A Frame-Semantic Approach to Semantic Annotation

John B. Lowe jblowe@garnet.berkeley.edu

Collin F. Baker collinb@icsi.berkeley.edu

Charles J. Fillmore fillmore@cogsci.berkeley.edu
Department of Linguistics
University of California
Berkeley, CA 94720

Abstract

The number and arrangement of semantic tags must be constrained, lest the size and complexity of the tagging sets (tagsets) used for semantic annotation become unwieldy both for humans and computers. The description of lexical predicates within the framework of frame semantics provides a natural method for selecting and structuring appropriate tagsets.

1 Motivation

The research present here is to be conducted under the *FrameNet* research product at the University of California.[1] On this project our primary aim is to produce frame-semantic descriptions of lexical items; our concern with semantically tagged corpora is at both ends of our research. That is, we expect to use partially semantically tagged corpora in the investigation stage—perhaps nothing more than having WordNet hypernyms associated with nouns—but we will produce semantically tagged corpus lines as a by-product of our work.

Most major grammatical theories now accept the general principle that some set of semantic roles ("case roles", "thematic roles", or "theta roles") is necessary for characterizing the semantic relations that a predicate can have to its arguments. This would seem to be one obvious starting-point for choosing a tag set for semantically annotating corpora, but there is no agreement as to the size of the minimal necessary set of "universal" roles. Also, when we examine particular semantic fields, it is obvious that each field brings to mind a new set of more specific roles. In fact, the more closely we look at individual predicates, the more specific the argument roles become, creating the specter of trying to define an unlimited number of very fine-grained tags and attributes. An adequate account of the syntax and semantics of a language will inevitably involve a fairly detailed set of semantic tags, but how can we find the right level of *granularity* of tags for each semantic area?

Consider the sentence:

(1) The waters of the spa cure arthritis.

A semantic annotation of the constituents must identify at least

- the action or state associated with the verb, possibly expressed in terms of primitives or some kind of metalanguage;
- the participants (normally expressed as *arguments*); and
- the *roles* of the participants in the action or state.

A basic parse will identify the sentence's syntactic constituents; from the point of view of the head verb *cure*, then, a semantic annotation should reveal the mapping between the syntactic constituents and the frame-semantic elements they instantiate. In sentence (1) above, for example, the grammatical *subject* "the waters of the spa" corresponds to the thematic *causer* of the curing effect on the entity expressed as "arthritis", the verb's syntactic *direct object* and its thematic *patient*.[2]

However, there is something incomplete about such an analysis: it fails to anchor the arguments of

[1]The work is housed in the International Computer Science Institute in Berkeley and funded by the National Science Foundation under NSF grant IRI 96-18838. The official name of the project is "Tools for lexicon building"; the PI is Charles J. Fillmore. Starting date March 1, 1997.

[2]Here we use the word *patient* (in italics) as the name of a case role; we will also use the word in the medical sense later in this paper. *Caveat lector!*

cure within a "generic medical event" where it would be understood that the disease (arthritis) must be borne by some *sufferer*, and that a sufferer undergoing a *treatment* is participating as a *patient* in such an event. We identify such "generic events" as *frames*, and express our understanding of the structure of such events and the relationship of linguistic material to them in terms of the theory of *frame semantics*.

2 Frame Semantics.

In frame semantics we take the view that word meanings are best understood in reference to the conceptual structures which support and motivate them. We believe, therefore, that any description of word meanings must begin by identifying such underlying conceptual structures.[3]

Frames have many properties of stereotyped scenarios — situations in which speakers expect certain events to occur and states to obtain.[4]

In general, frames encode a certain amount of "real-world knowledge" in schematized form. Consider the common scenario which exemplifies the *commercial transaction frame*: the elements of such frames are the individuals and the props that participate in such transactions (which we call FRAME ELEMENTS): the individuals in this case are the two protagonists in the transaction; the props are the two objects that undergo changes of ownership, one of them being money.

Some frames encode patterns of opposition that human beings are aware of through everyday experience, such as our awareness of the direction of gravitational forces; still others reflect knowledge of the structures and functions of objects, such as knowledge of the parts and functions of the human body. The study of the frames which enter into human cognition is itself a huge field of research – we do not claim to know in advance how much frame knowledge must be specifically encoded in frame descriptions to make them useful for either linguistic or NLP purposes. We expect to be able to draw tentative conclusions about this based on what we find in corpora.

