Frame Semantics for Text Understanding

Charles J. Fillmore and **Collin F. Baker**
International Computer Science Institute
1947 Center St.,Suite 600
Berkeley, California, 94704
{fillmore,collinb}@icsi.berkeley.edu*

Abstract

An introduction to knowledge representation using Frame Semantics, as is being carried out in the FrameNet Project. A short news article is analyzed, providing examples of many of the questions being dealt with and the proposed solutions, including semantic composition, text coherence, polysemy and WSD, and evidentiality.

1 Introduction

The FrameNet(FN) research project (Lowe et al., 1997; Baker et al., 1998; Fillmore et al., 2001) is building a lexical resource that aims to provide, for a significant portion of the vocabulary of contemporary English, a body of semantically and syntactically annotated sentences from which reliable information can be reported on the valences or combinatorial possibilities of each item targeted for analysis. Key aspects to the work of the project are a commitment to a descriptive framework based on semantic frames containing frame elements (semantic roles)[1], and a commitment to documenting its observations on the basis of carefully annotated attestations taken from large electronic corpora.[2]

The basic FrameNet data are stored in a MySQL database a portion of which is shown in Fig. 1. Most significant for our purposes are the tables showing the relation between lemmata and frames (polysemy is a one-many relation between lemmata and the frames that express their meanings), and tables showing the relations between frames. Frame-to-frame relations include (1) composition, by which a complex frame is shown to be decomposable as a temporal structure - often a structured procedural sequence of simpler frames - and (2) inheritance, by which a single frame can be seen as an elaboration of one or more other frames, with bindings between the elements of co-inherited frames. Lexical entries, including valence descriptions which summarize the attested combinability possibilities, are generated as reports derived algorithmically from the database.

One of the means chosen for demonstrating the relevance of the database to NLP research is a planned pilot effort at bringing FN data to bear on information extraction from newspaper accounts of crimes, criminal behavior, and instances of (low-level) criminal justice procedures. This commits us to selecting terminology that occurs frequently in such reports for detailed analysis touching on all aspects of FN research.

Since FN is a lexicographic project, our concern in its application to research on text understanding has to be limited to its potential service in other sorts of activities. We have to take for granted the existence of independent resources providing syntactic parsing (including mini-grammars for dates, addresses, proper names, titles, institutions, etc., such as FASTUS (Appelt et al., 1993)), as well as anaphora resolution, real-world connections, and discourse coherence. It should be possible for FN data to be called on for assistance with word-sense disambiguation, semantic composition (the integration of information associated with semantic depen-

* ~The authors are Principal Investigator and Project Manager, respectively, of the National Science Foundation sponsored project, FrameNet++. We are grateful to the National Science Foundation for funding the work of the project through two grants, IRI #9618838 "Tools for Lexicon Building" March 1997–February 2000, and ITR/HCI #0086132 "FrameNet++: An On-Line Lexical Semantic Resource and its Application to Speech and Language Technology" September 2000–August 2003. The other Principal Investigators of FrameNet++ are Dan Jurafsky (University of Colorado at Boulder), Srini Narayanan (SRI International), and Mark Gawron (San Diego State University).

[1] Semantic roles have played an important role in NLP for many years, from (Simmons, 1973) and (Schank, 1972), to (Riloff and Schmelzenbach, 1998), and in studies of human sentence processing, e.g. (Trueswell et al., 1994).

[2] For the first part of the project, the British National Corpus was used, courtesy of Oxford University Press; for continuing work, and especially for tasks of the kinds considered in this paper, FN are depending on both the BNC and the corpora of English news texts provided by the LDC (North American News Text Corpus, the Supplement, and AP Worldstream English). The project uses the Corpus Workbench software from Institut für Maschinelle Sprachverarbeitung of the University of Stuttgart for searching the corpora and selecting sentences for annotation. In the first phase of the project, the Alembic Workbench from MITRE was used for annotation; for the second phase, FN staff have written a custom Java front-end to the MySQL database.

