Natural Language Engineering **15** (3): 415–453. © 2009 Cambridge University Press
doi:10.1017/S1351324909005117 Printed in the United Kingdom

The FrameNet model and its applications†

BIRTE LÖNNEKER-RODMAN and
COLLIN F. BAKER

International Computer Science Institute, 1947 Center Street, Suite 600, Berkeley, CA 94704, USA
e-mail: loenneke@icsi.berkeley.edu, collinb@icsi.berkeley.edu

(*Received 28 March 2008; revised 25 December 2008; accepted 1 April 2009*)

Abstract

The FrameNet database comprises an English lexicon, organized in terms of semantic frames. Frames describe situations or entities, along with their participants and props, termed frame elements. The frames are organized in an ontology-like network. For the lexical units, corpus annotations illustrate which frame elements are typically realized, and how they behave syntactically. Texts where all content words are annotated with FrameNet information offer a detailed, structured semantic representation with a variety of uses in Natural Language Processing applications, in particular in retrieving and meaningfully organizing texts written by humans, or in making human–computer interaction more natural. Also, the FrameNet English lexicon can be replaced by lexical data from other languages, while maintaining frame information, so the model is attractive for cross-lingual resources and applications. Manual annotation produced by FrameNet and similar projects for other languages is used to train automatic frame semantic annotation systems, which add rich semantic information to any type of text, and are important components for more sophisticated semantic processing applications.

1 Motivation

While some natural language processing (NLP) applications, including spell checkers and stemmers, have matured into components of commercial products, the development of systems that incorporate a semantic layer is still an active research area. For example, given the evergrowing amount of textual information available in electronic format, organization by meaningful categories has become a central goal in information retrieval and related areas. The results of a key word search might thus be presented along with structured auxiliary data to facilitate the exploration

† The first author was supported by a fellowship within the postdoc programme of the German Academic Exchange Service (DAAD), September 2006–November 2008. We also gratefully acknowledge the support of the NSF (1997–2000, 2000–2003, and 2006–present), DARPA (2003–2004), and the AQUAINT project (2004–2005) for work on FrameNet. The opinions, findings, and conclusions expressed in this article are those of the author(s) and do not necessarily reflect the views of the National Science Foundation (NSF). We are also grateful for the helpful suggestions of the anonymous reviewers and of our colleagues at International Computer Science Institute and at FrameNet-related projects elsewere; any errors that remain are those of the authors.

https://doi.org/10.1017/S1351324909005117 Published online by Cambridge University Press

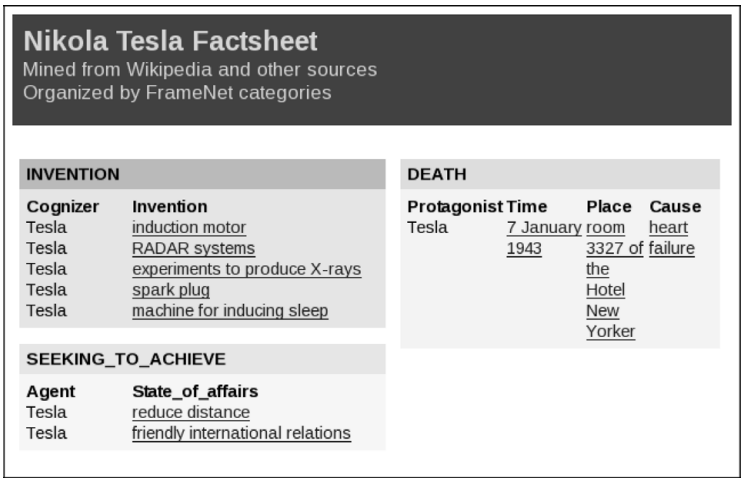


Fig. 1. Organizing information by meaningful categories.

of the result set (see e.g., Meredith and Pieper 2006). New systems, such as those in the recent 'faceted search' paradigm, integrate information retrieval approaches with information extraction and summarization, and this is supported by semantic representations derived from natural language texts and possibly other media.

To structure data mined from one or multiple documents, a system might produce several tables, as illustrated in Figure 1 for Nikola Tesla, the inventor of the induction motor. In the figure, each table represents a situation or an event, relevant to the result set for Tesla: `Invention`, `Seeking_to_achieve`, `Death`. Within the tables, each line contains information on the role Tesla played in a particular instance of these situations or events. So, Tesla is known as the Cognizer in five different invention events, each with a different value of the Invention, 'induction motor' being one of them. There can be more than two roles in an event; for example, the `Death` table lists information on Protagonist, Time, Place, and Cause.

The situations and events in the example have been described as **semantic frames** in the FrameNet database. FrameNet contains a structured network of more than 945 frames, including `Invention`, `Seeking_to_achieve`, and `Death`. For each frame, it also describes the relevant participants and roles (such as Agent, State_of_affairs). Since FrameNet was conceived as a *lexical database*, it offers rich linguistic information, useful for semantically describing and processing English texts. For example, invention situations can be referred to by the words *invention*, *invent*, *conceive*, *devise*, and others, listed in the `Invention` frame. FrameNet contains more than 11,500 such pairings between a word and a frame. The lines in the invention table in Figure 1 have been derived from sentences containing instances of several of the `Invention` words, found in the Wikipedia entry on Tesla. Finally, FrameNet offers approximately 170,000 English corpus sentences with marked-up frame-evoking words and additional manual annotation. The annotations map parts of a sentence (for example, 'Nikola Tesla,' 'induction motor,' or 'in room 3327 of the Hotel New Yorker') onto the correct roles in the relevant frame.

https://doi.org/10.1017/S1351324909005117 Published online by Cambridge University Press

FrameNet data could thus be used to automatically mine information from text, and to structure it by situations, similar to the manually created example in Figure 1. Since FrameNet frames are conceptual entities, linked not only to English words, but also to lexicons for other languages, such a system can be multilingual. Information could be derived from collections of texts in different languages, and even be summarized across languages by populating the tables with data from different collections. FrameNet data has already been used to train systems that automatically produce FrameNet-like annotation on previously unseen texts, known as automatic semantic role labeling (ASRL). ASRL is useful for interpreting the meaning of open-ended texts, such as in the information retrieval and extraction system sketched above, or for question answering (Narayanan and Harabagiu 2004; Sinha and Narayanan 2005; Shen and Lapata 2007), paraphrasing (Ellsworth and Janin 2007), textual entailment recognition and reasoning (Burchardt and Frank 2006; Scheffczyk, Baker and Narayanan 2006b), and, potentially, machine translation systems.

To facilitate wider and more effective use of the FrameNet resource in these and other application scenarios, this paper presents the project (Section 2) and provides an overview of the rich semantic data model (Section 3). Then, we lay out the implications of the model for two active application areas: Section 4 deals with multilingual resource creation, and Section 5 with automatic semantic role labeling. We conclude with an outlook on FrameNet extensions (Section 6).

2 The FrameNet project

The FrameNet approach to resource creation is based on the linguistic theory of frame semantics (2.1). This section also introduces the frame creation and annotation process (2.2) and describes the data releases (2.3).

2.1 Frame semantics

The theory of frame semantics (Fillmore 1976; Fillmore 1985) states that the meanings of many words are best understood in the context of a type of event or situation and the participants and 'props' involved in it. In frame semantics, a **lexical unit** (LU) (Cruse 1986) has a *form pole*, consisting of a lemma and the associated word forms, and a *meaning pole*, partially represented by a connection to a semantic frame. An LU is thus equivalent to a word sense.

Consider, for example, the `Leadership` frame. The central roles in it (called **frame elements** (FEs)) are the Leader, the Role (the title or position of the Leader), and the entity which is led, which is further differentiated into the Governed ('the group of people whose actions or beliefs the Leader directs in the context of a joint activity . . . ') or the Jurisdiction ('the political locale in which the Leader has control . . . '). The concept represented by the `Leadership` frame helps us explain the meaning of sentences like those shown in (1)–(4).

(1) [The Moors$_{Leader}$] had **ruled** [Granada$_{Jurisdiction}$] [for over seven hundred years $_{Duration}$] . . .

https://doi.org/10.1017/S1351324909005117 Published online by Cambridge University Press

(2) He has written to [Marstons $_{Governed}$] [managing **director** $_{Role}$] [Michael Hurdle $_{Leader}$], asking for a meeting . . .

(3) [A well-known Radio Forth RFM presenter $_{Leader}$] will **head** [the campaign $_{Governed}$].

(4) The [Low Church $_{Governed}$] [**bishops** $_{Leader}$] proved to be valuable Parliamentary allies to the Whigs.

The words shown in boldface, *ruled*, *director*, *head*, and *bishops*, all **evoke** the `Leadership` frame. Note that *head* and *rule* are verbs here, and the other two are nouns, but all can be frame-evoking expressions, or, in other words, all are LUs in the `Leadership` frame.

In addition to recognizing the sentences as instances of the `Leadership` frame, we can also recognize the other pieces of the sentence which fill in the FEs. In (1)–(4), FEs are shown by square brackets with subscript FE names. The same set of FEs are used for all LUs in a frame, whether they are nouns, verbs or other parts of speech, and regardless of the grammatical relation between the LU and the phrase filling the FE.

2.2 The frame creation and annotation process

FrameNet was initially a lexicographic project, engaged in building a lexicon with uniquely detailed information on the syntax and semantics of LUs. More recently, FrameNet has also been annotating continuous texts.

The lexicographic work at FrameNet is conducted on a frame-by-frame basis, rather than word by word as with most dictionaries. All of the LUs in one semantic frame will be defined (and in most cases, exemplified through annotation) before moving on the next frame. Consequently, if one sense of a polysemous lexeme is covered by FrameNet, this does not necessarily entail that frames for *all* of its senses have been created.

The stages in the lexicographic process to create frames and annotation are as follows:

(1) *Select a semantic domain and outline the frames involved, through a combination of introspection and careful study of corpus data.* Initially, FrameNet chose ten quite different semantic domains, to test if the Frame Semantic concepts were applicable to all of them. Later, frames were added that are semantically related to the original frames, or contain other senses (LUs) of the lexemes in the original frames.

(2) *Define the frames and their FEs and prepare a list of LUs, with a brief definition of each.* Again, this requires not only a basic understanding of the situation type, such as with the `Leadership` frame above, but also careful study of corpus examples to note the types of roles and props that commonly appear around the frame-evoking expression under study. Decisions about which LUs to include in a frame also impel the researcher to recognize other senses of the same lexeme and to consider whether they belong in already defined

https://doi.org/10.1017/S1351324909005117 Published online by Cambridge University Press

frames, or require new frames. Since frame creation starts from the linguistic data, rather than from some pre-existing ontology, the process is 'bottom-up' and data driven. More abstract frames and relations between frames (see Section 3.4 below) are defined only as needed to capture linguistic and conceptual generalizations.

(3) *For each LU, find the principal syntactic patterns and extract examples of each from a large corpus.*[1] The extraction process is governed by hand-written rules, which match against chunked sentences from the corpus, so that sentences matching each pattern are extracted. Each sentence is annotated separately, without regard to the context, as in (1)–(4), above.

(4) *Annotate enough examples to provide evidence for all the lexicographically relevant syntactic realizations of each FE.* By 'lexicographically relevant,' we mean patterns that vary at the level of the frame or the LU. Alternations that affect a very large proportion of the lexicon, such as the passive alternation for transitive verbs, are not lexicographically relevant, and therefore not necessarily exemplified.

(5) *Run report software to create a viewable version of the annotated LU.* There are two views, the 'annotation view' and the 'lexical entry view.' The annotation view shows all the annotated sentences for the LU, in a format similar to (1)–(4), above. A lexical entry report shows all the **valence patterns** of the LU. For each FE and FE combination, it displays syntactic information in terms of the phrase type of the FE (for example, noun phrase (NP) or prepositional phrase (PP)) and its grammatical function, such as Object or Dependent. For example, in (1) above, the Governed FE is realized as an NP ('Granada') with the grammatical function Object. The same FE might be expressed as a PP Dependent (e.g., 'over Saragossa') in other sentences. For each such syntactic pattern, the lexical entry report shows the number of times it has been annotated for the LU under study.

There has always been an expectation that the frames and FEs being developed would be used to 'densely' annotate texts (Fillmore and Baker 2001). Since 2004, FrameNet has also been annotating *continuous texts*, as a kind of benchmark for deep semantic annotation. In this **full-text annotation**, all the content words must be annotated (Baker 2008). New LUs in existing frames and (less often) new frames are created as dictated by the content; therefore, the number of frames and LUs increases more rapidly than the average number of annotations per LU. Full-text annotation also means that annotators cannot choose clear sentences as examples, but must deal with whatever syntactic complexities they encounter. Usually at least three or four frame-evoking expressions are annotated per sentence, and often there are more.

