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

Jacqueline Aguilar and **Charley Beller** and **Paul McNamee** and **Benjamin Van Durme**
Human Language Technology Center of Excellence
Johns Hopkins University
Baltimore, MD, USA

Stephanie Strassel and **Zhiyi Song** and **Joe Ellis**
University of Pennsylvania
Linguistic Data Consortium (LDC)
Philadelphia, PA, USA

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.

1 Overview

ACE and ERE are comprehensive annotation standards that aim to consistently 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 across annotators. ERE attempts to achieve this goal by consolidating some of the annotation type distinctions that were found to be the most problematic in ACE, as well as removing some more complex annotation features.

This paper provides an overview of the relationship between these two standards and compares them to the more restricted standard of the TAC-KBP slot-filling task and the more expansive standard of FrameNet. Sections 3 and 4 examine Relations and Events in the ACE/ERE standards, section 5 looks at TAC-KBP slot-filling, and section 6 compares FrameNet to the other standards.

2 ACE and ERE Entity Tagging

Many of the differences in Relations and Events annotation across the ACE and ERE standards stem from differences in entity mention tagging. This is simply because Relation and Event tagging relies on the distinctions established in the entity tagging portion of the annotation process. For example, since ERE collapses the ACE *Facility* and *Location* Types, any ACE Relation or Event that relied on that distinction is revised in ERE. These top-level differences are worth keeping in mind when considering how Events and Relations tagging is approached in ACE and ERE:

- Type Inventory: ACE and ERE share the *Person*, *Organization*, *Geo-Political Entity*, and *Location* Types. ACE has two additional Types: *Vehicle* and *Weapon*. ERE does not account for these Types and collapses the *Facility* and *Location* Types into *Location*. ERE also includes a *Title* Type to address titles, honorifics, roles, and professions (Linguistic Data Consortium, 2006; Linguistic Data Consortium, 2013a).
- Subtype Annotation: ACE further classifies entity mentions by including Subtypes for each determined Type; if the entity does not fit into any Subtype, it is not annotated. ERE annotation does not include any Subtypes.
- Entity Classes: In addition to Subtype, ACE also classifies each entity mention according

Proceedings of the 2nd Workshop on EVENTS: Definition, Detection, Coreference, and Representation, pages 45–53,
Baltimore, Maryland, USA, June 22-27, 2014.

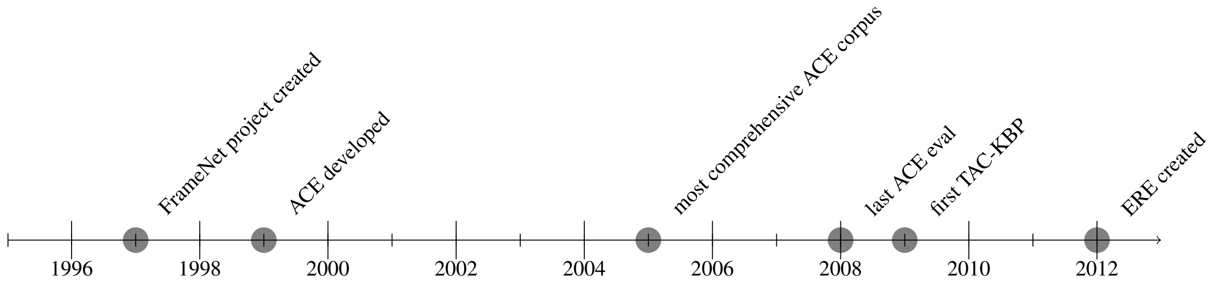

Figure 1: Important Dates for the ACE, ERE, TAC-KBP, and FrameNet Standards

to entity class (Specific, Generic, Attributive, and Underspecified).