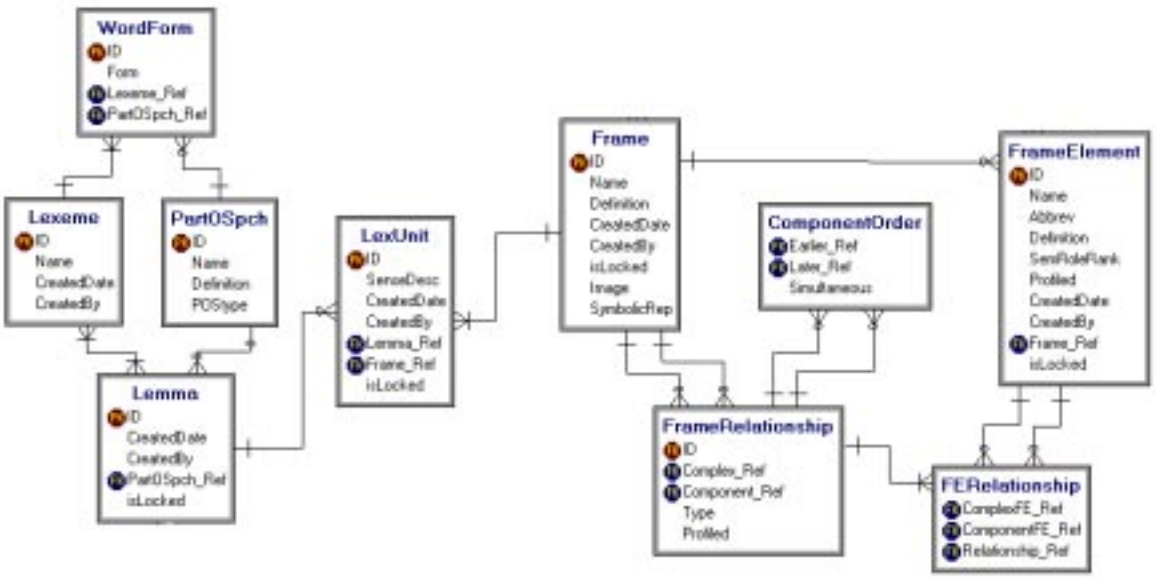

Figure 1: Entities and Relations for FN2 (Partial)

dents into frames evoked by their semantic governors), valence-justified choices among competing parses, and activation of topic-related vocabulary (through shared frame membership) for recognition and sense-selection in successive parts of a text. There are also preliminary studies (Gildea and Jurafsky, 2000) which suggest that the annotation data from FN can also be used to seed automatic recognition of frame elements, both to assist the manual annotation process and to generate a much larger body of annotated sentences.

2 An Example Text

The text we wish to examine here is the following, taken from CNN.com/LAW, dated 14 February 2001. In this text, as is common in journalistic writing, each paragraph contains one sentence. (Sentence numbering is ours.)

1. Washington (CNN)— Alleged White House gunman Robert Pickett was arraigned Wednesday at a federal court in Washington and ordered held without bond.
2. A federal magistrate informed Pickett of the charges against him - assaulting a federal officer with a deadly weapon, which carries a maximum of ten years in prison.
3. The magistrate set a preliminary hearing for next Tuesday and ordered Pickett held without bond.

Index	Dependents	Form	POS
8	none	Alleged	adj
7	none	White House	n
6	7,8	gunman	n
2	[6]	Robert Pickett	n
1	2,3,4,5	was	be
3	(2)9,10	arraigned	ppt
9	none	Wednesday	adv
10	11	at	prp
13	none	a	art
12	none	federal	adj
11	12,13,14	court	n
14	15	in	prp
15	none	Washington	n
4	none	and	cnj
5	(2)16	ordered	ppt
16	(2)17	held	ppt
17	18	without	prp
18	none	bond.	n

Table 1: Dependency Representation of Sentence 1

4. Pickett, who was shot in the knee by the Secret Service after allegedly firing two shots outside the White House, used crutches to walk into the court.
5. He did not enter a plea.