The first group of continuous texts were from PropBank, an annotation project discussed in Section 5.1.2, below. The texts, on various topics, were annotated by FrameNet in order to provide a direct comparison with the PropBank annotation

[1] FrameNet has taken most of its lexicographic examples from the British National Corpus.

https://doi.org/10.1017/S1351324909005117 Published online by Cambridge University Press

Table 1. *FrameNet data releases*

Release	LUs	Frames	FEs	Frame relations	Annotations
0.7	4,281	182	–	–	84,851
1.0	6,904	384	3,059	–	127,293
1.1	7,859	487	3,936	352	132,968
1.2	8,755	607	4,908	485	133,846
1.3	10,195	795	7,124	1,723	*Lexicographic:* 139,439
					Full-text: 11,671
					Total: 151,110

style. The second group of texts were 'country reports' from the Nuclear Threat Initiative (NTI) website (`http://www.nti.org`), giving the history of various countries with regard to their acquisition of weapons of mass destruction. The third group are texts from the American National Corpus (ANC) (Ide, Reppen and Suderman 2002), including articles from the online magazine *Slate* and Berlitz travel guides.[2]

2.3 Data releases

The primary product of the FrameNet project are the descriptions of the frames and FEs and the sentences annotated on the basis of these descriptions. A major activity has been collecting this data, putting it into a format suitable for use by others, and releasing copies, which are freely available from the FrameNet web site for research purposes, `http://framenet.icsi.berkeley.edu`. There have been five data releases since 2000, as shown in Table 1; the most recent has been downloaded more than 500 times.

Beginning with release 1.3, each release contains version difference files to help track changes and bring applications into alignment with new frames and LUs. Since release 1.2, efforts have been divided between lexicographic and full-text annotation, as shown in the table. The most recent data is always browsable on the FrameNet public website, updated monthly.

3 The FrameNet model

The previous sections have introduced the FrameNet notions of *LUs*, *frames*, and *FEs*. We have has also shown that FrameNet produces manual *annotation* of instances of LUs in terms of FEs, phrase types, and grammatical functions, thereby documenting semantic and syntactic *valence patterns*. All of this data is stored in a relational database that has three main parts: lexicon, annotation, and frames. The lexicon part covers *form information* pertaining to LUs, i.e., lemma,

[2] The PropBank texts were made available courtesy of Martha Palmer (U Colorado, Boulder) and the Linguistic Data Consortium at University of Pennsylvania. The NTI texts were made available courtesy of Gary Ackerman, of the Center for Nonproliferation Studies in Monterey, CA. The work on the ANC texts is being performed under a subcontract to develop a multiply-annotated subcorpus of the ANC with NSF support (CRI #0708952).

https://doi.org/10.1017/S1351324909005117 Published online by Cambridge University Press

part of speech, word forms, and word order for multiwords. The annotation part represents sentences and annotation with respect to a particular target LU. Finally, the frame part contains information on frames and FEs, much richer than what we have described in the introduction so far. In particular, FrameNet frames are linked to each other by different types of *relations*, resulting in a large frame-semantic knowledge base, similar to an ontology and useful for reasoning (Scheffczyk *et al.* 2006b). Importantly, the relations are directly applicable to NLP, since each LU makes reference to exactly one frame, thereby relating form information from the lexicon to conceptual knowledge in the frame hierarchy.

In the remainder of this section, a detailed model of the frame part of the FrameNet database will be presented. The model comprises the following components, discussed in the sections whose numbers are given in parentheses below:

- *Frames* (3.1);
- *FEs* (3.2);
- Frame-internal *relations between FEs* (3.3); and
- Directed *relations between frames* and their FEs (3.4).

3.1 Frame

We have seen in 2.1, above, that a semantic frame is a conceptual structure describing a particular type of situation along with its participants and props. The latter are referred to as FEs. In terms of the FrameNet database structure, the attributes of a frame are: its ID, its name (3.1.1), definition (3.1.2), and semantic type (3.1.3). Each frame has one or more FEs, discussed in Section 3.2.

3.1.1 Frame name

A frame name, such as `Leadership`, mentioned in Section 2.1, is a unique mnemonic that stands for the frame, a convenient 'handle' used in addition to its numeric ID. Some parts of frame names are conventionalized. For instance, the ending 'scenario' is part of the name of certain general frames at higher levels in the FrameNet frame hierarchy (see Section 3.4, below).

3.1.2 Frame definition

A frame definition describes the situation or object represented by the frame, often with example sentences. Definitions are free English text, and thus not subject to formal restrictions. Usually, the definition mentions at least the most central FEs and their relation to each other. Example (5) contains the beginning of the `Scrutiny` frame definition, with FE names being capitalized. The frame hosts LUs such as *examination.n*, *examine.v*, *rummage.v*, *search.n*, *search.v*, and many more.

(5) This frame concerns a Cognizer (a person or other intelligent being) paying close attention to something, the Ground, in order to discover and note its salient characteristics. The Cognizer may be interested in a particular

https://doi.org/10.1017/S1351324909005117 Published online by Cambridge University Press

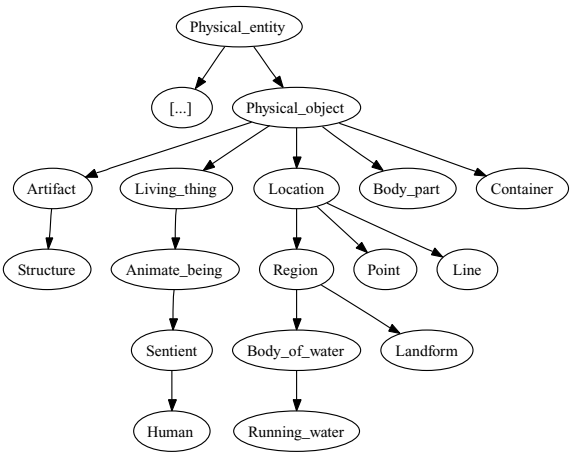


Fig. 2. FrameNet ontological semantic types (selection).

characteristic or entity, the Phenomenon, that belongs to the Ground or is contained in the Ground (or ensuring that such a property or entity is not present).

As mentioned above, FrameNet describes the syntactic realization patterns of FEs by summarizing annotations of particular LUs into lexical entry reports. Still, (primarily to guide the annotators) some frame definitions also explicitly summarize the syntactic behavior of LUs in the frame. For example, several words in the `Scrutiny` frame allow for alternative ways of syntactically realizing two FEs (Ground and Phenomenon), as in (6) *versus* (7).

(6) . . . [police and coastguards$_{Cognizer}$] were **searching** [the coastline$_{Ground}$] [for a man$_{Phenomenon}$] . . . .

(7) . . . [she$_{Cognizer}$] began to **search** [for a piece of land$_{Phenomenon}$] [in the Cotherstone area$_{Ground}$] . . . .

To make the annotators aware of this possibility, the `Scrutiny` definition continues as follows:

(8) Some words in this frame allow alternate expressions of the Ground and the Phenomenon.

If, on the other hand, syntactic alternations of this kind coincide with a difference in semantic frames, the definition lays out how the frames differ.

3.1.3 Semantic type of frames

Semantic type labels can be attached to a frame, to further characterize it. They fall into two categories: ontological semantic types and framal types. *Ontological semantic types* are organized into a hierarchy (see Figure 2). When attached to a frame, a semantic type indicates that all LUs evoking the frame should have this (or a more specific) semantic type. For example, the `Clothing` frame has the semantic

https://doi.org/10.1017/S1351324909005117 Published online by Cambridge University Press

type Artifact. So, all the LUs in it (*boot.n*, *cape.n*, *dress.n*, and many others) designate artifacts.

Framal types, however, are meta information about the frame, not the LUs. In particular, the framal type "Non-lexical" is used for frames that do not contain LUs. Non-lexical frames must appear as parent frames within the frame Inheritance relation (see Section 3.4). Their descendant frames will have LUs.

Frames are not required to have a semantic type. One hundred and thirty one frames in the current FrameNet database have been provided with a semantic type label: 104 have one or more framal types, and 29 have been assigned one or more ontological types.

3.2 Frame element

A FE is named and defined with respect to the general conceptual situation described by the frame. All LUs of a frame must have the same FEs.

3.2.1 Frame element name

A FE name is a mnemonic handle standing for the definition of a FE, besides its unique ID. In contrast to frame names, FE *names* are not unique. Rather, FEs in closely related frames are often named the same. This will be further discussed in Section 3.4 on frame inheritance.

3.2.2 Frame element definition

An FE definition describes, in free text, the meaning of the FE with respect to the frame it belongs to. Consider the Goal FE of the `Filling` frame. Against the background of the frame definition (9), the FE is further defined separately (10).

(9) These are words relating to filling containers and covering areas with some thing, things or substance, the Theme. The area or container can appear as the direct object with all these verbs, and is designated Goal because it is the goal of motion of the Theme. Corresponding to its nuclear argument status, it is also affected in some crucial way, unlike goals in other frames.

(10) The Goal is the area or container being filled.

To guide the annotators, FE definitions might generalize syntactic realization patterns for it, across LUs. The Goal definition from the `Filling` frame continues as in (11), describing the syntactic realization patterns illustrated in (12), where 'his suitcases' is a noun phrase (NP) with the grammatical function Object.

(11) Goal is generally the NP Object in this frame.

(12) [He$_{Agent}$] **packed** [his suitcases$_{Goal}$] [with clothes and books$_{Theme}$].

FE definitions do not have a fixed internal structure, but they may contain example sentences, different from those found in the annotation database. If they do, then these examples and their components are marked up in XML; tools for viewing FrameNet data convert the markup into presentational devices.

https://doi.org/10.1017/S1351324909005117 Published online by Cambridge University Press

A Communicator conveys a Message to an Addressee;
the Topic and Mediam of the communication also may be expressed.
This frame includes no specification of the method of communication
(speech, writing, gesture, etc.).

Fig. 3. `Communication` frame definition.

3.2.3 Semantic type of frame elements

In many cases, semantic type restrictions can be specified for the head word of a phrase that realizes or 'fills' the role of a particular FE. Consider (13), where *catechist* is the head word of a noun phrase filling the Communicator FE in the `Communication` frame, evoked by a word form of *communicate.v*.

(13) [The catechist$_{Communicator}$] **communicated** with Todd by way of simple communication and a flannel board.

In terms of the FrameNet ontological semantic type hierarchy presented in Figure 2, above, *catechist* can be described as Human. However, other types of Communicators are acceptable as well, including *horses* (14) and *aliens*.

(14) [Horses$_{Communicator}$] have a quick eye for slight muscular movements and changes in the posture of their companions and perhaps **communicate** many of their feelings by signalling in some such subliminal manner.

In fact, the situation described by `Communication` (Figure 3) is not restricted to human participants. So, the more general semantic type Sentient is attached to the Communicator FE. Therefore, in all occurrences of LUs evoking `Communication` (e.g., *communicate.v*), the Communicator should be sentient. Examples (13) and (14) meet this requirement. Exceptions can be found in figures of speech, such as metonymy in (15), where *letter* stands for the writer of the letter.

(15) . . . [her letter$_{Communicator}$] **communicated** nothing of her pleasure and love. . .

All semantic types on FEs are ontological. To facilitate reasoning in conjunction with external resources, the FrameNet semantic types have been mapped onto the Suggested Upper Merged Ontology (Scheffczyk, Pease and Ellsworth 2006a); this mapping is part of release 1.3. FEs are not required to have a semantic type and rarely have more than one. Currently, 28 distinct types are attached to 3,841 FEs (out of a total of 8,550 FEs).

3.2.4 Coreness status

Each FE must have a coreness status defined. This is meta-information pertaining to the *conceptual centrality* of the FE with respect to the frame. There are three coreness statuses: core, peripheral, and extra-thematic; see (Ruppenhofer *et al.* 2006)[24; 26–28; 127].

Core FEs are defined as instantiating 'a conceptually necessary component of a frame, while making the frame unique and different from other frames' (Ruppenhofer *et al.* 2006)[26]. To make the notion of conceptual necessity practically applicable,

https://doi.org/10.1017/S1351324909005117 Published online by Cambridge University Press

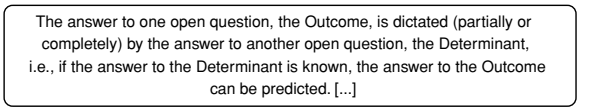


Fig. 4. `Contingency` frame definition.

some formal properties are considered as well: every instance of an LU must have an overt realization of each core FE or, if omitted, the core FE receives a definite interpretation. The semantics of core FEs are unpredictable from their form, in particular from any marking prepositions. For example, in (16), an example of the `Contingency` frame evoked by *depend (on)*, the semantics of the FE introduced by *on* are not predictable from the use of *on* itself.

(16) [Spotting a good school$_{Outcome}$] will **depend** [on using information and common sense$_{Determinant}$].

Consequently, the FE Determinant in the `Contingency` frame (Figure 4), evoked by *depend*, is core. Yet, in many other frames, the preposition *on* marks FEs that have the same semantics across frames, and this FE is called Place (or Time). Consider (17) in the `Cause_harm` frame and (18), in `Social_event`.

