


A Comparison of the Events and Relations Across ACE, ERE, TAC-KBP, and FrameNet Annotation Standards

Jacqueline Aguilar¹, Charley Beller¹, Paul McNamee¹, Benjamin Van Durme¹, Stephanie Strassel², Zhiyi Song², Joe Ellis²

¹Human Language Technology Center of Excellence, Johns Hopkins University & ²Linguistic Data Consortium, University of Pennsylvania

Abstract

The resurgence of effort within computational semantics has led to increased interest in various types of relation extraction and semantic parsing. While various manually annotated resources exist for enabling this work, these materials have been developed with different standards and goals in mind. In an effort to develop better general understanding across these resources, we provide a summary overview of the standards underlying ACE, ERE, TAC-KBP Slot-filling, and FrameNet.



ACE and ERE Overview

- ACE and ERE are comprehensive annotation standards that aim to annotate Entities, Events, and Relations within a variety of documents.
- The ACE (Automatic Content Extraction) standard was developed by NIST in 1999 and has evolved over time to support different evaluation cycles, the last evaluation having occurred in 2008.
- The ERE (Entities, Relations, Events) standard was created under the DARPA DEFT program as a lighter-weight version of ACE with the goal of making annotation easier, and more consistent.
- ERE consolidates the more problematic annotation distinctions in ACE and removes some complex annotation features.
- Many of the differences in relation and events tagging across ERE and ACE stem from the way in which entities are handled.

ACE and ERE Events

Event Tagging Similarities:

- Both schemas have same exact Event Types: LIFE, MOVEMENT, TRANSACTION, BUSINESS, CONFLICT, CONTACT, PERSONNEL, JUSTICE
- Both ontologies include 33 Subtypes for each Type
- Both use *triggers* (ACE restricts it to be a single word)
- Both methods annotate modifiers when they trigger events
- When there is ambiguity about which trigger to use, both methods have similar rules established (Stand-Alone Noun and Adj. Rules)
- Both tag Resultative Events
- Nominalized Events are tagged as regular events
- Reported Events are **not** tagged
- Implicit Events are **not** tagged
- Coreferential Events are tagged
- Tagging of multi-part triggers (only if they are contiguous)

Event Tagging Differences:

- Event Extent: ACE defines it as always being the entire sentence within which the Event is described. In ERE, the extent is the entire document unless an event is coreferenced.
- ERE does **not** tag negative, future, hypothetical, conditional, uncertain, or generic events.
- ACE allows for irrealis events (events that may have occurred or have some probability of occurring in the future. *E.g., Rumors of arrests circulated in Vancouver.*)

Argument Tagging Differences:

- ERE is limited to pre-specified arguments for each event and relation subtype. ACE's arguments are: Event participants, Event-specific attributes that are associated with a particular event type, and General event attributes (*e.g.,* time, place)
- ERE only tags asserted participants in the event
- The full noun phrase is marked in both ERE and ACE arguments, but the head is only specified in ACE

Event Type and Subtype Differences:

- Types of entities that can be moved in the Movement Type: in ACE, ARTIFACT entities (WEAPON or VEHICLE) as well as PERSON entities can be transported, whereas in ERE, only PERSON entities can be transported

ACE and ERE Relations

Similarities:

- Limiting relations to only those expressed in a single sentence
- Tagging only for explicit mention
- No *promoting* or *nesting* of taggable entities
- Tagging for past and former relations
- Two different Argument slots (Arg1 and Arg2) are provided for each relation to capture the importance of Argument ordering
- Arguments can be more than one token (although ACE marks the head as well)
- Using *templates* for each relation Type/Subtype
- Neither model tags for negative relations
- Both methods contain argument span boundaries. That is, the relations should only include *tagged* entities within the extent of a sentence.

Differences in Assertion, Modality, Tense, Syntactic Classes, and Triggers:

- ERE only annotates asserted events; ACE also annotates hypotheticals (with two modality attributes: ASSERTED and OTHER)
- ACE tags PAST, FUTURE, PRESENT, and UNSPECIFIED relations
- ACE includes Syntactic Classes to serve as a restraint for tagging: Possessive, Preposition, PreMod, Coordination, Formulaic, Participial, Verbal, Relations Expressed by Verbs, and Other
- Triggers: ACE does not have triggers; it annotates the full syntactic clause. ERE has an optional trigger word (defined as *the smallest extent of text that indicates a Type/Subtype relation*)

Physical Relation Differences:

- ACE only marks Location for PERSON entities. ERE marks Location for PERSON, GEO-POLITICAL, and LOCATION entities
- ACE includes Near as a Subtype
- ERE includes Origin as a Subtype (ACE accounts for this in *General Affiliation* Type, Citizen-Resident-Religion-Ethnicity Subtype)

Relation Type	Relation Subtype	ARG1 Type	ARG2 Type
<i>ERE</i>			
Physical	Located	PER, GPE, LOC	GPE, LOC
Physical	Origin	PER, ORG	GPE, LOC
<i>ACE</i>			
Physical	Located	PER	FAC, LOC, GPE
Physical	Near	PER, FAC, GPE, LOC	FAC, GPE, LOC

Part-Whole Relation Differences: ACE includes a Geographical Subtype. This sort of relation is covered by *Physical. Located* in ERE.