If we assume a simple dependency parse, a first pass in finding the ways in which semantically dependent elements are fitted into the frame struc-

tures evoked by the key words in the text is to trace the connections between governors and their dependents, recognizing the difference between syntactic heads and semantic heads when necessary. A sample of the kind of analysis we will need, for the first sentence, is given in Table 1.

The highest syntactic governor is the passive auxiliary *was*, whose dependents *arraigned* and *ordered* are the main frame-bearers in the sentence. The dependents of *arraigned* in this sentence give us the identity of the defendant (index 2), and the time and space coordinates of the arraignment (indices 9 and 10); the agencies behind the arraignment are unexpressed. The court's action (ordering Picket held without bond) is expressed by *ordered* (index 5) and its dependents. Syntactic interpretations associated with VP conjunction will make clear the role of the defendant in the ordering frame.

Before exploring the possible contributions of FN data, we can make some initial observations about the text as a whole. We notice first of all that the initial sentence informs us of an arraignment. From the meaning of this word we know that relevant activities within such a process involve informing the accused of the charges against him, setting a date for preliminary hearing, and making decisions to guarantee the defendant's future court appearances.

Details of the time of the arraignment and the scheduled time of the preliminary hearing need to be calculated by using the date of the news report (*Wednesday*, *next Tuesday*), together with the knowledge that a simple occurrence of *Wednesday* with past-time reference is understood as the most recent Wednesday, and *next Tuesday* is the Tuesday of the calendar week following the date of the report.

Anaphoric links must be established between *Robert Pickett* (sentence 1), *Pickett* (sentences 2, 3, 4), and *he* (sentence 5) in this text, and between *a federal magistrate* (sentence 2) and *the magistrate* (sentence 3).

The lexical items from the article for which FN needs to have analyses are *alleged*, *gunman*, *arraign*, *Wednesday*, *federal*, *court*, *order*, *hold*, *bond*, *magistrate*, *inform*, *charge*, *assault*, *officer*, *deadly*, *weapon*, *carry*, *maximum*, *year*, *prison*, *set*, *preliminary*, *hearing*, *next*, *Tuesday*, *shoot*, *knee*, *allegedly*, *fire*, *shot*, *outside*, *used*, *crutches*, *walk*, *court*, *enter*, and *plea*.

A general way of applying FN valence information to the analysis of the words in a text is to choose a word (starting from the highest semantically-relevant predicate in a given sentence), determining the frames that it is capable of evoking, noticing the semantic roles of the props and participants in each such frame, trying to match the semantic needs associated with each such frame (and hence with each sense of the word) with phrases found in the sentence at hand, choosing the one which makes the most coherent fit, and entering the semantic structures associated with the dependent constituents into slots provided by the selected frame.

Some of the structures in the text are non-problematic. In Sentence 2, the verb *inform* evokes a frame of one person providing another person with some information, and the syntactic valence possibilities include the pattern by which the communicator is expressed as the verb's subject, a phrase designating the addressee of the informing act follows the verb and the phrase indicating the transmitted message is introduced with an *of*-PP.

An understanding system that seeks coherence between portions of a text will have noted that the first sentence evokes the **Arraignment** frame and that one of the functions of a court appearance for an arraignment is for the judge to inform the accused of the nature of the charges against him; to the extent that the evoked arraignment structure is retained for predicting sentence-to-sentence connections, it seems clear that the mention of *a federal magistrate*, *inform*, and *charges against* can be quickly incorporated into the growing semantic representation of the text. In particular, the meaning selected for *magistrate* (as 'judicial officer' rather than as, say, 'the chief magistrate of the nation') and the meaning chosen for *charges* (as 'accusation' as opposed to 'agreements on delayed payment for merchandise'), will be determined by a coherence preference. It is not the work of FrameNet to line up such expectations, but it should be a service of the FN data to offer descriptions of the related frames and semantic connections among the words (*magistrate* as a synonym of *judge*, *charge* as a synonym of *accusation*, etc.) to facilitate the establishment of such coherence judgments.

