

FrameNet as a “Net”

Charles J. Fillmore*, Collin F. Baker*, and Hiroaki Sato†

*International Computer Science Institute
Berkeley, California, U.S.A.
{fillmore,collinb}@icsi.berkeley.edu

†Senshu University
Kawasaki, Japan
hiroaki@isc.senshu-u.ac.jp

Abstract

While FrameNet does not record the range of semantic relations found in thesaurus-style lexical resources like WordNet, it does provide a number of ways in which lexical units (LUs) can be seen as related to each other. This paper characterizes and motivates the networks of frame-to-frame relations that are being built on FrameNet’s frames, and introduces additional representational mechanisms needed for showing similarities among LUs that are independent of frame membership. LU-to-LU relationships are shown by shared membership in a single frame, by membership in frames that are themselves related to each other, and by shared semantic type information.

1. Introduction

FrameNet (FN) is a corpus-based¹ lexicon-building project that documents the links between lexical items and the semantic frame(s) they evoke; it accomplishes this by annotating sets of sentences that exemplify the items being described, and performing various operations on the resulting annotations. The basic units in FN descriptions are the **frame** and the **lexical unit (LU)**, the latter understood as the pairing of a “word” with just one of its meanings;² thus, a word with four meanings is treated as four lexical units. In most cases, for a word to have more than one meaning implies that it belongs to more than one frame.³

The main components of the FrameNet database are (1) the frame ontology, (2) the set of annotated sentences, and (3) the set of lexical entries. The basis of the ontology is the set of frames, each of which consists of an informal characterization of a situation type (the frame definition), together with a collection of **frame elements (FEs)**. The FEs are the semantic roles of the entities involved in each frame. FE names are used as labels for the words or phrases that are in grammatical construction with the LUs that evoke that particular frame. For example, the frame that includes the English verb *inform* has as its core FEs SPEAKER, ADDRESSEE and MESSAGE.

The example sentences are selected by FrameNet annotators as representing the typical uses of the LUs belonging to individual frames. Each set of annotations is centered around a particular LU; the sentence’s constituents are labeled (with FE names) according to the ways in which they fill in information about the frame. For example, sentences

(1) and (2) have SPEAKER appearing as subject, and ADDRESSEE as object; the MESSAGE FE appears as a *that*-clause in sentence (1), and as an event-naming nominalization introduced by *of* in sentence (2).

- (1) [SPEAKER We] **informed** [ADDRESSEE the press]
[MESSAGE that the prime minister has resigned]
- (2) [SPEAKER We] **informed** [ADDRESSEE the press]
[MESSAGE of the prime minister’s resignation]

The **lexical entry** for each LU is a summary of what has been recorded in its annotations, presented as **valence descriptions**, showing all the ways in which its frame elements can be realized, such as the alternative syntactic realizations of the MESSAGE just shown for the verb *inform*. The collection of annotated sentences is made available in the database as evidence for the analysis.

To date not much attention has been paid to the second syllable in the name “FrameNet”—the part that suggests a network of relations. The first and most obvious way in which LUs are related to each other is through membership in the same frame. Thus *inform* shares a frame with the verbs *notify* and *announce*, and also with the nouns *notification* and *announcement*, and the verb *resign* shares frame membership with its nominal partner *resignation*, and with verbal expressions like *abdicate*, *step down* and *stand down*.⁴ But LUs can also be related to each other in other ways, either because their frames are related to other frames, or through semantic properties (called **semantic types** in the FN database) assigned to LUs individually rather than through their frames.

¹Our corpus is the British National Corpus, used with permission of Oxford University Press, and, more recently, newswire text from the Linguistic Data Consortium.

²The term is from Cruse, 1986).

³Effectively this does not distinguish between **polysemy** and **homonymy**; FrameNet researchers are interested in the semantic links connecting the senses of a polysemous word (metaphor, metonymy, generalizing, narrowing, etc.—see Fillmore and Atkins, 2000) on the polysemy of *crawl*), but currently lack the means of identifying and displaying such relations.

