + ERE	ACE + ERE + FN	ACE + ERE + FN + VN
Trigger Total	102	306	375	389
Ex's	<i>Arrest, fine, prosecute</i>	<i>Behind bars, get_life</i>	<i>Bust, put_away, guillotine</i>	<i>Book, collar, impeach</i>

Table 1: Expansion of event trigger vocabulary using the REO class LEGAL ACTION.

The variety of triggers found across the resources is quite remarkable: only 17 of the 389 lexemes are duplicated from one resource to another. We see the data-driven resources, ACE and ERE, capturing much more informal expressions, such as *share a needle*, referring to an execution event. In contrast, FN and VN capture more formal expressions like *mulct* and *amerce*, referring to fining events. Furthermore, few nodes in the ontology have reference groups in all four resources. For example, only FN distinguishes events at a level of specificity fine-grained enough to have a specific frame for *Notification_of_charges*, which is a reference group for the CHARGE events node of the ontology. We feel

that this highlights the potential for the ontology to overcome data sparsity by combining resources.

3.2 Querying: From Events to Scenarios

Although a mapping (similar to SemLink (Palmer, 2009)) of the resources included in the ontology may be able to achieve the vocabulary expansion described in the previous section, a unique contribution of the ontology is the causal and temporal event relations included. With the exception of limited relations in FN, the linked lexical resources do not provide information on such relations. The ontology has adapted the RED relations, as described in section 2.5, and therefore allows insights into how events are typically related, both causally and temporally. This can enable an understanding of how individual events fit into more complex real-world scenarios. What's more, users can take advantage of the temporal and causal relations in addition to subclass 'is-a' relations to expand, refine, or alter their queries.

One area of the ontology where these relations are particularly rich and informative is the domain of conflict. PROTEST, ATTACK, and RECIPROCAL CONFLICT are three daughters of the SOCIAL INTERACTION class CONFLICT. As in other areas of the ontology, we drew upon domain expertise in the development of this area. We reviewed social science literature to establish the basic sub-events and preconditions of PROTEST. Combining research on both the psychology of protest (Van Stekelenburg and Klandermans, 2013) and the theory of planned behavior generally (Ajzen, 1991), we established subevents and stages of protest scenarios: PROTEST has as a precondition MOBILIZATION, which in turn has TAKE SIDE as a precondition; TAKE SIDE has GROUP IDENTITY as a precondition, as well as the typical precondition GRIEVANCE; a communication event is a sub-event of PROTEST. This excerpt of some of the relations to PROTEST captures social science theories suggesting that a protest is generally mobilized where there is a sense of a group identity and a grievance or trigger for intergroup conflict, and that protest by nature involves the communication of some claims calling for change. The event structure found in the ontology for PROTEST parallels the "stages" of protest outlined in Korolov et al. (2016), who find that trigger words associated with these stages can be used to predict social protest based on social media messaging.

REO users can take advantage of ontological relations in their queries. For example, a user interested in protest may start by querying for documents with PROTEST event trigger words (e.g., *boycott, burn, loot, march, occupation, take to the streets*, etc.), with accompanying SRL-annotated training data sentences, such as “*The events which unfolded over last week are still very unclear but peaceful protesters **took to the streets** in Tottenham Saturday to demand answers.*” If users decided they were interested in a broader range of events, including both physical attacks and arguments, they could broaden the search space using the CONFLICT node of the ontology. If users were interested in querying for events that may be indicators of protest to come, they could query for the preconditions of protest, including TAKE SIDE with associated triggers *endorse, oppose, pro, side*, etc. Thus, the ontology links the annotated resources in a way that uniquely allows for users to search for events that are related to others in higher-order scenarios.

4 Future Work

The modular architecture of the ontology was designed to allow efficient linking to other lexical resources, including those from other languages. We intend to pursue such expansion, as well as expansion of the main ontology through alignment with or importation of other ontologies, such as the Emotion Ontology (Hastings et al., 2011).

Although we have emphasized the ontology’s NLP applications, we have also begun testing the ontology’s usefulness for activity recognition in video. We are currently exploring the use of REO for understanding how complex activities can be decomposed into simpler events, and how those events are broken down into semantic components in the linked resource VN. We hypothesize that activities that share similar event semantics will likely have some similar visual components. The potential to detect similar visual components may allow for generalizing from the recognition of one activity type (e.g., baseball pitch) to another that is semantically similar (e.g., throw discus). Thus, we hope to leverage information from the ontology instead of seeking out greater amounts of training data specific to fine-grained activity types.

We are also exploring new types of event-to-event relations that could enhance the inferencing power of the ontology. The logic requirements of

OWL have prevented us from capturing relations that are not necessary but still highly probable. For example, a TRIAL event typically follows a CHARGE/INDICT event, but not always. We would like to explore ways to marry probabilistic methods with the ontology to allow for such common-sense (but not strictly logical) inferences.