It should be noted that frame structures needed for establishing text coherence clearly go beyond mere argument structures. For the **Arraignment** frame, for example, the "slots" needing to be filled are not generally going to be available in the same sentence, let alone among the syntactic dependents of the verb *arraign*.

3 Multi-word Units and Parsing

The analysis cannot simply proceed on the basis of frame information built on the text's words taken one at a time. Many word sequences must be identified as fixed phrases, the most obvious ones including *White House*, *the White House*, *Robert Pickett*, and *the Secret Service*. Tight collocations must be recognized for the following phrases: *held without bond*, *assaulting a federal officer with a deadly weapon*, *preliminary hearing*, *firing shots*, and *enter a plea*. All of these phrases are themselves analyzable from their parts, but *held without bond* is one of the standard

phrases for reporting one of the court decisions in an arraignment hearing, *assaulting a federal officer with a deadly weapon* is a named offense in American law, *preliminary hearing* is a named step in the criminal justice process, and *fire* and *enter* can be analyzed as standard support verbs for *shot* and *plea* respectively. All such information must take the form of separately listed lexical entries. Recognizing the ad hoc occasion-specific compound *White House gunman* and its reference will have to depend on the understander's being in touch with current news.

Sentence 3 speaks of *set(ting) a preliminary hearing*. Either the collocation between *set* and *hearing* or that between *preliminary* and *hearing* (or both) must be established in the lexicon or the understander must depend on knowledge of the steps in an arraignment hearing to choose among the possible senses of the highly polysemous word set. Since the thesaurus-like character of FrameNet is provided by linking words to the frames they evoke, all these words belong to a single (high-level) frame. (The FN lexicon will have to indicate that the valence possibilities for *set* in the **Appointment** frame include the phrasings *set a date for the hearing*, *set the hearing for February 20*, and *set a hearing date*.)

The text offers a few analytic challenges. One of these is in the the appositional relation between *the charges against him* (plural) and *assaulting a federal officer* (singular), and the association with the singular-agreement form *carries*, which violates some rather basic rules of English syntax. Another is the elliptical expression *carries a maximum ten years in prison* where the singular article *a* has to be connected with the grammatically anomalous *maximum ten years* (which appears to be plural). This could be seen as an ellipsis of *a maximum sentence of ten years*. And recognizing the correct interpretation of *carries* may not be straightforward, since its subject is not specifically identified as a crime and its object is not specifically identified as a punishment. The selectional needs for the correct sense of *carry* include a subject which names a crime and an object which can be construed as a penalty. Thus, the antecedent of *which* in this sentence, the *assaulting* phrase, has to be recognized as the name of a crime: the defendant is charged with committing a crime, and it is the crime which carries the sentence.

The parsers we have tested seem to have difficulty with *was ordered held without bond*, in which both the verb *order* and the secondary predicate that represents content of the order are in passive voice. The corresponding active matrix verbs do not strike one as rare (*The judge ordered the defendant shot at dawn*, *I ordered the package redelivered*). It is interesting that of the more than 1,300 passive instances of the verb *order* in the British National Corpus, only two were found that have the syntactic structure we see here, and one of them is: *He was ordered held for a bail hearing on Tuesday*. We suspect that this family of phrases should be treated as exhibiting a fixed pattern related to the criminal process frame.

4 Evidentiality

An important representation problem presents itself with the words *alleged* and *allegedly*. When the adjective *alleged* precedes a nominal constituent it is associated with the category signified by the nominal, i.e. the NP represents something that has been claimed to be an instance of the category, and this is not a problem. And when the adverb *allegedly* belongs to the highest predicate in an event predication, it means that the event in question is a report and protects the author from being accused of making the claim himself. But when the adverb is embedded in a fact-reporting phrase that identifies a sub-event, as in *who was shot in the knee after allegedly firing two shots outside the White House*, there arises a representational problem involving adverb scope. The shot in the defendant's knee did not occur after someone *alleged* that he fired shots outside the White House.