⁴Membership in a single frame is not equivalent to belonging to the same WordNet **synset**. Frame members do not need to belong to a single part of speech, and they can include words that express different values on a single scale, hence even antonymous words. Thus, such obviously non-synonymous pairs as *old* and *young* share frame membership, as do *heavy* and *light*, or *tall* and *short*.

2. Representation Mechanisms

We begin with a discussion of semantic types, and continue with three frame-to-frame relations, **inheritance**, **using** and **subframe**.

2.1. Semantic types

The FrameNet database allows the assignment of semantic types to LUs, FEs and frames. The perception verbs *hear* vs. *listen* are distinguished as passive versus active perception verbs, and so, respectively, are *see* vs. *look*. Hearing and seeing are things that happen to you, listening and looking are things that you do, and this difference is considered important enough to merit entry into separate frames.

In the FN database, *hear* and *see* and the passive-perception uses of other sensory words, such as *feel*, *taste* and *smell*, belong to the Perception_experience frame; the verbs *look* and *listen* belong to the Perception_active frame, along with the corresponding active uses of *feel*, *taste* and *smell*.⁵ A cross-cutting grouping of these verbs according to their sensory modality is handled indirectly, through markers of **semantic type** for Vision, Hearing, Smell, Taste and Touch.⁶

2.2. Inheritance

We define frame inheritance as an IS-A relation between a parent frame and a child frame which includes full inheritance of FEs and their semantic types. This means that if the parent frame has a semantic type, the child frame must have the same semantic type or a subtype (elaboration) of it. Also, for each FE in the parent frame there must be an FE in the child frame of the same semantic type or a subtype thereof. The FEs of the child may or may not have the same names as the FEs of the parent, and there may be additional FEs in the child for which there is no corresponding FE in the parent. Furthermore, if the parent frame has subframes (basically subevents, discussed below), its subframe structure is also inherited (and possibly elaborated) by the child frame. This is complete, monotonic inheritance.

2.3. Using

Because we found a number of relations among frames which do not quite fit the criteria for full inheritance, we have defined a second type of relation, similar to inheritance, not requiring full mapping of FEs from parent to child, or complete inheritance of the parents' subframe

⁵Examples of the passive vs. active uses of the polysemous verbs are *I just felt something touch me* vs. *Feel his leg to see if it's broken*; *Do you smell something?* vs. *Take time to smell the roses*; *I think I can taste rum in this mixture* vs. *Would somebody taste the soup to see if it's got enough salt?*

⁶The modality features are used to account for certain distributional properties of the perception verbs that are related to the difference between **distant** senses (Vision and Smell), **contact** senses (Taste and Touch), and **chemical** senses (Smell and Taste). Not all verbs in the Perception_experience frame bear a modality-specifying semantic type: *perceive* and *notice*, for example, are neutral in this respect.

structure. Like inheritance, there can be multiple using relations, so that a child may inherit and/or use multiple parents.

Defining such inheritance relations creates a lattice of frames, a directed acyclic graph. Working out the full details of multiple inheritance through several levels of the lattice can be complex and time-consuming, and we have not created all the links which we would like, but there are currently roughly 130 inheritance links and 160 using links in the FN database.

2.4. Subframes

Subframes are used for representing subevents; frames that represent complex processes have subframes representing their subparts.⁷ To take a simple example, the Motion_scenario frame has three subframes, Departing, Motion, and Arriving. In this case, the subframes are temporally ordered, but in general, subframes need not be a completely ordered with respect to each other. For example, the Commercial_transaction frame has two subframes Commerce_goods-transfer and Commerce_money-transfer, but these are not ordered with respect to each other. In some commercial transactions, you pay in advance, in others, only after receiving the goods or services.