- Taggability: ACE tags Attributive, Generic, Specific, and Underspecified entity mentions. ERE only tags Specific entity mentions.
- Extents and Heads: ACE marks the full noun phrase of an entity mention and tags a head word. ERE handles tagging based on the mention level of an entity; in Name mentions (NAM) the name is the extent, in Nominal mentions (NOM) the full noun phrase is the extent, in Pronoun mentions (PRO) the pronoun is the extent.
- Tags: ERE only specifies Type and Mention level (NAM, NOM, PRO). ACE specifies Type, Subtype, Entity Class (Attributive, Generic, Specific, Underspecified), and Mention Level (NAM, NOM, PRO, Headless).

3 Relations in ACE and ERE

In the ACE and ERE annotation models, the goal of the Relations task is to detect and characterize *relations* of the targeted *types* between *entities* (Linguistic Data Consortium, 2008; Linguistic Data Consortium, 2013c). The purpose of this task is to extract a representation of the meaning of the text, not necessarily tied to underlying syntactic or lexical semantic representations. Both models share similar overarching guidelines for determining what is *taggable*. For relations the differences lie in the absence or presence of additional features, *syntactic classes*, as well as differences in *assertion*, *trigger words*, and minor *subtype* variations.

3.1 Similarities in Relations Annotation

In addition to comprising similar Types (both models include *Physical* and *Part.Whole* Types as well as slightly different Types to address *Affiliation* and *Social* relations) used to characterize each relation, ACE and ERE share important similarities concerning their relation-tagging guidelines. These include:

- Limiting relations to only those expressed in a single sentence
- Tagging only for explicit mention
- No 'promoting' or 'nesting' of taggable entities. In the sentence, *Smith went to a hotel in Brazil*, (Smith, hotel) is a taggable *Physical.Located* relation, but (Smith, Brazil) is not. This is because in order to tag this as such, one would have to promote 'Brazil'.
- 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 (*e.g.,* in a *Physical.Located* relation, the Person that is located somewhere will always be assigned to Arg1 and the place in which the person is located will always be assigned to Arg2).
- 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.

3.2 Differences in Assertion, Modality, and Tense

A primary difference between these two annotation models is a result of ERE only annotating asserted events while ACE also includes hypotheticals. ACE accounts for these cases by including two Modality attributes: ASSERTED and OTHER

(Linguistic Data Consortium, 2008). For example, in the sentence, *We are afraid that Al-Qaeda terrorists will be in Baghdad*, ACE would tag this as an OTHER attribute, where OTHER pertains to situations in "some other world defined by counterfactual constraints elsewhere in the context", whereas ERE would simply not tag a relation in this sentence. Additionally, while both ACE and ERE tag past and former relations, ACE goes further to mark the Tense of each relation by means of four attributes: Past, Future, Present and Unspecified.

3.3 Syntactic Classes

ACE further justifies the tagging of each Relation through Syntactic Classes. The primary function of these classes is to serve as a sanity check on taggability and as an additional constraint for tagging. These classes include: Possessive, Preposition, PreMod, Coordination, Formulaic, Participal, Verbal, Relations Expressed by Verbs, and Other. Syntactic classes are not present in ERE relations annotation.

3.4 Triggers

Explicit trigger words do not exist in ACE relation annotation; instead, the model annotates the full syntactic clause that serves as the 'trigger' for the relation. ERE attempts to minimize the annotated span by allowing for the tagging of an optional trigger word, defined as "the smallest extent of text that indicates a relation Type and Subtype" (Linguistic Data Consortium, 2013c). These triggers are not limited to a single word, but can also be composed of a phrase or any extent of the text that indicates a Type/Subtype relation, left to the discretion of the annotator. It is common for prepositions to be triggers, as in *John is **in** Chicago*. However, sometimes no trigger is needed because the syntax of the sentence is such that it indicates a particular relation Type/Subtype without a word to explicitly signal the relation.

3.5 Types and Subtypes of Relations

There are three types of relations that contain varied Subtypes between ERE and ACE. These are the *Physical*, *Part-Whole*, *Social* and *Affiliation* Types. The differences are a result of ERE collapsing ACE Types and Subtypes into more concise, if less specific, Type groups.