5 Conclusion

The Rich Event Ontology is a freely available tool for semantic analysis of events, a key area in NLP tasks like question answering, information extraction, and knowledge representation. It provides an independent conceptual backbone that unifies valuable lexical resources and adds critical relational information in the form of event-to-event causal and temporal relations. Although this work is in the relatively early stages, we have shown how the ontology could be used to expand the number and variety of lexemes recognized as event denotations and to refine, expand or shift user queries using both subclass and temporal relations. We believe REO is unique among existing ontologies in combining in-depth representation of events with the ability to link valuable but disparate lexical resources and annotation schemes. REO is temporarily available by request, but we plan to migrate the ontology to an in-house server in the near future, where it will be freely available.

Acknowledgments

This project has deeply benefited from the expertise of the REO Advisory Board: Ann Bies, Diana McCarthy, Teruko Mitamura, James Pustejovsky, German Rigau, Roxane Segers, Clare Voss, Piek Vossen, Annie Zaenen, and additional members of the board. We sincerely thank them for their wise advice and the generous donation of their time.

In addition, we gratefully acknowledge the support of DARPA DEFT - FA-8750-13-2-0045 and DTRA HDTRA1 -16-1-0002/Project # 1553695, eTASC - Empirical Evidence for a Theoretical Approach to Semantic Components. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of DARPA or the US government.

Finally, we thank our anonymous reviewers for their constructive feedback.

References

- Ajzen, I., 1991. The theory of planned behavior. *Organizational behavior and human decision processes*, 50(2): 179-211.
- Bonial, Claire, Susan Windisch Brown and Martha Palmer. 2016. A lexically-Informed upper level event ontology. In *Proceedings of the LREC 2016 Workshop ISA-12 – 12th Joint ACL - ISO Workshop on Interoperable Semantic Annotation*, 28 May, Portorož, Slovenia.
- Brown, Susan Windisch. 2014. From visual prototypes of action to metaphors: extending the imagact ontology of action to secondary meanings. In *Proceedings of the Tenth Joint ISO – ACL SIGSEM Workshop on Interoperable Semantic Annotation*, Reykjavik, May.
- Chiarcos, C., C. Fäth, and M. Sukhareva. 2016. Developing and using the ontologies of linguistic annotation (2006-2016). In *Proceedings of the LREC 2016 Workshop “LDL 2016 – 5th Workshop on Linked Data in Linguistics: Managing, Building and Using Linked Language Resources”*.
- Davidson, D. 1980. *Essays on Actions and Events*. Oxford: Clarendon Press.
- Doddington, G. R., A. Mitchell, M. A. Przybocki, L. A. Ramshaw, S. Strassel and R. M. Weischedel. 2004. The Automatic Content Extraction (ACE) program-tasks, data, and evaluation. In *LREC 2*: 1.
- Fellbaum, Christiane, ed. 1998. *WordNet: An Electronic Lexical Database*. MIT Press.
- Fillmore, Charles J. 1976. Frame semantics and the nature of language*. *Annals of the New York Academy of Sciences*, 280(1): 20-32.
- Fillmore, Charles J., and Collin F. Baker. 2001. Frame semantics for text understanding. In *Proceedings of WordNet and Other Lexical Resources Workshop, NAACL*.
- Fillmore, Charles J., Christopher R. Johnson, and Miriam R. L. Petruck. 2002. Background to FrameNet. *International Journal of Lexicography*, 16(3): 235-250.
- Galton, Antony, and Riichiro Mizoguchi. 2009. The water falls but the waterfall does not fall: New perspectives on objects, processes and events. *Applied Ontology* 4: 71–107.
- Galton, A. 2012. The ontology of states, processes, and events. *InterOntology 2012*, 5: 35 - 45.
- Hastings, J., W. Ceusters, B. Smith, and K. Mulligan. 2011. The emotion ontology: Enabling interdisciplinary research in the affective sciences. In *Modeling and Using Context*, pp. 119-123. Springer.
- Hogenboom, F., F. Frasincar, U. Kaymak, and F. De Jong. 2011. An overview of event extraction from text. *Workshop on Detection, Representation, and Exploitation of Events in the Semantic Web (DeRiVE 2011)* at Tenth International Semantic Web Conference (ISWC 2011).
- Ikuta, R., W. F. Styler IV, M. Hamang, T. O’Gorman, and M. Palmer. 2014. Challenges of adding causation to richer event descriptions. In *Proceeding of ACL 2014*.
- Kipper, Karin, Anna Korhonen, Neville Ryant, and Martha Palmer. 2008. A large-scale classification of English verbs. *Language Resources and Evaluation Journal*, 42: 21– 40.
- Korolov, R., D. Lu, J. Wang, G. Zhou, C. Bonial, C. Voss, L. Kaplan, W. Wallace, J. Han, and H. Ji. 2016. On predicting social unrest using social media. In *Advances in Social Networks Analysis and Mining (ASONAM), 2016 IEEE/ACM International Conference on* (pp. 89-95), August.
- Lenat, D. B. 1995. Cyc: A large-scale investment in knowledge infrastructure, *CACM* 38(11): 33-38.
- Levin, B. 1993. *English Verb Classes and Alternations: A Preliminary Investigation*. University of Chicago Press.
- Lopez, Vanessa, Victoria Uren, Marta Sabou, and Enrico Motta. 2011. Is question answering fit for the semantic web?: A survey. *Semantic Web*, 2(2): 125-55.
- Maedche, Alexander, Günter Neumann, and Steffen Staab. 2003. Bootstrapping an ontology-based information extraction system. In *Intelligent Exploration of the Web*, pp. 345-359. Physica-Verlag HD.
- Maienborn, Claudia, Klaus von Heusinger, and Paul Portner, eds. 2011. *Semantics: An International Handbook of Natural Language Meaning*. Vol. 33. Walter de Gruyter.
- Maienborn, Claudia. 2011. Event semantics. In *Semantics: An international handbook of natural language meaning*. Vol. 33. Maienborn, Claudia, Klaus von Heusinger, and Paul Portner, eds., pp. 802-829. Walter de Gruyter,
- Masolo, C., S. Borgo, A. Gangemi, N. Guarino, A. Oltramari, and L. Schneider. 2003. DOLCE: a descriptive ontology for linguistic and cognitive engineering. WonderWeb Project, Deliverable D, 17.
- Mizoguchi, R., K. Kozakil, H. Kou, Y. Yamagata, T. Imai, K. Waki and K. Ohe. 2011. River flow model of diseases. In *Proceedings of International Conference on Biomedical Ontology*, pp. 63–70.