[3]For a discussion of these ideas, see (Fillmore, 1968); (Fillmore, 1977b); (Fillmore, 1977a); (Fillmore, 1982); (Fillmore and Atkins, 1992); (Fillmore and Atkins, 1994).

[4]The word *frame* has been much used in AI and NLP research. We wish to give the word a formal interpretation only to the extent that it helps us in our research and provides a container for the features and entities we describe. We do not, in this context, depend on any claims about the cognitive status of frames.

We will say that individual words or phrases *evoke* particular frames or *instantiate* particular elements of such frames. So, for example, if we are examining the "commercial transaction" frame, we will need to identify such frame elements as BUYER, SELLER, PAYMENT, GOODS, etc., and we can speak of such words as *buy, sell, pay, charge, customer, merchant, clerk*, etc., as capable of evoking this frame. In particular sentences, we might find such words or phrases as *John, the customer*, etc. instantiating the BUYER, or *a chicken, a new car*, etc., instantiating the GOODS.

3 Inheritance in Frame Semantics

Of course, speakers of a language know something about the differences and similarities among various types of commercial transactions, e.g. that buying a small item in a store often involves making change, etc. Strictly speaking, this is "world knowledge" rather than "linguistic knowledge", but this level of detail is required even to parse sentences correctly, e.g. to recognize the different functions of the PPs in "buy a candy bar with a red wrapper" and "buy a candy bar with a $20 bill" and thus to attach them appropriately.

frame (CommercialTransaction)
frame-elements{BUYER, SELLER, PAYMENT, GOODS}
scenes (BUYER gets GOODS,
SELLER gets PAYMENT)

frame (RealEstateTransaction)
inherits (CommercialTransaction)
link(BORROWER = BUYER, LOAN = PAYMENT)
frame-elements{BORROWER, LOAN, LENDER}
scenes (LOAN (from LENDER) creates PAYMENT,
BUYER gets LOAN)

Figure 1: A subframe can inherit elements and semantics from its parent.

More complicated cases require more elaborated frames. Thus, "buy a house with a 30-year mortgage" involves a different frame from buying a candy bar, and entails a slightly different interpretation of the PAYMENT element. The relationship between frames is frequently hierarchical; for example, the frame elements BUYER, SELLER, PAYMENT, and GOODS will be common to all commercial transactions; the purchase of real estate contains all of them and (typically) adds a LOAN and a bank (typically) as LENDER. In our database, these two frames might

be represented as shown in Figure 1.[5]

Corpus tagging for a sentence like sentence (2):

(2) Susan took out a huge mortgage to buy that new house.

would have to recognize Susan as playing slightly different roles in the two associated frames.

A similar problem in using labels from frame semantic descriptions in the tagging of corpus lines is due to the fact that separate parts of any single sentence can evoke different semantic frames. Consider the following sentence:

(3) George's cousin bought a new Mercedes with her portion of the inheritance.

In seeing this sentence merely as an expression evoking the commercial transaction frame, we could begin by tagging the subject of the sentence, "George's cousin", as the BUYER, and the object, "a new Mercedes" as the GOODS, and the oblique object, "her portion of the inheritance", marked by the preposition "with", as the PAYMENT. This could be done in a fairly natural and transparent way, as long as the tags were clearly seen as the names of frame elements specifically related to the head verb "bought" in that sentence. But since the words "cousin" and "inheritance" evoke frames of their own, the same sentence could easily come up in our exploration of the semantics of those words as well. In the case of "inheritance", for example, the information that it gets used for buying something will make clear that this is an instance of estate-inheritance rather than genetic inheritance (or frame inheritance!), and the phrasing "her portion" fits frame understandings about the distribution of an inheritance among multiple heirs. In other words, if we find ourselves tagging the frame elements of Inheritance in that same sentence, the phrase "George's cousin" would be tagged as an HEIR in that frame.