(17) . . . [a man$_{Victim}$] was **stabbed** [in the face$_{Body_part}$] [on a train platform$_{Place}$] [yesterday$_{Time}$].

(18) Drugs seized at [this weekend's$_{Time}$] **rave** [on the outskirts of Northampton$_{Place}$] include ecstasy, cannabis, ...

Peripheral FEs such as Place and Time, illustrated in the examples above, can be instantiated in any semantically appropriate frame. The FEs Manner, Means, and Degree also typically fall into this category. Sometimes, these notions are necessary to define the frame and to distinguish it from others; in those cases, the necessary FEs are core. As far as spatial notions are concerned, peripheral FEs are by convention called Place, and core FEs are usually given a different name, such as Location, Goal, etc. (The semantic distinction between core and peripheral FEs can be compared to, and usually coincides with, the division of syntactic complements into arguments and adjuncts, but since FrameNet is agnostic with respect to theories of syntax, it is not concerned with this distinction.)

Finally, **extra-thematic** FEs introduce new events or states of affairs, distinct from the main reported event, as a backdrop against which the main event is situated. Since this additional event or state is actually like evoking a different frame, extra-thematic FEs do not *conceptually* belong to the frames they appear in; they are *not frame-specific*. For example, there are currently 211 frames with an extra-thematic FE called Circumstances, including `Cause_harm` (19) and `Contingency` (20).

(19) [He$_{Victim}$] would be **poisoned** [if he accepted food or sweets from them$_{Circumstances}$].

(20) [In a close relationship$_{Circumstances}$], [the anger of one$_{Outcome}$] is a **function** [of both$_{Determinant}$].

https://doi.org/10.1017/S1351324909005117 Published online by Cambridge University Press

In English, extra-thematic FEs are usually expressed by prepositional phrases or separate clauses. Again, a FE that is extra-thematic in one frame might be peripheral or core in another frame.

These explanations should convey the general idea of coreness status as an information type within the FrameNet model. Yet, practical work on frame creation and LU annotation often brings up additional issues. One of them is related to the omission of core FEs in particular sentences. The omission might actually be licensed by the context or by a grammatical construction. For example, in (21), the `Chatting` frame is evoked by *spoke*, but the core FE Interlocutor_2 is omitted.

(21) She broke into passionate sobbing as [she$_{Interlocutor_1}$] **spoke**.

Interlocutor_2, the semantically less prominent participant in the conversation, must be understood as being given by the larger linguistic or situational context of (21). The phenomenon of core FE omission is termed **null instantiation (NI)** in FrameNet (Fillmore, Johnson and Petruck 2003)[245–246]. Annotators mark null-instantiated FEs on the sentence, but without an exact position. A core FE may also be incorporated into the target. Consider the `Clothing` frame. It has one core FE, Garment, identifying the clothing worn. So, in (22), the target LU *jumper* is annotated as Garment.

(22) Put [your$_{Wearer}$] [**jumper**$_{Garment}$] on.

The particular type of FE incorporation described above is called a **denoted** FE, since Garment is the FE the `Clothing` frame is about. Denoted FEs are found in frames hosting nouns referring to natural kinds. Finally, the omission of core FEs is sometimes possible because of certain relations between them, discussed in the next section.

3.3 Relations between frame elements

All LUs of a frame should have the same FEs. Yet, certain types of variation in FEs are captured not as frame distinctions, but by grouping alternative FEs within a frame. For example, situations such as `Killing` can be brought about either by a sentient Agent or a non-sentient Cause, as in Examples (23) and (24).

(23) [A person$_{Killer}$] **killed** two patients by applying corrosive plasters to their chests.

(24) [The fiery blast$_{Cause}$] **killed** everyone on deck instantly.

Rather than have frames for sentient and non-sentient cases of killing, we include both a (sentient) Agent FE and a (non-sentient) Cause FE in the frame `Killing`, and define a relation between them such that one or the other can appear in a sentence, but not both. This relation is called **Excludes**, formally expressed by an 'xor' connection. Typically, but not always, if FE A excludes FE B, then B also excludes A, forming a reciprocal exclusion relation. Excludes relations are found in 188 frames.

https://doi.org/10.1017/S1351324909005117 Published online by Cambridge University Press

In addition, FrameNet has defined two other types of frame-internal FE relations, 'coreness set' (or 'CoreSet') and 'Requires.' A **coreness set** is formally a disjunction of two or more interdependent FEs. Practically, this means that sentences are informationally complete and pragmatically felicitous with only a subset of them expressed. Consider the `Self_motion` frame, evoked by *walk.v*, *step.v* and many other LUs. It describes situations where a living being moves under its own power (Self_mover FE). The motion can be described in terms of its Source, Path, Goal, and Direction; these are FEs, grouped into a CoreSet. Every annotated sentence for any of the `Self_motion` LUs exhibits between one (25) and four of these CoreSet FEs; (26) contains three of them.

(25) I will just lock up and then [we$_{Self_mover}$] shall **stroll** [home$_{Goal}$] together.

(26) [Therese$_{Self_mover}$] **strolled** [from the back door$_{Source}$] [over$_{Path}$] [to the stove where the fish-kettle still stood. . . $_{Goal}$].

Likewise, a Killer uses a physical device (the Instrument: 'with an axe') and/or performs an action (Means: 'by drowning') to bring about the death of the Victim, but it is not necessary to mention each of them. The `Killing` frame therefore groups Killer, Instrument, and Means into a CoreSet. CoreSet relations appear in 257 frames.

A **Requires** relation between two FEs indicates that if one of them is realized in a sentence, the other must also be. Formally, they are connected via 'implies.' As such, this relation restricts rather than loosens the constraints on LU behavior, but usually, it involves FEs that also have an Excludes relationship to others. Typically, though not always, Requires relations come in pairs, which makes the relationship reciprocal rather than unidirectional. There are 56 frames with Requires relations.

3.4 The frame inheritance relation

FrameNet defines eight relations between frames, of which Inheritance is the strongest one. It links a more general parent frame to a more specific child frame. FrameNet Inheritance semantics is laid out in Subsection 3.4.1. Inheritance can also be compared to LU membership in a frame (3.4.2).

3.4.1 Inheritance semantics and representation

With the Inheritance relation, 'anything which is strictly true about the semantics of the parent must correspond to an equally or more specific fact about the child' (Ruppenhofer *et al.* 2006)[104–106]. This includes the inheritance of all FEs (except for extra-thematic ones) along with their semantic type (Baker, Fillmore and Cronin 2003)[286]. Corresponding FEs are explicitly related across frames. This is necessary because the semantics of the FEs in the child frame might be more specific, and so the FE names might differ. Consider the `Cooking_creation` frame, a child of the `Intentionally_create` frame; see Figure 5 for their definitions and frame elements. The Creator FE from `Intentionally_create` maps onto the more specific Cook

https://doi.org/10.1017/S1351324909005117 Published online by Cambridge University Press

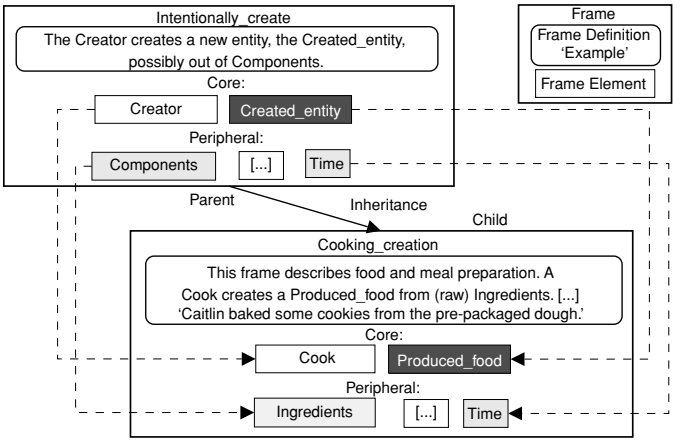


Fig. 5. Frame Inheritance and FE relations: `Intentionally_create` and `Cooking_creation`.

FE in the child frame, and this is documented by a relation between the FEs. Finally, the ontological semantic type of a frame is also inherited. Further formal restrictions in Inheritance, in particular with respect to FEs that are frame-internally linked via FE relations, have been expressed in formal rules (the work of Jan Scheffczyk), which are distributed as part of FrameNet Release 1.3.

Baker *et al.* (2003) characterize frame inheritance as *monotonic*; it has also been compared to **Is-a** relations in ontologies (Ruppenhofer *et al.* 2006)[8]. FrameNet allows multiple Inheritance: 37 frames have more than one parent.

3.4.2 Inheritance and LU membership

Frames are defined in such a way that they describe the behavior of all LUs in it. Conversely, this can be seen as a requirement for the LUs evoking the frame to fulfill all its conditions in terms of FEs, semantic types, and relations between FEs. (Of course, the behavior of an LU may be **more** restricted than the range allowed by the frame. For example, certain LUs in a frame might use exclusively one of two FEs related to each other via Excludes. This makes LU membership conditions very similar to child frame conditions (see Ruppenhofer *et al.* (2006)[127]). Still, an LU and a child frame are fundamentally different insofar as an LU is an *instance* of the semantic class described by the frame, whereas a child frame is itself a class and has its own instances. Figure 6 illustrates these relations.

Despite being conceptual in nature, a frame cannot be defined without knowledge of the LUs evoking it. An example of how the meaning of LUs and frame definitions depend on each other is the restriction that all LUs of a frame must have the same number of FEs. For example, the `Intentionally_create` frame, where a creator intentionally creates a new entity, has two core FEs: Creator and Created_entity. All LUs evoking this frame (including *create.v*, *establish.v*, *production.n*, and others) must allow the linguistic realization of both of these FEs in a sentence, as illustrated by Example (27).

https://doi.org/10.1017/S1351324909005117 Published online by Cambridge University Press

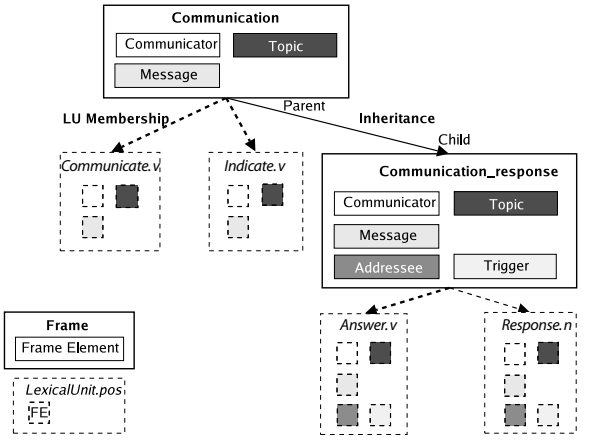


Fig. 6. Frame Inheritance and LUs: `Communication` and `Communication_response` with core FEs.

(27) [Breeders$_{Creator}$] have **established** [their own intelligence network$_{Created_entity}$] in a bid to combat the crime.

As the reader may have noticed, some of the decisions about how much variation between LUs is allowable within a frame or when it is necessary to divide LUs into separate frames may appear somewhat arbitrary. It is necessary to keep in mind the purpose of building such a lexical resource: to document the behavior of LUs and to make useful generalizations over them. In an absolute sense, each LU has slightly different patterns of occurrence, but no useful purpose would be served by creating a frame for each LU. Conversely, FrameNet is committed to documenting a rather precise semantics for each frame, and thus to going beyond the very general semantic roles of, e.g., Case Grammar (Fillmore 1968). In general, the best solution is a combination of dividing frames (with the full apparatus of frame and FE relations) and combining LUs with slightly different behavior in a frame, distinguishing them with semantic types, etc. and noting common patterns of FE alternation with frame-internal FE relations.

3.5 Summary and related work

We have covered frames, FEs, the relations that hold between them, and the notion of semantic type, which are the most important constituents of the frame and LU data structures.[3] We have also explained the notion of LU as a word sense, partially defined by its relation to a particular frame, and the interaction between LU behavior and constraints on the frame. To formalize the notions of the FrameNet model, Figure 7 provides an overview of its components in UML style. Other important components of FrameNet are additional lexical data types, such as

[3] Corresponding roughly to the frames.xml and frRelations.xml files of the data release.

https://doi.org/10.1017/S1351324909005117 Published online by Cambridge University Press

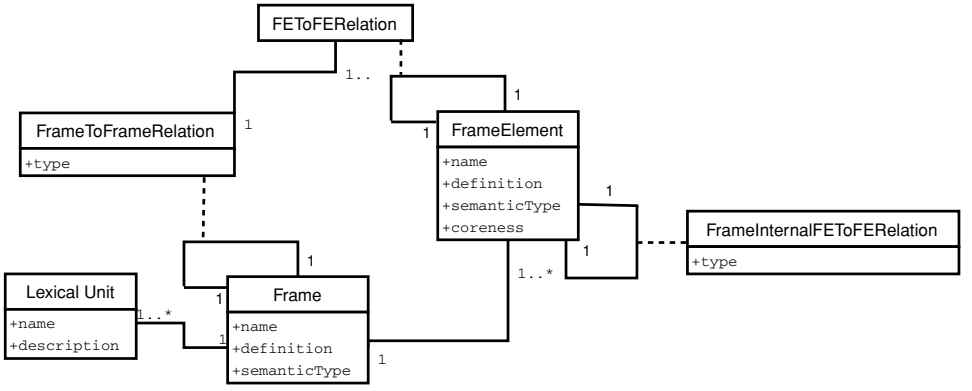


Fig. 7. FrameNet model of frames and related entities in UML style.

lemmas and word forms, as well as annotation data. For an account of the internal structure of these parts of the database (see Baker *et al.* 2003).