Physical Relation Type Differences The main differences in the handling of the physical relations between ACE and ERE are shown in Table 1. ACE only marks Location for PERSON entities (for Arg1). ERE uses Location for PERSON entities being located somewhere as well as for a geographical location being part of another geographical location. Additionally, ACE includes 'Near' as a Subtype. This is used for when an entity is explicitly near another entity, but neither entity is a part of the other or located in/at the other. ERE does not have an equivalent Subtype to account for this physical relation. Instead, ERE includes 'Origin' as a Subtype. This is used to describe the relation between a PER and an ORG. ACE does not have a *Physical* Type equivalent, but it does account for this type of relation within a separate *General Affiliation* Type and 'Citizen-Resident-Religion-Ethnicity' Subtype.

Part-Whole Relation Differences In Table 2, note that ACE has a 'Geographical' Subtype which captures the location of a FAC, LOC, or GPE in or at, or as part of another FAC, LOC, or GPE. Examples of this would be *India controlled the region* or a phrase such as *the Atlanta area*. ERE does not include this type of annotation option. Instead, ERE tags these regional relations as *Physical.Located*. ACE and ERE do share a 'Subsidiary' Subtype which is defined in both models as a "category to capture the ownership, administrative and other hierarchical relationships between ORGs and/or GPEs" (Linguistic Data Consortium, 2008; Linguistic Data Consortium, 2013c).

Social and Affiliation Relation Differences The most evident discrepancy in relation annotation between the two models lies in the *Social* and *Affiliation* Relation Types and Subtypes. For social relations, ACE and ERE have three Subtypes with similar goals (Business, Family, Unspecified/Lasting-Personal) but ERE has an additional 'Membership' Subtype, as shown in Table 3. ACE addresses all 'Membership' relations in its *Affiliation* Type. ERE also includes the 'Social.Role' Subtype in order to address the *TITLE* entity type, which only applies to ERE. However, both models agree that the arguments for each relation must be PERSON entities and that they should not include relationships implied from interaction between two entities (*e.g., President*

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

Table 1: Comparison of Permitted Relation Arguments for the *Physical* Type Distinction in the ERE and ACE Guidelines

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

Table 2: Comparison of Permitted Relation Arguments for the *Part-Whole* Type and Subtype Distinctions in the ERE and ACE Guidelines

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

Table 3: Comparison of Permitted Relation Arguments for the *Social* Type and Subtype Distinctions in the ERE and ACE Guidelines

Relation Type	Relation Subtype	ARG1 Type	ARG2 Type
	ERE		
Affiliation	Employment/Membership	PER, ORG, GPE	ORG, GPE
Affiliation	Leadership	PER	ORG, GPE
	ACE		
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, ORG, GPE	ORG, GPE
ORG-Affiliation	Membership	PER, ORG, GPE	ORG
Agent-Artifact	User-Owner-Inventor-Manufacturer	PER, ORG, GPE	FAC
Gen-Affiliation	Citizen-Resident-Religion-Ethnicity	PER	PER.Group, LOC, GPE, ORG
Gen-Affiliation	Org-Location-Origin	ORG	LOC, GPE

Table 4: Comparison of Permitted Relation Arguments for the *Affiliation* Type and Subtype Distinctions in the ERE and ACE Guidelines

Clinton met with Yasser Arafat last week would not be considered a social relation).

As for the differences in affiliation relations, ACE includes many Subtype possibilities which can more accurately represent affiliation, whereas ERE only observes two Affiliation Subtype options (Table 4).

4 Events in ACE and ERE

Events in both annotation methods are defined as 'specific occurrences', involving 'specific participants' (Linguistic Data Consortium, 2005; Linguistic Data Consortium, 2013b). The primary goal of Event tagging is to detect and characterize events that include tagged entities. The central Event tagging difference between ACE and ERE is the level of specificity present in ACE, whereas ERE tends to collapse tags for a more simplified approach.