We wish to represent the kind of meaning shown here as **evidential**, by which the author of a report is adding some sort of evidential qualification to a part of the description. The frame in question, **Evidentiality**, is a metalinguistic frame that has as frame elements the ASCRIBER (of an evidential label) and the DESCRIPTION (of an event or a category). The word list for this frame includes *alleged*, *allegedly*, *reported*, *reportedly*, *known*, *certified certifiably*, *authentic*, *suspected*, *self-described*, *admitted*, and perhaps a few others.

5 Slot-Fillers

Most of the work in FrameNet to date has focussed on the verbs (and some nouns and adjectives) which we call **frame-bearing** or **frame-evoking**—those lexical heads which evoke a frame whose frame elements are typically expressed by the dependent NPs, PPs, VPs and Ss. But now the FN database also has the means of targeting frame-relevant dependent nouns for a separate kind of annotation. This would include the names of artifacts when they appear in sentences exhibiting information about the functions for which they have been manufactured. In our text we have examples of instruments in a *with*-phrase, as in *assaulting with a deadly weapon*, and as the object of the verb *use*, as in *use crutches to walk into the court*. The FN project intends to provide "reverse" information about these lexical items, characterizing the slots they fill in the frames in which they most naturally occur and identifying the manner of their syntactic marking.

The FN treatment of weapons and weapon use will include subframes involving the discharging of firearms, where distinctions for direct object roles include the WEAPON, the PROJECTILE and the TARGET (*shoot a gun*, *shoot a bullet*, *shoot a person*, cf. *shoot at a person*); other verbs that specifically deal with using a firearm are *discharge* and *fire* (*fire a gun*, **fire a person*, *fire a bullet*). (In the case of *fire a shot* we can treat the verb as a support verb for the event noun *shot*.)

6 Conclusion

Table 2 shows a summary of some of the frames discussed in this paper and the relations among them. In addition to inheritance (elaboration) relations among frames, we define corresponding relations between elements of frames which are so related; some of these are shown by means of equals signs in the "Elements" column. For example, **Court-date-setting** elaborates **Appointment**, which has a subframe **Action**. The JUDGE of **Court-date-setting** is the PROT1 of **Appointment**, the Defense and Prosecution (collectively) are the PROT2 of **Appointment**, the ACTION of **Appointment** is constrained to be one of the subframes of the **Criminal Process** frame, and the location of this action is a courtroom. These frame descriptions are still very preliminary, and have been greatly compressed to fit into a readable chart. The 16 frames shown here represent the minimum needed to understand the criminal proceedings described in this news story, including a few high-level, abstract frames such as **Action** and **Event**. Although we are just beginning work in this content domain, we estimate that less than 200 frames will suffice to represent most of the semantics of the vast majority of newspaper crime stories, covering hundreds of frame-evoking lexical units and their valences. (We are not attempting to represent expert knowledge of law, fraudulent accounting, stock market manipulation, etc.)

The FrameNet project is not itself dedicated to NLP efforts as such, but but we hope the information it makes available to the research community is of the sort which suitably fits any of several kinds of NLP activities. Evidence on the syntactic and collocational environments of polysemous words in different senses should be an aid to Word Sense Disambiguation. The association of lemmata with the frames they evoke, and hence with other lemmata belonging to the same frame, should aid in topic recognition and hence coherence establishment.

The provision of a large body of semantically annotated sentences (that is, annotated with respect to single key words within them) should amount to a training corpus for automatic semantic tagging. In the extent to which the FN database is capable of exhibiting all of the major valence possibilities for each sense of each word, it should provide material for enhanced statistical surveys of word sense frequencies for polysemous words, and preferred subcategorization frames for given word senses.