A slightly different and more simplified version of the FrameNet model is presented in Francopoulo (2005), who maps it onto lexical markup framework (LMF) (Francopoulo *et al.* 2006). LMF provides a common standardized framework for the construction of lexicons for NLP. The LMF document (ISO-24613:2008), published as an ISO standard in November 2008, describes an NLP semantics extension that covers many of the central components of the FrameNet lexicon model. Efforts to produce a standard that represents annotations produced by FrameNet and other groups are underway in a different sub-group of the ISO technical committee (Ide and Suderman 2007).

4 Multilingual FrameNets

The FrameNet model can be used to build lexical resources for any natural language. In theory, when creating a FrameNet-like resource for a new language, all instances of every data type could be created from scratch, replicating the entire process laid out in Section 2. However, FrameNet is especially appealing for cross-lingual resource creation because much of the content of the FrameNet database can be reused. Given that semantic frames are language independent to a fair degree (Boas 2005; Ruppenhofer *et al.* 2006), it is possible to take the data from the frame part of the database (excluding LUs) as a conceptual starting point for a non-English FrameNet. Instead of creating from scratch all frames and the relevant relations, FrameNets for other languages can thus start with the English frames and determine to what extent they are adequate for LUs from their language.

From an application point of view, this procedure has some practical advantages. First, the resulting resources use the same data model, so they can, for example, share the same API. More interestingly, the resources are also semantically interrelated (aligned) at the level of frames. In the most straightforward case, the same frames are used by FrameNet and a resource for a different language; these two resources are thus aligned by (implicit) cross-lingual equivalence relations between frames

https://doi.org/10.1017/S1351324909005117 Published online by Cambridge University Press

of the same name. This characteristic can be exploited to algorithmically compare language descriptions and to project information across languages. In the future, it might prove useful for semantically enriched machine translation.

Several FrameNets for languages other than English are being created using the same model and reusing the frames. After introducing the most important of these initiatives (4.1), we present a common approach to cross-lingual FrameNet creation (4.2). We discuss cross-lingual equivalence and divergence (4.3), and point out consequences for applications using the resources (4.4).

4.1 Cross-lingual resources using the frameNet model

The most important lexical resources being created along the lines of FrameNet cover Spanish, Japanese, and German; among these, Spanish FrameNet (SFN) (Subirats-Rüggeberg and Petruck 2003; Subirats and Sato 2004), is the most advanced and follows the English FN most closely. The SFN 'starter lexicon' (published on the web in June, 2008) contains 111 frames covering more than 1,000 LUs; 627 LUs are illustrated by corpus annotation.

Japanese FrameNet (JFN) is documenting basic Japanese content words. By April 2008, 22 frames had been populated with 77 Japanese LUs, and illustrative sentences annotated for each of these. The project also investigates to what extent the FrameNet frames are applicable when characterizing Japanese LUs (Ohara *et al.* 2004; Ohara 2008).

The SALSA project is producing FrameNet-style annotation of German newspaper texts (Burchardt *et al.* 2006). For the 2007 data release, SALSA annotated roughly 20,000 instances of 500 predicates, primarily verbs. The annotation is exhaustive in that every instance of these verbs in the corpus is annotated, using the frames from FrameNet release 1.2. As hoped, the vast majority of these frames can be used to describe the syntax and semantics of German (Burchardt *et al.* 2006). To describe word senses not yet covered by FrameNet, SALSA creates additional, predicate-specific 'proto-frames.'

4.2 Cross-lingual resource creation: expanding FrameNet frames

In populating existing FrameNet frames with LUs from their language, non-English FrameNets follow a top-down 'Expand' approach to multilingual resource creation (Vossen 1999; Lönneker-Rodman 2007b). This section illustrates the three steps of the approach, using Spanish FrameNet (SFN) as an example.

(1) **Expand.** SFN imports the frame part of the original FrameNet database (frames, FEs, relations, etc.). Of course, the English LUs are not imported; instead, corresponding Spanish LUs are added to each frame. Subirats and Sato (2004) assume that frames that *share the same name* in English and Spanish also have the same characteristics. SFN is thus implicitly aligned to the English FrameNet at the level of frames. When SFN developers find that the English frames are not appropriate for their language, the differences are dealt with by restructuring.

https://doi.org/10.1017/S1351324909005117 Published online by Cambridge University Press

(2) **Restructure.** Explicit restructuring takes place when new language-specific frames ('Spanish frames') are introduced into SFN; more on this in Section 4.3.1, below. If changes are applied only to attributes or elements of a frame while keeping the original frame name, SFN does not currently perform any restructuring.
(3) **Realign.** Realignment would be the explicit cross-lingual linking of frames that have been newly added or have undergone changes in Spanish FrameNet. Even if there is no English frame directly corresponding to the newly created target language frame, a link could be made to a related English frame, similar to frame Inheritance within a language. (This has not yet been done.)

4.3 Cross-lingual divergence at the frame Level

Frames are conceptual structures and, as such, language independent to a certain degree. Many of them have already been successfully used for other languages in the projects mentioned above. However, neither does (English) FrameNet seek to define semantic primitives, nor does defining a frame in FrameNet guarantee that it is universal. Consequently, adding frames to a non-English FrameNet, or having slightly different frames in FrameNets of different languages does not invalidate work that has been done on frames for one language. On the contrary, the FrameNet model provides a way for describing these cross-lingual differences, often not purely conceptual, but rather resulting from an interdependence between (conceptual) frames and the behavior of (language-specific) LUs.

This section discusses some types of observed cross-lingual differences, from frames as conceptual entities (4.3.1) and their textual definitions (4.3.2) to the set of FEs in a frame (4.3.3) and coreness (4.3.4). The FrameNet model defines frame equivalence and provides mechanisms for documenting cross-lingual differences (4.3.5).

4.3.1 New frames

There are at least two possible reasons for the need to define a completely new frame in a non-English FrameNet: **1.** existing frame definitions in the corresponding semantic domain of the English FrameNet are inadequate; or **2.** the corresponding English frame and LUs have not been defined yet, though they might be fully equivalent to those needed for the non-English FN, once defined.

In Spanish and Japanese FrameNet, some existing English frames were found to be inadequate for the respective target language. For example, the Motion-related frames in FrameNet 1.3 do not cover some Spanish LUs, for which frames such as `Motion_manner` and `Return` were therefore added to SFN.

Figure 8 compares the English `Self_motion` frame to Spanish `Motion_manner`. The `Self_motion` frame (see 3.3, above) describes situations where a Self_mover moves, either along a Path from a Source to a Goal in a directed fashion (Direction), or within an area, but with no single linear path (Area). Spanish does not use the `Self_motion` frame; the corresponding Spanish LUs are found in several other frames instead. One of them is the Spanish-only `Motion_manner` frame, with LUs

https://doi.org/10.1017/S1351324909005117 Published online by Cambridge University Press

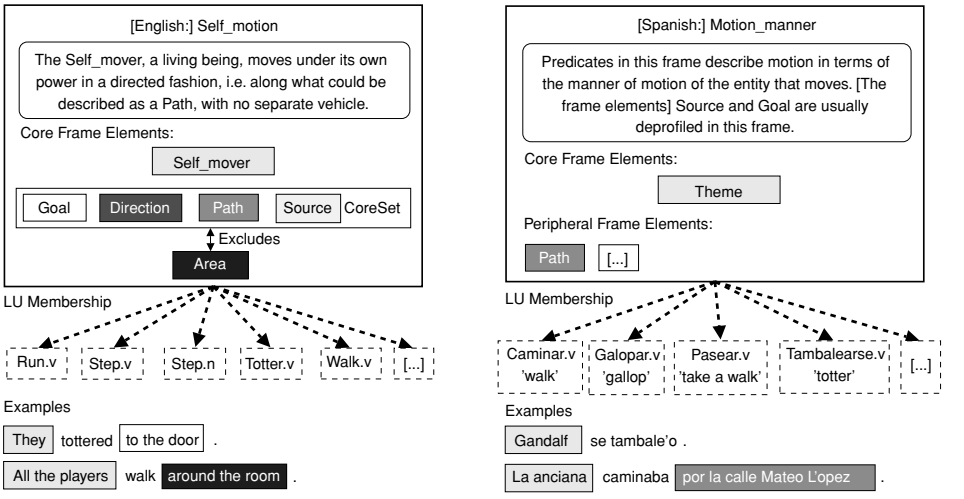


Fig. 8. The English Self_motion frame and the Spanish Motion_manner frame.

including *caminar.v* 'walk' and *pasear.v* 'take a walk,' describing a scene where an entity moves in a particular manner. Stressing this aspect of the motion event, the Spanish LUs imply very little with respect to path and direction. Consequently, there is only one core FE, the moving entity (Theme FE). Other FEs with core status in English `Self_motion`, such as Path, are peripheral in Spanish `Motion_manner`. Spanish sentences illustrating `Motion_manner` are syntactically and semantically well-formed with only Theme being explicitly realized, such as (28) and (29) from SFN; the Path expression in (29) is optional.

(28) [Gandalf$_{Theme}$] se tambaleó. *Gandalf tottered.*

(29) [La anciana$_{Theme}$] caminaba [por la calle Mateo López$_{Path}$]. *The old lady was walking along Mateo López Street.*

If LUs are not created following the Expand approach but in some other way, it is likely that frames will be needed that have not yet been covered by FrameNet. This is the case with the proto-frames created by the SALSA project. Given the way SALSA annotation proceeds, LUs are determined by the text to be annotated and do not necessarily pertain to domains with existing frames. While most of the proto-frames created by SALSA are likely to correspond to English frames not yet defined by FrameNet, some of them might be specific for German.

4.3.2 Language-specific information within definitions

As we have seen in Section 3.1.2, frame definitions may contain embedded information, such as example sentences and generalizations about syntactic realization patterns. In FrameNets for other languages, English examples are generally replaced by sentences in the respective languages, and the summaries of syntactic realization patterns also usually need to be adapted. To illustrate, the definition of the `Perception_body` frame, which has to do with bodily sensations and includes

https://doi.org/10.1017/S1351324909005117 Published online by Cambridge University Press

LUs such as *ache*, *hurt*, and *tingle*, states that the Experiencer typically appears as a possessive determiner (30).

(30) The body part affected... is typically expressed by the noun heading the external argument [i.e., subject], and this noun is typically accompanied by a possessive determiner that refers to the possessor of the body part: *My head hurts!*.

But this generalization does not hold for Slovenian, as illustrated by (31).

(31) [Noge$_{Body_part}$] [me$_{Experiencer}$] bolijo.
Legs-NOM me-ACC hurt.
My legs hurt.

In (31), the Experiencer is realized as direct object of the (non-reflexive) verb *boleti* 'hurt,' the natural way to express this state of affairs with this Slovenian word. Further words that should be considered LUs of `Perception_body`, in Slovenian as well as other languages including Japanese and Italian, exhibit a range of different syntactic valences (see e.g., (Lambrecht 1994)[137]).

The comments about syntactic patterns and the example sentences found in English frame definitions are not intended to express language universals or norms. Therefore, if we adapt an EFN frame to another language and intend that the frame be the same in both languages, we cannot modify the the conceptual information in the frame definition, but we can freely change embedded statements about the syntax, or the example sentences.

4.3.3 Differences in frame elements

When comparing frames across languages, differences in the number of FEs sometimes become obvious. For example, the `Motion` frame in SFN contains an extra-thematic FE Intention, which is not present in the English `Motion` frame. Example (32) adapted from (Subirats and Sato 2004) illustrates that the Spanish Intention FE is semantically and syntactically different from Purpose, which is found in both English and Spanish. English does not syntactically distinguish between the information rendered as Intention *versus* Purpose in Spanish, and cannot realize this distinction in two separate dependents of the frame-evoking word.

(32) Juan **fue** a San Francisco [a visitar a un amigo$_{Intention}$] [para pedir-le
Juan went to San Francisco to visit to a friend to ask-him
dinero$_{Purpose}$].
money.
John went to San Francisco [to visit a friend and ask him for money$_{Purpose}$].

Similar cases, in which there are a larger number of FEs in certain frames for a non-English FrameNet as compared to the most similar EFN frame, have also been reported in Japanese (Ohara *et al.* 2004) and German (Burchardt *et al.* 2006).

https://doi.org/10.1017/S1351324909005117 Published online by Cambridge University Press

4.3.4 Different coreness status

Formal and checkable conditions on coreness status such as obligatory overt instantiation for core FEs depend on the syntactic behavior of the LUs in the frame; therefore, coreness might be language-specific, as suggested in the following example from German and Slovenian.