- Moens, Marc, and Marc Steedman. 1988. Temporal ontology and temporal reference. *Computational Linguistics*, 14: 15-28.
- R. Navigli and S. Ponzetto. 2012. BabelNet: The automatic construction, evaluation and application of a wide-coverage multilingual semantic network. *Artificial Intelligence*, 193: 217-50.
- Niles, I., and A. Pease. 2001. Towards a standard upper ontology. In *Proceedings of the 2nd International Conference on Formal Ontology in Information Systems (FOIS-2001)*, Chris Welty and Barry Smith, eds. Ogunquit, Maine, October.
- O’Gorman, Tim, Kristen Wright-Bettner, and Martha Palmer. 2016. Richer Event Description: Integrating Event Co-reference with temporal, causal and bridging annotation. In *Proceedings of 2nd Workshop on Computing News Storylines*, pp. 47-56.
- Palmer, Martha. 2009. Semlink: Linking propbank, verbnet and framenet. In *Proceedings of the Generative Lexicon Conference (GenLex-09)*, pp. 9-15. Pisa, Italy, September.
- Pease, A., I. Niles, and J. Li. 2002. The suggested upper merged ontology: A large ontology for the semantic web and its applications. In *Working Notes of the AAAI-2002 Workshop on Ontologies and the Semantic Web* (Vol. 28), July.
- Rovetto, R. J. and R. Mizoguchi. 2015. Causality and the ontology of disease. *Applied Ontology* 10: 79–105.
- Scheffczyk, J., C. F. Baker, and S. Narayanan. 2006. Ontology-based reasoning about lexical resources. In *Proc. of OntoLex*, pp. 1-8.
- Smith, B., and P. Grenon. 2002. Basic formal ontology. *Draft. Downloadable at <http://ontology.buffalo.edu/bfo>.*
- Song, Z., A. Bies, S. Strassel, T. Riese, J. Mott, J. Ellis, ... and X. Ma. 2015. From light to rich ERE: Annotation of entities, relations, and events. In *Proceedings of the 3rd Workshop on EVENTS at the NAACL-HLT*, pp. 89-98, June.
- Van Stekelenburg, J., and B. Klandermans. 2013. The social psychology of protest. *Current Sociology*, p.0011392113479314.
- Vendler, Z. 1957. Verbs and times. *The Philosophical Review* 66(2): 143-60.
- Walker, Christopher, et al. 2005. *ACE Multilingual Training Corpus LDC2006T06*. DVD. Philadelphia: Linguistic Data Consortium.
- Wimalasuriya, Daya C., and Dejing Dou. 2010. Ontology-based information extraction: An introduction and survey of current approaches. *Journal of Information Science*, 36 (3): 306–323.