4.1 Event Tagging Similarities

Both annotation schemas annotate the same exact Event Types: LIFE, MOVEMENT, TRANSACTION, BUSINESS, CONFLICT, CONTACT, PERSONNEL, and JUSTICE events. Both annotation ontologies also include 33 Subtypes for each Type. Furthermore, both rely on the expression of an occurrence through the use of a 'Trigger'. ACE, however, restricts the trigger to be a single word that most clearly expresses the event occurrence (usually a main verb), while ERE allows for the trigger to be a word or a phrase that instantiates the event (Linguistic Data Consortium, 2005; Linguistic Data Consortium, 2013b). Both methods annotate modifiers when they trigger events as well as anaphors, when they refer to previously mentioned events. Furthermore, when there is any ambiguity about which trigger to select, both methods have similar rules established, such as the Stand-Alone Noun Rule (*In cases where more than one trigger is possible, the noun that can be used by itself to refer to the event will be selected*) and the Stand-Alone Adjective Rule (*Whenever a verb and an adjective are used together to express the occurrence of an Event, the adjective will be chosen as the trigger whenever it can stand-alone to express the resulting state brought about by the Event*). Additionally, both annotation guidelines agree on the following:

- Tagging of Resultative Events (states that result from taggable Events)
- Nominalized Events are tagged as regular events
- Reported Events are **not** tagged
- Implicit events are **not** tagged
- Light verbs are **not** tagged
- Coreferential Events are tagged
- Tagging of multi-part triggers (both parts are tagged only if they are contiguous)

4.2 Event Tagging Differences

One of the more general differences between ERE and ACE Event tagging is the way in which each model addresses *Event Extent*. ACE defines the extent as always being the 'entire sentence within which the Event is described' (Linguistic Data Consortium, 2005). In ERE, the extent is the entire document unless an event is coreferenced (in which case, the extent is defined as the 'span of a document from the first trigger for a particular event to the next trigger for a particular event.' This signifies that the span can cross sentence boundaries). Unlike ACE, ERE does not delve into indicating Polarity, Tense, Genericity, and Modality. ERE simplifies any annotator confusion engendered by these features by simply not tagging negative, future, hypothetical, conditional, uncertain or generic events (although it does tag for past events). While ERE only tags attested Events, ACE allows for irrealis events, and includes attributes for marking them as such: Believed Events; Hypothetical Events; Commanded and Requested Events; Threatened, Proposed and Discussed Events; Desired Events; Promised Events; and Otherwise Unclear Constructions. Additionally both ERE and ACE tag Event arguments as long as the arguments occur within the event mention extent (another way of saying that a taggable Event argument will occur in the same sentence as the trigger word for its Event). However, ERE and ACE have a diverging approach to argument tagging:

- ERE is limited to pre-specified arguments for each event and relation subtype. The possible arguments for ACE are: Event participants (limited to pre-specified roles for each event type); Event-specific attributes that are associated with a particular event type (*e.g.,* the victim of an attack); and General event attributes that can apply to most or all event types (*e.g.,* time, place).

- ACE tags arguments regardless of modal certainty of their involvement in the event. 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. This is because ACE handles entity annotation slightly differently than ERE does; ACE marks the full noun phrase with a head word for entity mention, and ERE treats mentions differently based on their syntactic features (for named or pronominal entity mentions the name or pronominal itself is marked, whereas for nominal mentions the full noun phrase is marked).

Event Type and Subtype Differences Both annotation methods have almost identical Event Type and Subtype categories. The only differences between both are present in the Contact and Movement Event Types.

A minor distinction in Subtype exists as a result of the types of entities that can be transported within the Movement Type category. In ACE, ARTIFACT entities (WEAPON or VEHICLE) as well as PERSON entities can be transported, whereas in ERE, only PERSON entities can be transported. The difference between the Phone-Write and Communicate Subtypes merely lies in the definition. Both Subtypes are the default Subtype to cover all Contact events where a 'face-to-face' meeting between sender and receiver is not explicitly stated. In ACE, this contact is limited to written or telephone communication where at least two parties are specified to make this event subtype less open-ended. In ERE, this requirement is simply widened to comprise electronic communication as well, explicitly including those via internet channels (*e.g.,* Skype).