The German and Slovenian verbs corresponding to English *be about* are tentatively considered members of the `Topic` frame, which has to do with the relation between a text and what it is about. The frame has three core FEs: Communicator, Topic, and Text. The Text FE is defined as 'a set of propositions that is coherent in being about a Topic.' However, the German and Slovenian verbs, *gehen (um)* and *iti (za)*, rarely occur in combination with the Text FE, illustrated in (33) for German and (34) for Slovenian.

(33) Es **geht** [mir$_{Communicator}$] [um zwei Probleme$_{Topic}$].
It-EXPLETIVE goes I-DAT around two problems.

(34) **Gre** [mi$_{Communicator}$] [za dva problema$_{Topic}$].
Goes I-DAT for two problems.
I am concerned with two problems.

The Communicator FE is omissible in both languages. However, when omitted, the Communicator must be someone already understood in the linguistic or discourse context, as in German (35); this is a case of null instantiation.

(35) Es geht [um mein Kind$_{Topic}$]. [(NI)$_{Communicator}$]
It-EXPLETIVE goes around my child.
It is about my child.

These observations suggest that Topic and Communicator are core FEs of `Topic` or some similar frame in these languages, but the Text FE is peripheral. While it might be possible to explain the observations by coverage limitations (4.3.1), i.e., the necessary frame is not yet defined for English either, the example at least illustrates the *method* of comparing FE coreness when adapting a frame across languages.

Cross-lingual differences in coreness status have not yet been reported in the literature on other languages, but when comparing FrameNet release 1.3 and Spanish FrameNet (as of October, 2008), we found that such differences are not uncommon. Spanish FrameNet has edited a total of 103 frames from Release 1.3, and coreness changes occur in 26 of them. Table 2 shows the differences by type. We notice a tendency to have fewer FEs towards the 'core' of a frame in Spanish; moving from English to Spanish, 72 FEs move toward lower coreness status, while only 17 move higher.

4.3.5 Summary and evaluation

Two frames should be regarded as *strictly equivalent* only when all relations, FEs (except extra-thematic ones), and semantic types are the same. The frames of two FrameNets in different languages do not always fulfill these requirements.

https://doi.org/10.1017/S1351324909005117 Published online by Cambridge University Press

Table 2. *FE coreness status: differences between English and Spanish*

Spanish → **English ↓**	**Core**	**Peripheral**	**Extra-thematic**
Core	N.A.	18 FEs in 12 frames	1 FE in 1 frame
Peripheral	2 FEs in 1 frame	N.A.	53 FEs in 21 frames
Extra-thematic	2 FEs in 2 frames	13 FEs in 11 frames	N.A.

Differences can arise, not only due to the introduction of new frames, but also because of divergent FEs and their attributes, naturally following from the behavior of (language specific) LUs. There is no general rule as to which and how many minor divergences from a frame created for English are allowed in a non-English FrameNet, or at what point it is necessary to create a new frame instead. On the one hand, it is desirable to maintain direct links to the English FrameNet frames by reusing their names and FEs, even if modified. On the other hand, each FrameNet should accurately describe the actual behavior of the LUs in its language, which might deviate from that of English LUs, and this might be best reflected by introducing a language-specific frame. Minimally, one would expect the *conceptual* information in the frame definition to be identical for frames of the same name, across languages; however, this would include also identity of coreness status, at least for core FEs. Current practice shows that this is generally the case, though not always.

Having a shared data model can facilitate the comparison of English and non-English instantiations of the components and attributes (frames, FEs, coreness status, semantic type). In fact, the program that was written to compare successive versions of English FN[4] can also be used to compare EFN release 1.3 and Spanish FN, to find cross-linguistic differences in coreness status. The report thus generated can feed into the process of defining the relations between EFN and new or conceptually different frames in the target language.

4.4 Multilingual natural language processing with FrameNet

With FrameNet-style resources of a substantial size available for English, Spanish, and German, and similar resources for other languages under construction, FrameNet is part of a linguistic data inventory for multilingual applications, including machine translation (MT) and cross-lingual information retrieval.

Statistical MT has begun to incorporate tree-based models, empowering it to capture hierarchical syntactic structure, as exemplified by Graehl, Knight and May (2008). Since FrameNet annotation can be converted into tree form (Baker, Ellsworth

[4] This program produces the difference file included with the FN data distribution.

https://doi.org/10.1017/S1351324909005117 Published online by Cambridge University Press

and Erk 2007), both semantic and syntactic knowledge from FrameNet could be incorporated in this way into statistical MT systems. However, it may take a while before syntactic knowledge is fully incorporated in statistical MT systems, so frame-semantic-aware statistical MT systems are not likely to be developed in the immediate future. This section therefore presents some other current cross-linguistic uses of the FrameNet model and data, thereby illustrating the merit of the semantic frame model for multi-lingual NLP in general.

4.4.1 Describing translation divergences

The FrameNet model handles many common types of *translation divergences* (Dorr 1994) by treating them as different linguistic realizations of the same frame. The general idea can be illustrated by monolingual paraphrasing. Consider the `Emotion_directed` frame, which describes an Experiencer feeling an emotional response to a Stimulus. The frame contains nouns and adjectives, such as *anger* and *angry*, which can frequently convey the same concept, as in (36).

(36) a. [Christina$_{Experiencer}$] blushed, **angry** [at being spoken to like a member of staff$_{Stimulus}$].

b. [Christina$_{Experiencer}$] blushed *with* **anger** [at being spoken to like a member of staff$_{Stimulus}$].

Examples (36a) and (36b) are monolingual paraphrases of each other. The frame is evoked by words of different parts of speech, entailing differences in syntactic structure and possibly emphasis, but the examples are still semantically equivalent; the same set of FEs is realized in both versions of (36).

Similarly, frames and FEs can represent the common semantics of parallel bilingual texts, even in the case of translation divergences. For example, in a parallel version of the novel *The Hound of the Baskervilles*, the situation of `Being_at_risk` (where an Asset is exposed to a Harmful_event or Dangerous_entity) is evoked by an adjective in the English original (37), and by a noun in German (38).

(37) [English] Is [he$_{Asset}$] **safe**? [(NI)$_{Harmful_event}$]

(38) [German] Ist [er$_{Asset}$] in **Sicherheit**? [(NI)$_{Harmful_event}$]
lit. Is he in security?

FrameNet lexical entry reports also suggest possible ways to translate a sentence, using different LUs. Consider the English verbs *judge* and *rate* in the `Assessing` frame, where an Assessor examines a Phenomenon to figure out its Value (Value is not a core FE). The most frequent syntactic pattern with `Assessing` verbs is for the Assessor to appear as a noun phrase in external argument (subject) position, and for the Phenomenon to be the object; the lexical entry reports in Tables 3 and 4 illustrate this. However, with the German verb *bedeuten* (in the `Assessing` frame), the Assessor is always expressed as a noun phrase in the dative case, as in (39), taken from the SALSA annotated data.

https://doi.org/10.1017/S1351324909005117 Published online by Cambridge University Press

Table 3. *Lexical entry report for the verb* judge *in the* `Assessing` *frame*

Judge.v 6 TOTAL	Pattern	
	Assessor	Phenomenon
4	NP External	NP Object
1	(NI) –	NP External
1	(NI) –	NP Object

Table 4. *Lexical entry report for the verb* rate *in the* `Assessing` *frame*

Rate.v 5 TOTAL	Pattern	
	Assessor	Phenomenon
3	NP External	NP Object
1	NP External	PP(as) Dep.
1	PP(with) Dep.	NP External

(39) [Im Korps zu sein$_{Phenomenon}$] **bedeutet** [ihnen$_{Assessor}$]
[In.ART.DAT.SG corps to be$_{Phenomenon}$] means [they.DAT.PL$_{Assessor}$]
mehr als die Mitgliedschaft in einer Partei.
more than the membership in a party.
To be in the corps means more to them than party membership.

Using the most frequent syntactic realization pattern observed for English verbs to translate (39), thus results in a cross-lingual *thematic divergence*, with different FEs as subjects, as in (40):

(40) [They$_{Assessor}$] **rate** [corps membership$_{Phenomenon}$] higher than party membership.

However, FrameNet also documents a less frequent realization pattern of the verb *rate*, which does not involve this divergence with respect to German (41). The lexical entries also show that the theme-preserving pattern in (41) is not possible with the verb *judge* (cf. 42).

(41) [Corps membership$_{Phenomenon}$] **rates** higher [with them$_{Assessor}$] than party membership does.

(42) *[Corps membership$_{Phenomenon}$] **judges** higher [with them$_{Assessor}$] than party membership does.

Many translation divergences can thus be captured by using the same frame and by semantically aligning the FEs, abstracting away from syntactic variations. The mark-up of FEs also facilitates user-friendly visualization of cross-lingual semantic correspondences despite syntactic divergences, for example, in bilingual dictionaries

https://doi.org/10.1017/S1351324909005117 Published online by Cambridge University Press

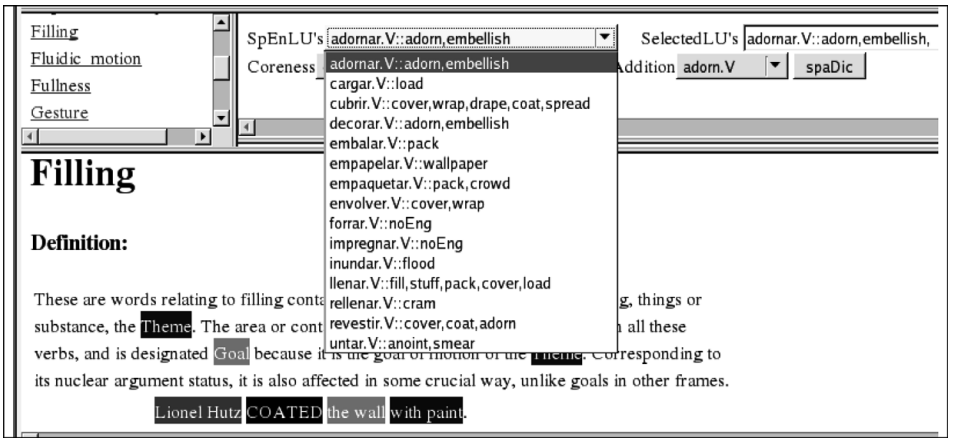


Fig. 9. Spanish LUs of the `Filling` frame with English translations in FrameSQL.

(Lönneker-Rodman 2007a). Finally, in natural language generation, FrameNet data can be used to prevent (or to encourage) the production of divergences.

4.4.2 A viewer for comparing cross-lingual FrameNet data

The FrameSQL application is a viewer for mono- and multilingual FrameNet data (Sato 2008). An automatic summary function for equivalent LUs has been implemented within FrameSQL, to compare annotation data from two FrameNets. When aligning frames and LUs from different languages, FrameSQL considers frames of the same name as equivalent; to align LUs, FrameSQL relies on translation equivalents from a bilingual electronic dictionary.[5] All translations proposed by the dictionary are treated as candidate senses in the target language. To select the most likely correct sense or senses, all frames hosting candidate translations in the target language are compared with the frame of the source language LU, and only senses belonging to the same frame are retained. Figure 9 shows the alignment of Spanish and English LUs of the `Filling` frame (defined in Section 3.2.2 above).

FrameSQL then facilitates the comparison of semantic valence patterns, summarizing the annotation for aligned LUs. Consider sentences that evoke the `Filling` frame and instantiate a combination of Agent, Goal (the area being filled), and Theme (the physical object which changes location), with a null instantiation of the latter, as in (43).

(43) . . . the power of [the emperor$_{Agent}$] to **embellish** [the city and court$_{Goal}$] as he wished. [(NI)$_{Theme}$]

Figure 10 shows a FrameSQL contrastive summary of the valences annotated for Spanish *adornar* from the `Filling` frame and the English verbs *adorn* and *embellish*. The particular FE combination of Agent, Goal, and a null-instantiated Theme (indicated by enclosing parentheses) has been annotated three times for

[5] FrameSQL uses the *Grupo Anaya Spanish-English Dictionary* (H. Sato, p.c.).

https://doi.org/10.1017/S1351324909005117 Published online by Cambridge University Press

	Adornar	Adorn	Embellish
(Agent) + (Goal) + Theme	01		
(Agent) + Goal + Theme	06	02	03
Cause + Goal + (Theme)	01		
(Agent) + Goal + (Theme)	02		
Agent + Goal + Theme	03	07	05
Agent + Goal + (Theme)	03		03
(Agent) + (Goal) + (Theme)	01		

Fig. 10. Semantic valences of the Spanish LU *adornar* and two likely English equivalents (null-instantiated FEs in parentheses).

adornar, but only for one of the two English LUs it is aligned with *embellish*. The annotated sentences summarized in the table, including (43), can be viewed directly in FrameSQL.