2.5. The Implementation of Frame-to-Frame Relations in the Database

The conceptual relations described above are implemented as a relational database, using MySQL. So far as possible, the tables of the database and the links between them directly mirror the conceptual structure. For example, there is a table for frames and another for FEs, with a one-to-many relation between them. The lemma table is linked to the frame table via the lexical unit table, each entry of which has a pointer to a lemma and a pointer to a frame; this is a many-to-many relation—frames typically include many lemmas and the same lemma can appear in several frames, representing polysemy or homonymy. The higher-order relations are handled similarly. For a full discussion of the database structure, see Baker et al., 2003.

3. Using the Relations

There are two broad uses of the relations we have been discussing: (1) capturing linguistic generalizations, and (2) projecting inferences from frame-annotated sentences.

3.1. Capturing Generalizations

Capturing linguistic generalizations requires defining the features or structures with respect to which the generalizations can be stated. If LUs across frames are assigned the same semantic types, there should be something that these LUs have in common, where the generalization is expressed with reference to the shared semantic type. If some frames are seen as subtypes of another frame, then statements about the more general frame will be applicable to the more specific frames.

⁷Note that **subtypes** of frames are represented by inheritance and using relations, which are quite distinct from the **subframe** relation.

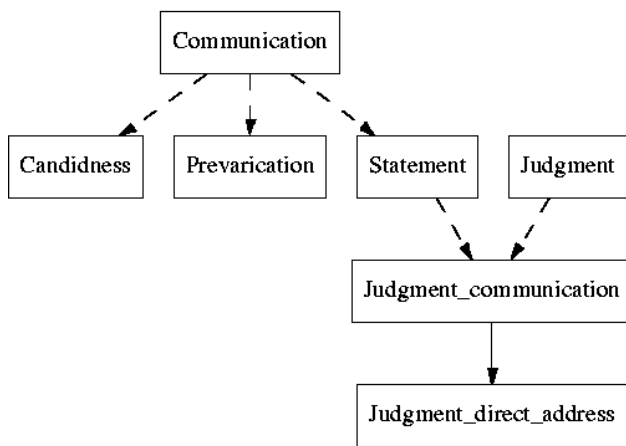


Figure 1: Examples of Frame Relations

In Figure 1 we see Judgment_communication as using both Judgment and Statement, which in turn uses Communication (shown by the dashed arrows). We also see Judgment_direct_address inheriting from (i.e. a subtype of, shown by the solid arrow) Judgment_communication. Not shown in the figure is the information that the FEs of the mother, in each case, are bound to particular FEs in the daughter.

For example, the subtypes of the Communication frame (with LUs like *communicate*, *signal*) all presuppose a situation in which someone is conveying wishes or information to someone else. The subtypes of the Judgment frame (LUs *admire*, *critical*, *contempt*) involve situations in which someone is evaluating something or someone else. Communicating such judgments is a combination of these two situations, where the Communication.COMMUNICATOR is linked to the Telling.SPEAKER, and thence to the Judgment.COGNIZER, captured in the Judgment_communication frame (LUs *condemn*, *extol*, *scathing*). This is in turn used by a frame called Judgment_direct_address (LUs *chastise*, *compliment*, *scold*) in which the EVALUÉE is also identified with the ADDRESSEE. The frame-to-frame relations are shown in Figure 1

Note that the LUs in the Judgment frame and its descendants are also further differentiated by semantic types for Positive_judgment and Negative_judgment. (In Judgment, *value* vs. *condemn*; in Judgment_communication, *praise* vs. *criticize*; in Judgment_direct_address, *flatter* vs. *scold*.) The main motivation behind the assignment of semantic types to LUs is the assumption that in some part of the grammar of the language, the LUs bearing the type participate in some generalization. For example, we say that words designating colors have a color feature, meaning that they participate in the Color_qualification frame whose members are adjectives like *light*, *pale*, *dark*, *deep*, etc. This allows us to recognize that color adjectives in English are regularly modifiable by adjectives, and we need a way of stating this.