5 TAC-KBP

After the final ACE evaluation in 2008 there was interest in the community to form an evaluation explicitly focused on knowledge bases (KBs) created from the output of extraction systems. NIST had recently started the Text Analysis Conference series for related NLP tasks such as Recognizing Textual Entailment, Summarization, and Question Answering. In 2009 the first Knowledge Base Population track (TAC-KBP) was held featuring two initial tasks: (a) Entity Linking — linking entities to KB entities, and (b) Slot Filling — adding information to entity profiles that is missing from the KB (McNamee et al., 2010). Due to its generous license and large scale, a snapshot of English Wikipedia from late 2008 has been used as the reference KB in the TAC-KBP evaluations.

5.1 Slot Filling Overview

Unlike ACE and ERE, Slot Filling does not have as its primary goal the annotation of text. Rather, the aim is to identify knowledge nuggets about a focal named entity using a fixed inventory of relations and attributes. For example, given a focal entity such as former Ukrainian prime minister Yulia Tymoshenko, the task is to identify attributes such as schools she attended, occupations, and immediate family members. This is the same sort of information commonly listed about prominent people in Wikipedia Infoboxes and in derivative databases such as FreeBase and DBpedia.

Consequently, Slot Filling is somewhat of a hybrid between relation extraction and question answering — slot fills can be considered as the correct responses to a fixed set of questions. The relations and attributes used in the 2013 task are presented in Table 5.

5.2 Differences with ACE-style relation extraction

Slot Filling in TAC-KBP differs from extraction in ACE and ERE in several significant ways:

- information is sought for *named* entities, chiefly PERs and ORGs;
- the focus is on values not mentions;
- assessment is more like QA; and,
- events are handled as uncorrelated slots

In traditional IE evaluation, there was an implicit skew towards highly attested information such as $leader(Bush, US)$, or $capital(Paris, France)$. In contrast, TAC-KBP gives full credit for finding a single instance of a correct fill instead of every attestation of that fact.

Slot Filling assessment is somewhat simpler than IE annotation. The assessor must decide if provenance text is supportive of a posited fact about the focal entity instead of annotating a document with all evidenced relations and events for any entity. For clarity and to increase assessor agreement, guidelines have been developed to justify when a posited relation is deemed adequately supported from text. Additionally, the problem of

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

Table 5: Relation and attributes for PERs and ORGs.

slot value equivalence becomes an issue - a system should be penalized for redundantly asserting that a person has four children named Tim, Beth, Timothy, and Elizabeth, or that a person is both a cardiologist and a doctor.

Rather than explicitly modeling events, TAC-KBP created relations that capture events, more in line with the notion of Infobox filling or question answering (McNamee et al., 2010). For example, instead of a criminal event, there is a slot fill for charges brought against an entity. Instead of a founding event, there are slots like org:founded_by (who) and org:date_founded (when). Thus a statement that “Jobs is the founder and CEO of Apple” is every bit as useful for the org:founded_by relation as “Jobs founded Apple in 1976.” even though the date is not included in the former sentence.

5.3 Additional tasks

Starting in 2012 TAC-KBP introduced the “Cold Start” task, which is to literally produce a KB based on the Slot Filling schema. To date, Cold Start KBs have been built from collections of O(50,000) documents, and due to their large size, they are assessed by sampling. There is also an event argument detection evaluation in KBP planned for 2014.

Other TAC-KBP tasks have been introduced including determining the timeline when dynamic slot fills are valid (*e.g.,* CEO of Microsoft), and targeted sentiment.

6 FrameNet

The FrameNet project has rather different motivations than either ACE/ERE or TAC-KBP, but shares with them a goal of capturing information about events and relations in text. FrameNet stems from Charles Fillmore’s linguistic and lexicographic theory of Frame Semantics (Fillmore, 1976; Fillmore, 1982). Frames are descriptions of event (or state) types and contain information about event participants (*frame elements*), information as to how event types relate to each other (*frame relations*), and information about which words or multi-word expressions can trigger a given frame (*lexical units*).