Cross-lingual comparisons of the kind exemplified by FrameSQL are intended to support decisions in translation and natural language generation, both within systems assisting human translators and in fully automated systems.

4.4.3 Projecting frame-semantic annotation

The Expand approach to manual cross-lingual FrameNet creation is motivated by the assumption that most frames can be preserved across languages, in spite of the minor differences pointed out in Section 4.3. Similar assumptions underlie *automatic projection efforts*, where information from the English FrameNet lexicon and annotation is transferred to other languages. These approaches predominantly use statistical methods on bilingual parallel corpora, where FrameNet-style annotation labels on the English side are transferred to the corresponding words and phrases within the parallel foreign-language sentence.

Padó (2007) describes a general approach, based on word alignment information, various filters, and syntactic information for FE projection. Padó works with a manually annotated subset of the EuroParl corpus (Koehn 2005) in English, German, and French. Projecting frame labels of English LUs, he automatically derives German and French lemmas presumably evoking the same frame. For foreign language frames comparable in size to the original FrameNet (ten LUs on average), a precision of 68 per cent for German and 50 per cent for French is achieved.

For the `Filling` frame, Tables 5 and 6 contain the ten most frequently occurring French candidate lemmas, produced with two different filter combinations.[6] Each table row shows an LU candidate, followed by a manual evaluation provided by us, based on corpus searches and native speaker judgments: If the French lemma does not belong to `Filling`, we propose a more appropriate frame. The third column

[6] Thanks to Sebastian Padó for sharing this data.

https://doi.org/10.1017/S1351324909005117 Published online by Cambridge University Press

Table 5. *French candidate lemmas for* `Filling`*, no predominant frame filter*

French Lemma	Closest Frame	Links	English Lemma
couvrir	Filling	1327	Cover
combler	Filling	115	Fill
concerner	Topic	85	Cover
viser	Aiming	74	Cover
aborder	Topic	72	Cover
étendre	Placing	68	Spread
remplir	Filling	67	Fill
traiter	Topic	56	Cover
englober	Inclusion	44	Cover
inclure	Inclusion	44	Cover

Table 6. *French candidate lemmas for* `Filling`*, predominant frame filter*

French Lemma	Closest Frame	Links	English Lemma
semer	Filling	22	Sow
planter	Filling	11	Plant
plantation	Filling	11	Plant
emballer	Filling	5	Wrap
tracer	Create_representation	4	Pave
peindre	Filling	4	Paint
submerger	Adorning	2	Flood
semis	Filling	2	Sow
entasser	Placing	2	Cram
embellir	Filling	2	Embellish

contains the absolute frequency of word alignment links involving this lemma, after filtering. Finally, the English lemma most frequently aligned to the French candidate is provided.

In Table 5, some of the proposed French lemmas do not refer to `Filling` or any other concrete actions; these results reflect the polysemy of English lemmas. So, for Table 6, a predominant frame filter was added, which treats French lemmas as monosemous, always assigning them to the frame with the highest probability for this lemma (Padó 2007). With the filter, the algorithm proposes more correct lemmas for `Filling`, and some that are LUs of frames closely related in the frame hierarchy (`Adorning`, `Placing`). Yet many concrete frames, such as `Filling`, are rarely evoked in the EuroParl corpus, with the result that the candidate lemmas filtered by predominant frame have a low frequency, and some high-frequency `Filling` LUs, correctly detected without this filter, are not found.

https://doi.org/10.1017/S1351324909005117 Published online by Cambridge University Press

In a different experiment, Johansson and Nugues (2006) transfer frame-semantic annotation from English to Swedish, also relying on information from word alignment. The subsequent correction, however, involves a number of heuristics specifically developed for the language pair English–Swedish. Since there was no English–Swedish parallel corpus with manual FrameNet annotation, Johansson and Nugues first automatically added frame-semantic annotation to the English side and then projected the annotation onto aligned constituents in Swedish. Automatic annotation is required for a number of other applications as well, since most of them deal with texts different from those covered by the manual FrameNet annotation. The design and implementation of automatic systems that create frame-semantic annotation is dealt with in the next section.

5 Automatic Frame-Semantic Annotation

For frame semantic analysis to be used on any text of interest, a way must be found to add frame semantic annotation automatically. This task can be subsumed under 'automatic semantic role labeling' (ASRL, discussed in Subsection 5.1). Subsection 5.2 discusses implications of the FrameNet model for automatic frame-semantic annotation. To investigate how ASRL systems might be improved by taking into account certain constraints following from the FrameNet model, we analyze the output of an available system (Subsection 5.3).

5.1 Automatic semantic role labeling

Automatic semantic role labeling enriches a text with a representation of its predicate-argument structure. For each predicate, the system tries to determine the relationship to its semantic arguments or roles (in FrameNet terms, the relationship of each *LU* to its *FEs*). ASRL systems use machine learning methods to perform this task automatically, which are trained on large data sets, such as the annotation part of the FrameNet database, or the PropBank corpus (Palmer, Gildea and Kingsbury 2005). We will first describe ASRL in terms of automatic frame semantic annotation, and then briefly compare this to PropBank-style semantic labeling.

5.1.1 Automatic Frame Semantic Annotation

Early systems for automatic frame semantic annotation (AFSA) were designed to tag predefined sentence constituents with FE labels pertaining to a prelabeled LU. Later, the detection of sentence constituents bearing FEs was included into the task (Gildea and Jurafsky 2002). In the SemEval-2007 task on frame semantic structure extraction (Baker *et al.* 2007), AFSA systems were supposed to perform tasks similar to the actual (manual) FrameNet annotation process.

Thus, ideally, a full-fledged system should first decide, for each word or multi-word, whether it is a form of a frame-evoking LU, and if so, which frame it evokes. (For the same lemma, several LUs might be available in FrameNet, making it necessary to disambiguate between frames. Frame disambiguation is basically the same as

https://doi.org/10.1017/S1351324909005117 Published online by Cambridge University Press

word sense disambiguation (Erk 2005).) With respect to each LU, the system next identifies sentence constituents representing FEs, and then labels the constituents with the appropriate FE name.

A small number of LUs and even frames in the SemEval-2007 test data were entirely new; the new LUs might properly belong to an existing frame, or to a frame not yet created (but which was created in the gold-standard used for scoring). This represented a realistic test of a semantic analysis system, since any new text is likely to contain a certain number of previously unseen LUs, but it also made the task considerably harder than standard WSD exercises, in which all senses are known.

For each instance of each predicate, the systems participating in SemEval had to decide whether it represented a known LU, a new sense of a lemma that already had an LU for another sense, or a previously unseen lemma that should be an LU. When frames were missing, the systems were not expected to propose new ones (which would have been almost impossible), but they were given credit for identifying a semantically close existing frame. The detection of new senses (missing LUs) has been modeled as outlier detection by Erk (2006) and was tackled by Johansson and Nugues (2007) by adding new senses to semantically close frames. For the full task of frame and FE recognition and labeling, the best SemEval-2007 system achieved a recall of 0.46 and a precision of 0.67, for an F_1 score of 0.55 on the 'easiest' text in the test set. Perhaps the closest comparison to this score would be the F_1 score of 0.69 attained by the highest-scoring system on the closed-challenge semantic dependency measure in the CoNLL-2008 shared task (Surdeanu *et al.* 2008).[7] For the more limited task of frame identification, participants managed to reach F-scores as high as 0.75 (Baker *et al.* 2007), comparable to scores in the WSD tasks in SenseEval-3, where F_1 scores of 0.59 were reported for WordNet senses, where the baseline most frequent sense would yield 0.51. An idealized program flow for treating a frame-evoking word is given in Figure 11.

5.1.2 *PropBank*

The Proposition Bank ('PropBank') (Palmer *et al.* 2005) is also a large-scale research project engaged in annotating English text with semantic role labels and creating a lexicon in the process which is of the same order of magnitude as the FrameNet lexicon, although PropBank annotates only verbs.[8] The two projects start from rather different theoretical assumptions, so it is instructive to compare their results.

As Palmer *et al.* (2005) point out, PropBank was conceived from the beginning as providing data for machine learning. The project therefore started from a set of texts, all of which were to be annotated, rather than having the goal of creating a lexicon, which was to be exemplified by examples from texts. Encountering the same difficulty as FrameNet in choosing a small set of semantic roles that would

[7] However, in that task partial credit was given for getting the (PropBank) semantic dependency relations right, even if the predicator was given the wrong sense; in the FN task, both the LU and the relations had to be right.

[8] The NomBank project (Meyers *et al.* 2004) provided annotation of nouns and their argument structure in the same texts and using broadly the same principles as PropBank.

https://doi.org/10.1017/S1351324909005117 Published online by Cambridge University Press

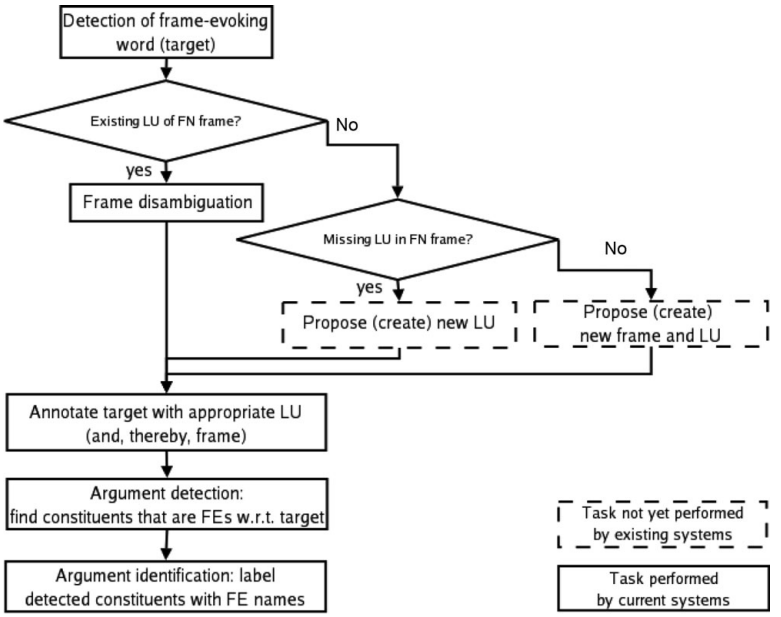


Fig. 11. Program flow of sample frame-semantic annotation system.

cover all types of eventualities, PropBank chose 20 labels for all semantic roles; these are divided into a set of core semantic roles Arg0–Arg5, which are defined on a verb-by-verb basis, and a larger set of 'modifiers' for notions such as time, place, and manner, common to all predicates. A 'descriptor' was written for each core argument label for each verb, describing how the general semantic role is to be interpreted in the case of the specific sense of the verb, but it 'does not have any theoretical standing' (Palmer *et al.* 2005)[76]. As to the broader interpretation of the core argument labels, Arg0–Arg5, Arg0 is generally the Proto-Agent (Dowty 1991) and Arg1 is generally the Proto-Patient, but 'No consistent generalizations can be made across verbs for the higher-numbered arguments, though an effort has been made to consistently define roles across members of VerbNet classes,' (Palmer *et al.* 2005)[75].

These groupings by VerbNet classes (ultimately derived from Levin's (1993) classes) are the nearest equivalent to FrameNet frames. They apparently do not play a major role in PropBank annotation, but PropBank has followed the Levin classes 'quite closely,' even though this means that they have 'propagated some of the errors of that work' (Kingsbury, Palmer and Marcus 2002). PropBank staff have made major efforts to relate the Levin classes, intersections of the Levin classes (Dang *et al.* 1998), and FrameNet frames, producing the VerbNet resource, which has been used to populate PropBank with groups of related verbs (Kipper, Dang and Palmer 2000). Work on the VerbNet-FrameNet mapping is continuing with the cooperation of FrameNet staff.

The difference in lexical coverage of FrameNet and PropBank arises mainly from their different approaches: PropBank has annotated all the verbs in a section of

https://doi.org/10.1017/S1351324909005117 Published online by Cambridge University Press

the Wall Street Journal, while FrameNet has worked on the most general English words, including nouns, verbs, adjectives, and some adverbs and conjunctions. Thus PropBank has a rich collection of verbs related to financial affairs that are not found in FrameNet (e.g. *accrue.v, arbitrage.v, bankroll.v, bottom.v, capitalize.v, churn.v*), while FrameNet has more verbs for notions such as body movement, etc.

FN often splits senses, creating two or more LUs for the same lemma in different frames, where PropBank has a single sense. For example, *confirm.v* has only one roleset in PropBank, with the gloss 'confirm, attest the truth of a proposition;' the same verb appears in four frames in FrameNet:

- Verification (*The trustees contacted the Commission to* **confirm** *this…*)
- Rite (*The bishops will be prepared to* **confirm** *children…*)
- Statement (*The Chancellor* **confirmed** *in his pre-Budget report that…*)
- Evidence (*Tests of spatial visualisation…* **confirm** *that this is the case.*)

In some cases, FrameNet treats metaphorical senses in a separate frame from the literal (source) sense, where PropBank does not separate them–indeed, the metaphorical use is often the only one in PropBank, since financial events are predominantly described metaphorically. The fact that inter-annotator agreement is somewhat higher for PropBank annotation than for FrameNet may be explained in part by finer sense divisions in FrameNet. For further information, Palmer *et al.* (2005) contains an excellent discussion of the differences and similarities of PropBank and FrameNet, and Màrquez *et al.* (2008) is a volume on ASRL with PropBank.