In a number of frames indicating scalar or gradable concepts there is a distinction between **plain** and **end-of-scale** gradables. Examples of the contrast are shown in

Dimension	Plain	End of Scale
SIZE	large, big	enormous, immense
SIZE	small	minuscule, tiny
BEAUTY	attractive, pretty	gorgeous, ravishing
TASTE	tasty	delicious
INTEREST	interesting	fascinating, enchanting

Table 1: Semantic Features for Scalar Adjectives

Table 1; the elements on the right are of semantic type **end-of-scale**. Although these word groups belong to different frames, they have something in common that is independent of their frame membership. They participate as arguments in modification structures expressing the Extreme_value frame whose LUs include *absolutely*, *utterly*, *perfectly* and a few others.

Polar evaluative semantic types, which could be called simply Negative and Positive, are needed for describing various contexts that preferentially select expressions of desirable or undesirable phenomena. The verb *cause* is famously⁸ the creator of contexts favoring negatively evaluated concepts like *disaster*, *infection*, *death*, *injury*.⁹

As we have seen with Communication and Judgment, **multiple inheritance** can show how the general properties of two independently-needed frame can apply to a single daughter frame. Multiple inheritance also covers situations in which one frame provides information about what is happening in the world, and the other presents a particular point of view on the event and determines certain aspects of the structuring of the clause, such as which frame elements are realized in the nucleus of a clause (as subject or direct object) and which are expressed obliquely or left unexpressed. For example, sentences about buying and selling can give the same information about a particular exchange event, but the buying frame is simultaneously an acquiring frame and the selling frame is simultaneously a giving frame. Buying something from somebody is getting something from somebody; selling something to somebody is giving something to somebody, both in terms of syntax and of semantic profiling. Thus, in sentences about buying, selling, paying, spending, charging, etc., inferences about who has what before and after the transaction are more dependent on who is identified as the buyer or seller, what the goods and the payment are, than on who is seen as the agent and what is seen as the theme.

3.2. Projecting Inferences

The more “situational” frames in such relationships can shade into the kinds of cognitive structures that have been treated in a number of fields under such names as **schema** (Bartlett, 1932; Piaget, 1952), **frame** (Minsky, 1975; Goffman, 1974; Tannen, 1985), **script** (Schank and Abelson, 1977), **idealized cognitive models** (Lakoff, 1987, 68-76 & ff.) and **memes** (Dawkins, 1976). Some frames identify events or states of affairs that are seen as part of larger situation types, and evoke in the mind of the language user an

⁸cf. Stubbs, 1995a; Stubbs, 1995b; Sinclair, 1991

⁹This is a case where a needed sortal property cannot believably correspond to a node in any of the existing ontologies.

understanding of that larger situation that shapes the coherence of the ongoing discourse. Recording such relationships will provide information usable in tracking events, detecting situational backgrounds, making predictions, etc. These can be referred to as **cognitive frames**, though this is not a term of art recognized in the FN database.

It would be impossible to describe all of the cognitive frames, or even all of those that the lexicon clearly impinges on, in a purely lexicographic project. But for a research project that is reaching toward language understanding and full text semantic parsing within such domains, it would be important to work through a small number of such frames characteristic of a particular class of domain types, if only to pave the way for others wishing to adapt FrameNet methods to specialist language.

The relation between linguistic frames and background frames are of the types we refer to as sub-frame and using, typically both. Tipping, for example, is interpreted as one temporal part of a general service script, and is capable of invoking a variety of service contexts—restaurants, taxi rides, hotel services, etc., where the relationship could be spoken of as of the using type.

There are many nouns in English that designate a “gift” or a “giving” act in some close-up event where one person gives (usually) money to another person within a particular cultural setting, but the individual nouns trigger or evoke very different experiential or institutional settings: among them are *tip*, *bribe*, *ransom*, *honorarium*, *fine*, and *tax*, each of which can be used in a support construction (with *pay* or *give*, for example) to profile the moment of money transfer, while separately evoking very different schematic histories and outcomes. At play here is not merely the cultural fact that money payments are expected or demanded under particular conditions, but that certain lexical units simultaneously identify¹⁰ the payment event and evoke the setting.