Relation Type	Relation Subtype	ARG1 Type	ARG2 Type
<i>ERE</i>			
Part-Whole	Subsidiary	ORG	ORG, GPE
<i>ACE</i>			
Part-Whole	Geographical	FAC, LOC, GPE	FAC, LOC, GPE
Part-Whole	Subsidiary	ORG	ORG, GPE

Social Relation Differences:

- ACE and ERE have three Subtypes with similar goals (Business, Family, Unspecified/Lasting-Personal), but ERE has an additional Membership Subtype.
- ERE also includes the Social.Role Subtype in order to address the *TITLE* entity type, which only applies to ERE.

Relation Type	Relation Subtype	ARG1 Type	ARG2 Type
<i>ERE</i>			
Social	Business	PER	PER
Social	Family	PER	PER
Social	Membership	PER	PER
Social	Role	TTL	PER
Social	Unspecified	PER	PER
<i>ACE</i>			
Personal-Social	Business	PER	PER
Personal-Social	Family	PER	PER
Personal-Social	Lasting-Personal	PER	PER

Affiliation Relation Differences:

- ACE addresses all Membership relations in its *Affiliation* Type.
- ACE includes many Subtype possibilities which can more accurately represent affiliation, whereas ERE only observes two *Affiliation* Subtype options.

Relation Type	Relation Subtype	ARG1 Type	ARG2 Type
<i>ERE</i>			
Affiliation	Employment/Membership	PER, GPE	ORG, ORG, GPE
Affiliation	Leadership	PER	ORG, GPE
<i>ACE</i>			
ORG-Affiliation	Employment	PER	ORG, GPE
ORG-Affiliation	Ownership	PER	ORG
ORG-Affiliation	Founder	PER, ORG	ORG, GPE
ORG-Affiliation	Student-Alum	PER	ORG, Educational
ORG-Affiliation	Sports-Affiliation	PER	ORG
ORG-Affiliation	Investor-Shareholder	PER, GPE	ORG, GPE
ORG-Affiliation	Membership	PER, GPE	ORG, ORG
Agent-Artifact	User-Owner-Inventor-Manufacturer	PER, GPE	ORG, FAC
Gen-Affiliation	Citizen-Resident-Religion-Ethnicity	PER	PER, Group, LOC, GPE, ORG
Gen-Affiliation	Org-Location-Origin	ORG	LOC, GPE

TAC-KBP

Differences with ACE-style Relation Extraction:

- Information is sought for *named* entities, chiefly PERs and ORGs
- The focus is on values, not mentions
- Assessment is more like QA
- Events are handled as uncorrelated slots
- In IE evaluation, there is a skew toward highly attested information; TAC gives full credit for finding a single instance of a correct fill instead of every attestation of that fact
- Somewhat simpler than IE annotation. The assessor must decide if the text supports the posited fact, instead of annotating a document with all the evidenced relations and events for an entity
- Instead of explicitly modeling events, TAC-KBP creates relations that capture events

Relations		Attributes	
per:children per:other_family per:parents per:siblings per:spouse per:employee_or_member_of per:schools_attended perc:city_of_birth per:stateorprovince_of_birth perc:country_of_birth per:city:of_residence per:stateorprovinces_of_residence per:countries_of_residence perc:city_of_death per:stateorprovince_of_death perc:country_of_death	org:shareholders org:founded_by org:top_members_employees org:member_of org:members org:parents org:subsidiaries org:city_of_headquarters org:stateorprovince_of_headquarters org:country_of_headquarters	per:alternate_names per:date_of_birth per:age per:origin per:date_of_death per:cause_of_death per:title per:religion per:charges	org:alternate_names org:political_religious_affiliation org:number_of_employees_members org:date_founded org:date_dissolved org:website

Relation and attributes for PERs and ORGs

FrameNet

FrameNet Motivations:

- Shares same goal of capturing information about events and relations in text
- Designed with text annotation in mind, but unlike ACE/ERE, it prioritizes lexicographic and linguistic completeness over ease of annotation
- Frames tend to be finer grained than ACE/ERE events and more numerous
- Used as a resource for Semantic Role Labeling

<i>Relations</i>			
FrameNet	ACE	ERE	TAC-KBP
Kinship	Personal-Social.Family	Social.Family	per:children per:other_family per:parents per:siblings per:spouse
Being_Employed	ORG-Affiliation.Employment	Affiliation.Employment/Membership	per:employee_or_member_of
Being_Located	Physical.Located	Physical.Located	org:city_of_headquarters org:stateorprovince_of_headquarters org:country_of_headquarters

<i>Events</i>			<i>Attributes</i>	
FrameNet	ACE	ERE	FrameNet	TAC-KBP
Contacting	Phone-Write	Communicate	Being_Named	per:alternate_names
Extradition	Justice-Extradition	Justice-Extradition	Age	per:age
Attack	Conflict-Attack	Conflict-Attack		
Being_Born	Life-Bc-Born	Life-Bc-Born		