References

Douglas E. Appelt, Jerry R. Hobbs, John Bear, David Israel, and Mabry Tyson. 1993. Fastus: A finite-state processor for information extraction from real-world text. In *Proceedings. IJCAI-93*, Chambery, France.

Collin F. Baker, Charles J. Fillmore, and John B. Lowe. 1998. The Berkeley FrameNet Project. In *COLING-ACL '98: Proceedings of the Conference, held at the University of Montréal*, pages 86–90. Association for Computational Linguistics.

Charles J. Fillmore, Charles Wooters, and Collin F. Baker. 2001. Building a large lexical databank which provides deep semantics. In Benjamin Tsou and Olivia Kwong, editors, *Proceedings of the 15th Pacific Asia Conference on Language, Information and Computation*, Hong Kong.

Daniel Gildea and Daniel Jurafsky. 2000. Automatic labeling of semantic roles. In *ACL 2000: Proceedings of ACL 2000, Hong Kong*. Association for Computational Linguistics.

John B. Lowe, Collin F. Baker, and Charles J. Fillmore. 1997. A frame-semantic approach to semantic annotation. In Marc Light, editor, *Tagging Text with Lexical Semantics: Why, What and How?* Special Interest Group on the Lexicon, Association for Computational Linguistics.

Ellen Riloff and Mark Schmelzenbach. 1998. An empirical approach to conceptual case frame acquisition. In *Proceedings of the Sixth Workshop on Very Large Corpora*, pages 49–56, Montreal, Canada.

Roger C. Schank. 1972. Conceptual dependency: A theory of natural language processing. *Cognitive Psychology*, 3:552–631.

Robert F. Simmons. 1973. Semantic networks: Their computation and use for understanding English sentences. In Roger C. Schank and Kenneth Mark Colby, editors, *Computer Models of Thought and Language*, pages 61–113. W.H. Freeman and Co., San Francisco.

John C. Trueswell, Michael K. Tanenhaus, and Susan M. Garnsey. 1994. Semantic influences on parsing: Use of of thematic role information in syntactic ambiguity resolution. *Journal of Memory and Language*, 33:285–318.

Frame	Elaborates	Elements	Sub-Frames
Court		Judge(s) Officer(s) Courtroom	
Criminal Process		Court Prosecution (group) Prosecutor Defense (group) Defendant(s) Defense Attorney(s) Charges Law	Court Appearance Arrest Accusation Arraignment Preliminary hearing Trial Verdict
Arraignment		Defendant Court	Stating of Charges Entering of Plea
Confinement		Jailer Prisoner	
Pre-trial Confinement	Confinement	Bail	Confinement Posting of bail Release on bail Return to court Flight
Appointment		Prot1 Prot2	Promise Action
Court-date-setting	Appointment	Judge = Prot1 Defense & Prosecution = Prot2 Action = subframe of criminal process Action.Place = courtroom	
Event		Theme ("affected object") (Cause) (Result) Place Time	
Action	Event	Actor (Means) (Manner)	Action = Event
Action-Intentional	Action	Actor.type = Sentient	Forming of Intention Action = Action Result = Result
Action on Bodily target	Action	Prot1 = Actor Prot2 Prot2_body_part = Theme	
Crime-against-people	Action-Intentional	Perpetrator = Actor Victim = Theme Means/Weapon = Means	Malice = Intention Action = "harming victim"
Crime-law		Definition Penalty Jurisdiction	
Assault	Crime-against-people Action on bodily target	Assailant = Perpetrator Intention = "bodily harm" (Threat) = (of Action) (Action) = Crime-against-people.Action = Action-on-Bodily-Target	
Assault with deadly weapon	Assault	Weapon.type = "deadly"	
Shooting	Action	Shooter = Actor Gun = Means Projectile = Theme (Point of impact)	

Table 2: Summary of Frames and Relations