The FrameNet team has undertaken a fine-grained analysis of a number of concepts in Criminal process, though still more or less restricted to the level of newspaper reports of crimes and the law. One of our goals in this work has been to be able to make predictive and postdictive inferences based on lexically tagged events. Arraignment is a temporal portion of a full script of criminal justice procedure, presupposing arrest and expecting trial and adjudication, but also presupposing the general cultural expectations associated with wrong-doing and rewards-and-punishment. Someone who knows the meaning of the word will conclude from a report that X has been *arraigned* (1) that X is engaged in a criminal procedure (an “outward” inference), (2) that X was suspected of committing a crime and has been arrested (a “backward” inference), (3) that a report of the event should mention a criminal charge, the setting of bail, the setting of a trial date, the plea (“inward” inferences), and (4) that X can be expected to participate in fu-

ture events within the same scenario, including a trial with its own set of possible outcomes (a “forward” inference).

4. References

- Baker, Collin F., Charles J. Fillmore, and Beau Cronin, 2003. The structure of the FrameNet database. *International Journal of Lexicography*, 16(3).
- Bartlett, Frederick, 1932. *Remembering*. Cambridge, U.K.: Cambridge University Press.
- Cruse, D.A., 1986. *Lexical Semantics*. Cambridge Textbooks in Linguistics. Cambridge and New York: Cambridge University Press.
- Dawkins, Richard, 1976. *The Selfish Gene*. New York: Oxford University Press.
- Fillmore, Charles J. and Sue Atkins, 2000. Describing polysemy: the case of ‘crawl’. In Yael Ravin and Claudia Leacock (eds.), *Polysemy: Theoretical and Computational Approaches*. Oxford: Oxford University Press, pages 91–110.
- Goffman, Erving, 1974. *Frame Analysis: An Essay on the Organization of Experience*. New York: Harper and row.
- Lakoff, George, 1987. *Women, Fire, and Dangerous Things*. Chicago: University of Chicago Press.
- Langacker, Ronald, 1987. *Foundations of Cognitive Grammar*, volume I. Stanford University Press.
- Minsky, Marvin, 1975. A framework for representing knowledge. In Patrick Winston (ed.), *The Psychology of Computer Vision*. McGraw-Hill.
- Piaget, Jean, 1952. *The Origins of Intelligence in Childre*. New York: International Universities Press.
- Schank, Roger C. and Robert P. Abelson, 1977. *Scripts, Plans, Goals and Understanding: an Inquiry into Human Knowledge Structures*. Hillsdale, NJ: Lawrence Erlbaum.
- Sinclair, John, 1991. *Corpus, Concordance, Collocations*. Oxford: Oxford University Press.
- Stubbs, Michael, 1995a. Collocations and semantic profiles: on the cause of the trouble with quantitative studies. *Functions of Language*, 2(1).
- Stubbs, Michael, 1995b. Corpus evidence for norms of lexical collocation. In Guy Cook and Barbara Seidlhofer (eds.), *Principle and Practice in Applied Linguistics. Studies in Honour of H. G. Widdowson*, chapter 16. Oxford: Oxford University Press.
- Tannen, Deborah, 1985. Frames and schemas in interaction. *Quaderni di Semantica*, 6(2):326–355. Round Table discussion on frame/script semantics, ed. by Victor Raskin.

¹⁰Here we might use the Langacker term **profile** (Langacker, 1987), a term he introduces with the word *hypotenuse*. Everything known about an act of bribery would be the same even if there were no special term for the money that gets paid, just as everything known about right angle triangles would be the same even if we only had an analytic way to identify the side opposite the right angle.