5.2 Implications of the FrameNet model for automatic annotation

When the manually annotated FrameNet data is used to train a model or classifier for automatic frame-semantic annotation, FrameNet sentences are first parsed with a syntactic parser, and a variety of training features is extracted from the data (Gildea and Jurafsky 2002; Erk and Padó 2006; Bejan and Hathaway 2007; Johansson and Nugues 2007). Three groups of features can be distinguished:

(1) **syntactic features**, e.g. the path through the syntactic parse tree, from a frame-evoking word to a constituent that might represent a FE;
(2) **lexical features**, e.g. the head word of a phrase possibly representing a FE;
(3) **semantic features**, e.g. whether a phrase represents a 'Human;' see (Bejan and Hathaway 2007)[462] for details.

The use of syntactic features is motivated by the observation that FEs tend to be realized by similar syntactic patterns, across contexts. Lexical features capture the fact that certain FEs of some frames are restricted with respect to their lexical realization. For example, the Text FE of the `Topic` frame prefers head nouns such as *book*, *article*, *chapter*, and *section*. To cover more flexible lexical realizations, a generalization of this feature is usually added; this can be done by clustering head words (Gildea and Jurafsky 2002) or by resorting to the head word's part of speech instead of its word form or lemma. Finally, the use of semantic features is motivated by the same considerations as the labeling of FEs with semantic type information

https://doi.org/10.1017/S1351324909005117 Published online by Cambridge University Press

in FrameNet, although to our knowledge no system directly exploits the semantic type labels.

The FrameNet model specifies that the same set of FEs is available for all LUs in a frame. This implies that machine learning systems can (to some extent) be trained on all the data in a frame, generalizing over the LUs. This is especially useful when limited training data is available for some LUs in the frame (Gildea and Jurafsky 2002). The FrameNet database contains many LUs that are similar in behavior to others in the same frame (e.g., *Monday*, *Tuesday*, etc. in the `Calendric_unit` frame); not all of these are illustrated by annotated sentences. It is therefore currently standard to train models for argument identification per frame, instead of per LU (Erk and Padó 2007; Bejan and Hathaway 2007).

The FrameNet model offers further possibilities for training automatic semantic role labelers which are not yet commonly exploited. Models for extra-thematic FEs might be trained on the entire FrameNet annotation database, given that their interpretation is not frame specific. Training of FE recognizers and labelers could also be fine-tuned, using argument-specific feature selection as implemented by Pradhan *et al.* (2005) for ASRL with PropBank.

5.3 Checking automatic annotation output against the FrameNet model

The specifications of the FrameNet model offer the possibility of checking AFSA output against formal requirements. We will discuss how the implications behind coreness and FE relations can be exploited for post-processing automatic annotations.

5.3.1 Experiment design and results

We used Shalmaneser 1.1[9] (Erk and Padó 2006) with classifiers pre-trained on FrameNet 1.3 to label 200 sentences from the EuroParl corpus and 269 sentences from *The Hound of the Baskervilles* (A. C. Doyle). The system produced exactly 700 frame annotations on EuroParl and 900 on the novel. These were checked against the FrameNet database as of September 2007 for the existence of core FEs and for FE relation violations.

For each annotation of a frame-evoking LU in the analyzed text, the post-processor checks whether annotations exist for all core FEs, given the implications of CoreSet, Excludes, and Requires relations. Missing core FEs are counted and their labels are returned. Additionally, violations of Excludes and Requires relations are recorded.

Table 7 gives a quantitative overview of the results. Across corpora and across part-of-speech of the frame-evoking LU, approximately 60 per cent of the annotations violate at least one of the constraints. Next, we were interested in what caused the constraints of the FrameNet model to be violated.

[9] Available from the SALSA project web site: `http://www.coli.uni-saarland.de/projects/salsa/shal/` [October 17, 2007].

https://doi.org/10.1017/S1351324909005117 Published online by Cambridge University Press

Table 7. *Checking ASRL output against frame definitions in FrameNet database*

Corpus	POS of LU	Frames	No violations	Frames missing core FEs	Frames violating Excludes	Frames violating Requires
EuroParl	Adj	43	16 (37.21%)	27 (62.79%)	0	0
	Noun	276	84 (30.43%)	189 (68.48%)	5	11
	Verb	381	145 (38.06%)	235 (61.68%)	2	3
	Total	700	245 (35.00%)	451 (64.43%)	7	14
Hound of Baskervilles	Adj	71	27 (38.03%)	43 (60.56%)	0	2
	Noun	292	154 (52.74%)	138 (47.26%)	0	5
	Verb	537	198 (36.87%)	338 (62.94%)	2	4
	Total	900	379 (42.11%)	519 (57.67%)	2	11

5.3.2 *Reasons for missing core FEs in automatic frame-semantic annotation*

To investigate the reasons for missing core FEs in AFSA output, we manually examined 5 per cent of the noun and verb annotations reported as missing core FEs, in total 15 nouns and 28 verbs. The results point to a variety of reasons:

Null instantiation: Omission of core FEs licensed by null instantiation (see Section 3.2.4, above) occurs with seven of the 15 analyzed frame-evoking nouns and with four of the 28 analyzed frame-evoking verbs. To our knowledge, none of the currently implemented frame-semantic annotation systems provides information on null-instantiated FEs, which makes checking for core FEs (as a constraint on the output) problematic.

Missed FE: With five of the analyzed frame-evoking nouns and with 11 of the analyzed verbs, the core FE is expressed in the analyzed sentence, but the system fails to identify and classify it (Examples (44) and (45)).

(44) . . . , my mind paralyzed by the dreadful shape which had **sprung** [out$_{Source}$] [upon us$_{Area}$] [from the shadows of the fog$_{Source}$]. (`Self_motion` frame with FE Self_mover not annotated)

(45) The steps passed along the path on the other side of the wall under which [I$_{Agent}$] **crouched**. (`Posture` frame with FE Location not annotated)

Full-text annotation requires the systems to handle complex syntactic structures that might be infrequent in the lexicographic training data, such as the relative clauses in the examples above. To alleviate the problem, future AFSA systems might be trained on both lexicographic and full-text annotation. (Of course, this will first require more manual or semi-automatic full-text annotation.)

Wrong sense: The tagged word form is not an LU of the frame selected by the system, and therefore the context does not exhibit the core FEs of this frame. This

https://doi.org/10.1017/S1351324909005117 Published online by Cambridge University Press

is an issue at the sense detection and disambiguation stage. Since the Shalmaneser system is optimized for a slightly differently annotated data set, a common class of sense errors occur with support verbs, such as *pay* in (46), which is wrongly assigned to the frame `Commerce_pay`; instead, *attention* should be considered the main predicator, with *pay* as just a support verb.

(46) . . . reasons why [we$_{Buyer}$] need to **pay** [particular attention$_{Seller}$] [to this$_{Goods}$]. (frame: `Commerce_pay`)

Wrong senses occur only with verbs in the analyzed data. There are two support verbs, at least 10 occurrences of *be* as copula but tagged as evoking `Categorization`, and two wrong sense assignments for other reasons.

Differences between LUs with respect to their core FEs: FrameNet frames can host LUs of different parts of speech. Sometimes, a frame contains non-verbal LUs that under normal circumstances *cannot* realize a core FE. An example is the LU *today* in the `Calendric_unit` frame. As opposed to nouns such as *evening* (47), the expression of the core FE Whole is not normally possible with *today* (48).

(47) He doesn't normally work on [Sunday$_{Whole}$] **evenings** . . . . (FrameNet annotation)

(48) Considering that it is only **today** that we are dealing with a . . . proposal . . . . (automatic annotation of EuroParl corpus: `Calendric_unit` frame with FE Whole not annotated)

Fortunately, this phenomenon is rare. The reported example is the only case found in the manually analyzed data.

By checking for the presence of all core FEs in the annotation output, a post-processing procedure might thus find some instances of wrong sense assignment. However, in current AFSA systems, missing core FEs do not necessarily indicate that the word sense is wrong. In particular, to facilitate a more reliable detection of wrong senses, the frame-semantic annotation system should try to detect null-instantiated FEs and address more complicated syntactic constructions.

6 Future extensions of the FrameNet model and annotation process

FrameNet researchers are constantly updating the inventory of frames, frame relations, LUs, and annotations in the database. The next data release will contain more conceptual information for aligning multi-lingual lexicons, as well as additional English lexical data and gold standard annotation that can be used to develop automatic frame-semantic annotation systems. Directions for future extensions of the FrameNet model and for complementing the manual annotation process are suggested by several initiatives that are now being tested at FrameNet.

6.1 Extending the FrameNet model

The following improvements are being studied to extend the model: creating a 'constructicon,' with definitions and annotations of syntactic constructions; alignment

https://doi.org/10.1017/S1351324909005117 Published online by Cambridge University Press

of WordNet and FrameNet; and making the definitions of frames and FEs more structured.

Syntactic constructions: In many natural language understanding (NLU) systems, e.g., question answering applications, there is an assumption that the semantics represented by the annotations on various parts of a full-text annotated sentence will be composed by the operation of a parser into a representation of the sentence as a whole. But one frequently encounters bits of syntax that are not handled correctly by current parsers, constructions which are outside of the 'core syntax' of English, yet common enough to impede NLU. Some of the constructions contain specific lexical material, such as 'just because X doesn't mean Y' (Bender and Kathol 2001), and 'What's X doing Y' (Kay and Fillmore 1999), and others do not, such as '-ed compound adjectives' (*long-legged folk, leather-covered sofas*) and rate expressions (*100 km. an hour, $5 a pair*). Many have been discussed in linguistic literature. A pilot study called 'Beyond the core' is now underway to annotate and document some of these syntactic constructions in a way analogous to FrameNet's treatment of LUs. Examples of the constructions are extracted from corpora, their **construction elements** are identified and labeled, and the syntactic and semantic restrictions on their fillers are carefully described.

WordNet-FrameNet alignment: Another study is underway on the question of how FrameNet and WordNet can be aligned, so that users will be better able to combine the often complementary information from these two resources (Fellbaum and Baker 2008). WordNet has a vast taxonomy of nouns, while FrameNet has very few non-event nouns, as they do not have a very 'interesting' frame semantics; e.g. FrameNet contains *destruction* and *application*, but not *table, chair*, or *cat*. Also, WordNet has more-or-less separate hierarchies for nouns, verbs, adjectives, and adverbs, while FrameNet frames include all parts of speech. Completely merging the resources would create an unwieldy mess, but properly aligning them should improve the usefulness of both. Several algorithms are being tested at FrameNet to measure the semantic distances between FrameNet LUs and WordNet word senses, and to propose mappings between them.

Structure of FN frame definitions: Finally, a number of problems arise because the FrameNet frame and FE definitions are rather unstructured. For example, a number of researchers have mined them to measure semantic distance, but have not always been able to separate the examples from the definition proper – clearly the two parts should be handled differently. Likewise, for those developing or using FrameNets for languages other than English, it will be difficult to automatically distinguish frame-changing modifications from modifications of supportive material, as long as language-specific example sentences and syntactical pattern descriptions in the definitions cannot be reliably detected. In the future, FrameNet might add more structure to the definitions.

6.2 Extending the annotation process

To add full-text FrameNet annotation to a subsection of the ANC, some of the texts will continue to be manually annotated in the usual FrameNet manner, but

https://doi.org/10.1017/S1351324909005117 Published online by Cambridge University Press

an additional goal is to train an ASRL system to label others. The method involves training several existing ASRL systems on the current FrameNet data, comparing their output on the new texts, finding where they disagree (or the confidence in their output is low), manually annotating more examples from the region of the errors, and then retraining the ASRL systems with the new data.

The present article, in conjunction with these on-going efforts, demonstrates that the FrameNet model is becoming more formal and complete, the relation of FrameNet to other NLP resources (English and non-English) is becoming better defined, and its coverage of word senses and the amount of gold-standard annotation are both increasing. All this will continue to improve the contribution FrameNet can make to any NLP research and development project which requires a rich semantic representation.

References

Baker, Collin. 2008. FrameNet, present and future. In *Proceedings of First International Conference on Global Interoperability for Language Resources (ICGL-2008)*, 12–17, Hong Kong, January.

Baker, Collin, Ellsworth, Michael and Erk, Katrin. 2007. SemEval-2007 Task 19: frame semantic structure extraction. In *Proceedings of the Fourth International Workshop on Semantic Evaluations (SemEval-2007)*, 99–104, Prague, Czech Republic.

Baker, Collin F., Fillmore, Charles J. and Cronin, Beau. 2003. The structure of the FrameNet database. *International Journal of Lexicography* **16**(3): 281–296.