Rough mappings between subsets of FrameNet, ACE, ERE, and TAC-KBP

Similarity with TAC-KBP and ACE/ERE:

- Relations and attributes in TAC-KBP and the relation and event types in the ACE/ERE standards can be mapped to FrameNet frames (although often, the mapping is one-to-many)

FrameNet	ACE	ERE
Authorities	Agent-Arg	Agent-Arg
Crime_jurisdiction	Destination-Arg	Destination-Arg
Current_jurisdiction	Origin-Arg	Origin-Arg
Suspect	Person-Arg	Person-Arg
Reason	Crime-Arg	
Time	Time-Arg	
Legal_Basis		
Manner		
Means		
Place		
Purpose		
Deictive		

Mapping between frame elements of Extradition (FrameNet) and arguments of Justice-Extradition (ACE/ERE): A line divides core frame elements (above) from non-core (below).

Differences with TAC-KBP and ACE/ERE:

- FrameNet frames are more fine-grained than ACE/ERE categories
- FrameNet frames are more complex objects than ACE/ERE events and considerably more complex than TAC-KBP relations

Conclusion

The ACE and ERE annotation schemas have closely related goals of identifying similar information across various possible types of documents, though their approaches differ due to separate goals regarding scope and replicability. ERE differs from ACE in collapsing different Type distinctions and in removing annotation features in order to eliminate annotator confusion and to improve consistency, efficiency, and higher inter-annotator agreement. TAC-KBP slot-filling shares some goals with ACE/ERE, but is wholly focused on a set collection of questions (slots to be filled) concerning entities to the extent that there is no explicit modeling of events. At the other extreme, FrameNet seeks to capture the full range of linguistic and lexicographic variation in event representations in text. In general, all events, relations, and attributes that can be represented by ACE/ERE and TAC-KBP can be mapped to FrameNet representations, though adjustments need to be made for granularity of event/relation types and granularity of arguments.

References

- Collin F Baker, Charles J Fillmore, and John B Lowe. 1998. The berkeley framenet project. In *Proceedings of the 17th international conference on Computational linguistics Volume 1*, pages 86–90. Association for Computational Linguistics.
- Collin Baker, Michael Ellsworth, and Katrin Erk. 2007. Semeval2007 task 19: Frame semantic structure extraction. In *Proceedings of the Fourth International Workshop on Semantic Evaluations (SemEval-2007)*, pages 99–104, Prague, Czech Republic, June. Association for Computational Linguistics.
- Dipanjan Das, Nathan Schneider, Desai Chen, and Noah A Smith. 2010. Probabilistic frame-semantic parsing. In *Proceedings of NAACL HLT*, pages 948–956. Association for Computational Linguistics.
- Charles J Fillmore. 1976. Frame semantics and the nature of language. *Annals of the New York Academy of Sciences*. 280(1):20–32.
- Charles Fillmore. 1982. Frame semantics. In *Linguistics in the morning calm*, pages 111–137. Hanshin Publishing Co.
- Daniel Gildea and Daniel Jurafsky. 2002. Automatic labeling of semantic roles. *Computational linguistics*, 28(3):245–288.
- Linguistic Data Consortium. 2005. ACE (automatic content extraction) English annotation guidelines for events. <https://www ldc upenn edu/ collaborations/ pastprojects/ ace>. Version 5.4.3 2005.07.01.
- Linguistic Data Consortium. 2006. ACE (automatic content extraction) English annotation guidelines for entities. <https://www ldc upenn edu/ collaborations/ pastprojects/ ace>. Version 5.6.6 2006.08.01.
- Linguistic Data Consortium. 2008. ACE (automatic content extraction) English annotation guidelines for relations. <https://www ldc upenn edu/ collaborations/ past-projects/ ace>. Version 6.0 2008.01.07.
- Linguistic Data Consortium. 2013a. DEFT ERE annotation guidelines: Entities v1.1. 05.17.2013.
- Linguistic Data Consortium. 2013b. DEFT ERE annotation guidelines: Events v1.1. 05.17.2013.
- Linguistic Data Consortium. 2013c. DEFT ERE annotation guidelines: Relations v1.1. 05.17.2013.
- Ken Litkowski. 2004. Senseval-3 task: Automatic labeling of semantic roles. In Rada Mihalcea and Phil Edmonds, editors, *Senseval-3. Third International Workshop on the Evaluation of Systems for the Semantic Analysis of Text*, pages 9–12, Barcelona, Spain, July. Association for Computational Linguistics.
- Paul McNamee, Hoa Trang Dang, Heather Simpson, Patrick Schone, and Stephanie Strassel. 2010. An evaluation of technologies for knowledge base population. In *Proceedings of LREC*.

Acknowledgement

This material is partially based on research sponsored by the NSF under grant IIS-1249516 and DARPA under agreement number FA8750-13-2-0017 (the DEFT program).