FrameNet is designed with text annotation in mind, but unlike ACE/ERE it prioritizes lexicographic and linguistic completeness over ease of annotation. As a result Frames tend to be much finer grained than ACE/ERE events, and are more numerous by an order of magnitude. The Berkeley FrameNet Project (Baker et al., 1998) was developed as a machine readable database of distinct frames and lexical units (words and multi-word constructions) that were known to trigger specific frames.[1] FrameNet 1.5 includes 1020 identified frames and 11830 lexical units.

One of the most widespread uses of FrameNet has been as a resource for Semantic Role Labeling (SRL) (Gildea and Jurafsky, 2002). FrameNet related SRL was promoted as a task by the SENSEVAL-3 workshop (Litkowski, 2004), and the SemEval-2007 workshop (Baker et al., 2007). (Das et al., 2010) is a current system for automatic FrameNet annotation.

The relation and attribute types of TAC-KBP and the relation and event types in the ACE/ERE standards can be mapped to FrameNet frames. The mapping is complicated by two factors. The first is that FrameNet frames are generally more fine-grained than the ACE/ERE categories. As a result the mapping is sometimes one-to-many. For example, the ERE relation **Af-**

[1]This database is accessible via webpage (`https://framenet.icsi.berkeley.edu/fndrupal/`) and as a collection of XML files by request.

Relations

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

Events

FrameNet	ACE	ERE
Contacting	Phone-Write	Communicate
Extradition	Justice-Extradition	Justice-Extradition
Attack	Conflict-Attack	Conflict-Attack
Being_Born	Life-Be_Born	Life-Be_Born

Attributes

FrameNet	TAC-KBP
Being_Named	per:alternate_names
Age	per:age

Table 6: Rough mappings between subsets of FrameNet, ACE, ERE, and TAC-KBP

filiation.Employment/Membership covers both the **Being_Employed** frame and the **Membership** frame. At the same time, while TAC-KBP has only a handful of relations relative to FrameNet, some of these relations are more fine-grained than the analogous frames or ACE/ERE relations. For example, the frame **Kinship**, which maps to the single ERE relation **Social.Family**, maps to five TAC-KBP relations, and the **Being_Located**, which maps to the ACE/ERE relation **Being.Located**, maps to three TAC-KBP relations. Rough mappings from a selection of relations, events, and attributes are given in Table 6.

The second complication arises from the fact that FrameNet frames are more complex objects than ERE/ACE events, and considerably more complex than TAC-KBP relations. Rather than the two entities related via a TAC-KBP or ACE/ERE relation, some frames have upwards of 20 frame elements. Table 7 shows in detail the mapping between frame elements in the Extradition frame and ACE's and ERE's Justice-Extradition events. The "core" frame elements map exactly to the ERE event, the remaining two arguments in the ACE event map to two non-core frame elements, and the frame includes several more non-core elements with no analogue in either ACE or ERE standards.

FrameNet	ACE	ERE
Authorities Crime_jursidiction Current_jursidiction Suspect	Agent-Arg Destination-Arg Origin-Arg Person-Arg	Agent-Arg Destination-Arg Origin-Arg Person-Arg
Reason Time Legal_Basis Manner Means Place Purpose Depictive	Crime-Arg Time-Arg	

Table 7: 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).

7 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-KPB 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 standards can be mapped to FrameNet representations, though adjustments need to be made for granularity of event/relation types and granularity of arguments.

Acknowledgements

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).

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. Semeval-2007 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.

George Doddington, Alexis Mitchell, Mark Przybocki, Lancec Ramshaw, Stephanie Strassel, and Ralph Weischedel. 2004. The automatic content extraction (ace) program- tasks, data, and evaluation. In *Proceedings of LREC 2004: Fourth International Conference on Language Resources and Evaluation*, Lisbon, May 24-30.

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/past-projects/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/past-projects/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*.