Bejan, Cosmin Adrian and Hathaway, Chris. 2007. UTD-SRL: A pipeline architecture for extracting Frame Semantic Structures. In *Proceedings of the Fourth International Workshop on Semantic Evaluations (SemEval-2007)*, 460–463, Prague, Czech Republic.

Bender, Emily M. and Kathol, Andreas. 2001. Constructional effects of 'Just because... doesn't mean'. In *Proceedings of the 27th Annual Meeting of the Berkeley Linguistics Society*, 13–25, Berkeley, CA.

Boas, Hans C. 2005. Semantic Frames as interlingual representations for multilingual lexical databases. *International Journal of Lexicography* **18**(4): 445–478.

Burchardt, Aljoscha and Frank, Anette. 2006. Approximating textual entailment with LFG and FrameNet frames. In *Proceedings of the second PASCAL Recognizing Textual Entailment Workshop*, 92–97, Venice, Italy.

Burchardt, Aljoscha, Erk, Katrin, Frank, Anette, Kowalski, Andrea, Padó, Sebastian and Pinkal, Manfred. 2006. The SALSA corpus: a German corpus resource for lexical semantics. In *Proceedings of the Fifth International Conference on Language Resources and Evaluation (LREC-06)*, 969–974, Genoa, Italy.

Cruse, D. Alan. 1986. *Lexical Semantics*. Cambridge, England, UK: Cambridge University Press.

Dang, Trang Hoa, Kipper, Karin, Palmer, Martha and Rosenzweig, Joseph. 1998. Investigating regular sense extensions based on intersective Levin classes. In: *Proceedings of the 36th Annual Meeting of the Association for Computational Linguistics and 17th Conferences on Computational Linguistics (COLING-ACL 1998)*, 293–99, Montréal, Quebec, Canada.

Dorr, Bonnie. 1994. Machine translation divergences: a formal description and proposed solution. *Computational Linguistics* **20**(4): 597–633.

Dowty, David R. 1991. Thematic proto-roles and argument selection. *Language* **67**: 547–619.

Ellsworth, Michael and Janin, Adam. 2007. Mutaphrase: paraphrasing with FrameNet. In *Proceedings of the ACL-PASCAL Workshop on Textual Entailment and Paraphrasing*, 143–150, Prague, Czech Republic.

https://doi.org/10.1017/S1351324909005117 Published online by Cambridge University Press

Erk, Katrin. 2005. Frame assignment as word sense disambiguation. In *Proceedings of the sixth International Workshop in Computational Semantics (IWCS 6)*, 362–364, Tilburg, The Netherlands.

Erk, Katrin. 2006. Unknown word sense detection as outlier detection. In *Proceedings of the Human Language Technology Conference of the North American Chapter of the ACL (HLT-NAACL 2006)*, 128–135, New York City, NY.

Erk, Katrin and Padó, Sebastian. 2006. Shalmaneser - a toolchain for shallow semantic parsing. In *Proceedings of the Fifth International Conference on Language Resources and Evaluation (LREC-06)*, 527–532, Genoa, Italy.

Erk, Katrin and Padó, Sebastian. 2007. User's guide to the Shalmaneser system. `http://www.coli.uni-saarland.de/projects/salsa/shal/`.

Fellbaum, Christiane and Baker, Collin F. 2008. Can WordNet and FrameNet be made "interoperable"? In *Proceedings of The First International Conference on Global Interoperability for Language Resources (ICGL-2008)*, 67–74, Hong Kong.

Fillmore, Charles J. 1968. The case for case. In: Emmon Bach and Robert T. Harms (eds.), *Universals in Linguistic Theory*. New York: Holt, Rinehart and Winston.

Fillmore, Charles J. 1976. Frame semantics and the nature of language. *Annals of the New York Academy of Sciences* **280**: 20–32.

Fillmore, Charles J. 1985. Frames and the semantics of understanding. *Quaderni di Semantica* **6**(2): 222–254.

Fillmore, Charles J. and Baker, Collin F. 2001. Frame semantics for text understanding. In *Proceedings of the NAACL-2001 Workshop on WordNet and Other Lexical Resources: Applications, Extensions and Customizations*, Pittsburgh, PA.

Fillmore, Charles J., Johnson, Christopher R. and Petruck, Miriam R. L. 2003. Background to FrameNet. *International Journal of Lexicography* **16**(3): 235–250.

Francopoulo, Gil. 2005. Extended examples of lexicons using LMF (auxiliary working paper for LMF). Technical Report. INRIA-Loria.

Francopoulo, Gil, George, Monte, Calzolari, Nicoletta, Monachini, Monica, Bel, Nuria, Pet, Mandy and Soria, Claudia. 2006. Lexical markup framework (LMF). In *Proceedings of the Fifth International Conference on Language Resources and Evaluation (LREC-2006)*, 233–236, Genoa, Italy.

Gildea, Daniel and Jurafsky, Daniel. 2002. Automatic labeling of semantic roles. *Computational Linguistics* **28**(3): 245–288.

Graehl, Jonathan, Knight, Kevin and May, Jonathan. 2008. Training tree transducers. *Computational Linguistics* **34**(3): 391–427.

Ide, Nancy and Suderman, Keith. 2007. GrAF: A graph-based format for linguistic annotations. In *Proceedings of the ACL-2007 Workshop on Linguistic Annotation*, 1–8, Prague, Czech Republic.

Ide, Nancy, Reppen, Randi and Suderman, Keith. 2002. The American National Corpus: More than the web can provide. In *Proceedings of the Third Language Resources and Evaluation Conference (LREC-2002)*, 839–844, Las Palmas, Canary Islands, Spain.

Johansson, Richard and Nugues, Pierre. 2006. A FrameNet-based semantic role labeler for Swedish. In *Proceedings of the COLING/ACL-06 Main Conference Poster Sessions*, 436–443, Sydney, Australia.

Johansson, Richard and Nugues, Pierre. 2007. LTH: Semantic structure extraction using nonprojective dependency trees. In *Proceedings of the Fourth International Workshop on Semantic Evaluations (SemEval-2007)*, 227–230, Prague, Czech Republic, June.

Kay, Paul and Fillmore, Charles J. 1999. Grammatical constructions and linguistic generalizations: The What's X Doing Y? construction. *Language* **75**: 1–33.

Kingsbury, Paul, Palmer, Martha and Marcus, Mitch. 2002. Adding semantic annotation to the Penn TreeBank. In *Proceedings of the Human Language Technology Conference (HLT-2002)*, San Diego, CA.

https://doi.org/10.1017/S1351324909005117 Published online by Cambridge University Press

Kipper, Karin, Dang, Hoa Trang and Palmer, Martha. 2000. Class-based construction of a verb lexicon. In *Proceedings of the Seventeenth National Conference on Artificial Intelligence (AAAI-2000)*, Austin, TX.

Koehn, Philipp. 2005. Europarl: A parallel corpus for Statistical Machine Translation. In *Machine Translation Summit X*, 79–86, Phuket, Thailand.

Lambrecht, Knud. 1994. *Information structure and sentence form. topic, focus and the mental representations of discourse referents*. Cambridge Studies in Linguistics, no. 71. Cambridge, England, UK: Cambridge University Press.

Levin, Beth. 1993. *English verb classes and alternations: A preliminary investigation*. Chicago, IL: University of Chicago Press.

Lönneker-Rodman, Birte. 2007a. Beyond syntactic valence: FrameNet markup of example sentences in a Slovenian-German online dictionary. In Levická, Jana and Garabík, Radovan (eds.), *Slovko 2007. Computer Treatment of Slavic and East European Languages. Proceedings of the Fourth International Seminar, Bratislava, Slovakia*., 152–164, Brno, Czech Republic, October. Tribun.

Lönneker-Rodman, Birte. 2007b. Multilinguality and FrameNet. Technical Report TR-07-001. International Computer Science Institute, Berkeley, CA.

Màrquez, Lluís, Carreras, Xavier, Litkowski, Kenneth C. and Stevenson, Suzanne (eds). 2008. *Special Issue on Semantic Role Labeling*. Computational Linguistics, vol. 34. MIT Press for the Association for Computational Linguistics.

Meredith, Daniel N. and Pieper, Jan H. 2006. BETA: better extraction through aggregation. In: *Proceedings of the SIGIR-2006 Workshop on Faceted Search*, Seattle, WA.

Meyers, Adam, Reeves, Ruth, Macleod, Catherine, Szekely, Rachel, Zielinska, Veronika, Young, Brian and Grishman, Ralph. 2004. Annotating noun argument structure for NomBank. In *Proceedings of the Fourth International Conference on Language Resources and Evaluation (LREC-2004)*, 803–806, Lisbon, Portugal.

Narayanan, Srini and Harabagiu, Sanda. 2004. Question Answering based on semantic structures. In *Proceedings of the 20th International Conference on Computational Linguistics (COLING 2004)*, 693–701, Geneva, Switzerland.

Ohara, Kyoko Hirose. 2008. Lexicon, grammar, and multilinguality in the Japanese FrameNet. In *Proceedings of the 6th International Conference on Language Resources and Evaluation (LREC-08)*, Marrakech, Morocco.

Ohara, Kyoko Hirose, Fujii, Seiko, Ohori, Toshio, Suzuki, Ryoko, Saito, Hiroaki and Ishizaki, Shun. 2004. The Japanese FrameNet project: An introduction. In *Proceedings of the LREC-04 Workshop on Building Lexical Resources from Semantically Annotated Corpora*, 9–11, Lisbon, Portugal.

Padó, Sebastian. 2007. Translational equivalence and cross-lingual parallelism: The case of FrameNet frames. In *Proceedings of the NODALIDA-07 Workshop on Building Frame Semantics Resources for Scandinavian and Baltic Languages*, 39–46, Tartu, Estonia.

Palmer, Martha, Gildea, Dan and Kingsbury, Paul. 2005. The Proposition Bank: An annotated corpus of semantic roles. *Computational Linguistics* **31**(1): 71–105.

Pradhan, Sameer, Ward, Wayne, Hacioglu, Kadri, Martin, James and Jurafsky, Daniel. 2005. Semantic role labeling using different syntactic views. In *Proceedings of the 43rd Annual Meeting of the Association for Computational Linguistics (ACL-05)*, 581–588, Ann Arbor, MI.

Ruppenhofer, Josef, Ellsworth, Michael, Petruck, Miriam R. L., Johnson, Christopher R. and Scheffczyk, Jan. 2006. *FrameNet II: Extended theory and practice*. `http://framenet.icsi.berkeley.edu/`.

Sato, Hiroaki. 2008. New functions of FrameSQL for multilingual FrameNets. In *Proceedings of the 6th International Conference on Language Resources and Evaluation (LREC-2008)*, Marrakech, Morocco.

https://doi.org/10.1017/S1351324909005117 Published online by Cambridge University Press

Scheffczyk, Jan, Pease, Adam and Ellsworth, Michael. 2006a. Linking FrameNet to the Suggested Upper Merged Ontology. In *Proceedings of the International Conference on Formal Ontology in Information Systems (FOIS-2006)*, Baltimore, MD.

Scheffczyk, Jan, Baker, Collin F. and Narayanan, Srini. 2006b. Ontology-based reasoning about lexical resources. In *Proceedings of the LREC-2006 Workshop on Interfacing Ontologies and Lexical Resources for Semantic Web Technologies (OntoLex-2006)*, 1–6, Genoa, Italy.

Shen, Dan and Lapata, Mirella. 2007. Using semantic roles to improve question answering. In *Proceedings of the 2007 Joint Conference on Empirical Methods in Natural Language Processing and Computational Natural Language Learning (EMNLP-CoNLL)*, 12–21, Prague, Czech Republic.

Sinha, Steve and Narayanan, Srini. 2005. Model based answer selection. In *Proceedings of the AAAI-2005 Workshop on Textual Inference for Question Answering*, 24–31, Pittsburgh, PA.

Subirats, Carlos and Sato, Hiroaki. 2004. Spanish FrameNet and FrameSQL. In *Proceedings of LREC-2004 Workshop on Building Lexical Resources from Semantically Annotated Corpora*, 13–16, Lisbon, Portugal.

Subirats-Rüggeberg, Carlos and Petruck, Miriam R. L. 2003. Surprise: Spanish FrameNet! In *Proceedings of the Workshop on Frame Semantics, XVII International Congress of Linguists (CIL)*, Prague, Czech Republic.

Surdeanu, Mihai, Johansson, Richard, Meyers, Adam, Màrquez, Lluís and Nivre, Joakim. 2008. The CoNLL-2008 shared task on joint parsing of syntactic and semantic dependencies. In *Proceedings of the 12th Conference on Computational Natural Language Learning (CoNLL-2008)*, 159–177, Manchester, UK.

Vossen, Piek. 1999. EuroWordNet General Document. Version 3. Technical Report. University of Amsterdam.

https://doi.org/10.1017/S1351324909005117 Published online by Cambridge University Press