4 Applied frame semantics: a sample frame description.

Tagsets for semantic annotation would be derivable from a database of frame descriptions like the ones in Figure 1 above. We can move to another frame to illustrate how frame-based annotation would be accomplished by considering a few words from the language of health and sickness and showing how the elements and structure of this frame would be identified and described. First, appealing to common, unformalized knowledge of health and the body, the frame semanticist identifies the typical elements in everyday health care situations and scenarios, a process involving the interaction of linguistic intuition and the careful examination of corpus evidence.

label	meaning
HEALER	individual who tries to bring about an improvement in the PATIENT
PATIENT	individual whose physical well-being is low
DISEASE	sickness or health condition that needs to be removed or relieved
WOUND	tissue damage in the body of the PATIENT
BODYPART	limb, organ, etc. affected by the DISEASE or WOUND
SYMPTOM	evidence indicating the presence of the DISEASE
TREATMENT	process aimed at bringing about recovery
MEDICINE	substance applied or ingested in order to bring about recovery

Table 1: Part of Frame-semantic "Tagset" for the Health Frame

The first product of this analysis is a preliminary list of frame elements (FEs) from this domain, such as, for instance, those shown in Table 1.

We have found it necessary to include all of these elements for our purposes, even though some of them are so closely related that they are unlikely to be given separate instantiation in the same clause. Our justification for distinguishing them is based on the results of corpus research and on comparison of the elements of this frame with those of other related frames. Corpus examples in which WOUND and DISEASE are both instantiated are of course rare, and given this complementary distribution we might be tempted to identify these as variants of a single frame element (which we might call AFFLICTION). But this would prevent us from being able to express certain syntactic and semantic generalizations, such as the fact that while we speak of curing diseases, we do not speak of curing wounds, and we speak of wounds but not diseases as healing.[6]

[5]We leave out of this account the inheritance of a higher-level EXCHANGE frame in the COMMERCIAL-TRANSACTION frame, and the means for showing that a completed instance of the REALESTATETRANSACTION scene is a prerequisite to the enactment of the associated COMMERCIALTRANSACTION scene.

[6]There might be alternative ways of considering such data. It is conceivable that a description with, say, AFFLICTION as a single role element could be maintained

In the specific case of the contrast between WOUND and DISEASE we find in metaphor further support for our decision to keep them separate. Metaphoric uses of "cure" and "heal" tend to take direct objects which are target-domain analogues of DISEASE and WOUND respectively. One of the most common instantiations of the DISEASE complement in metaphorical uses of *cure* is the word *ills*, a word which in fact appears to be used only in such metaphorical contexts (in talk about "curing society's ills", for example); and the direct objects of metaphorical *heal* tend to be based on the notion of a tear or cut or separation, the words *wound* and *scar* first of all, but also such words as *rift*, *schism*, and *breach*.

For each semantic frame, the process of elucidation involves a series of steps:

1. Identification of the most frequent lexical items which can serve as predicates in this frame,
2. Formulation of a preliminary list of frame elements (encoded we expect as a TEI-compliant SGML document using *feature structures* (Sperberg-McQueen and Burnard, 1994),
3. Annotation of examples from a corpus by tagging the predicate with the name of the frame and its arguments with the names of the FE's designating their roles relative to the predicate (also using SGML markup introduced with software developed for this purpose),
4. Revision of the frame description — specification of the co-occurrence constraints and possible syntactic realizations in the light of the corpus data, and,
5. Retagging of the corpus examples to fit the revised frames.[7]

The last two steps will be repeated as needed to refine the frame description.

by describing certain distinctions between "cure" and "heal" as involving selectional restrictions. Our inclination, however, is to maximize the separation of frame elements at the beginning, and to postpone the task of producing a parsimonious and redundancy-free description until after we have completed our analysis.

[7]In the context of the FrameNet project, the question of how much text will be tagged is a practical one. Our direct purpose is not to create tagged corpora, but to tag enough corpus lines to allow us to make reliable generalizations on the meanings and on the semantic and syntactic valence of the lexical entries we have set out to describe. Whether we choose to tag more than what we need for our analysis will depend on the extent to which the process becomes automated and the resources available.

Identifying the semantic frame associated with a word and the FEs with which it constellates does not, of course, constitute a complete representation of the word's meaning, and our semantic descriptions will not be limited to just this. However, we believe that such an analysis is a prerequisite to a theoretically sound semantic formalization.[8] While any given frame description could be made more precise for other NLP/AI purposes (such as inference-generation), the development of such a formalism is not a central part of our current work.

For our present purposes, the adequacy of lists of frame elements such as what we present in Table 1 for the vocabulary domain of health care can be established only if precisely these elements are the ones that are needed for distinguishing the semantic and combinatorial properties of the major lexical items that belong to that domain. An initial formulation of the combinatorial requirements and privileges of a frame's lexical members — here we concentrate on verbs — can be presented as a list of the groups of FEs that may be syntactically expressed or perhaps merely implied in the phrases that accompany the word.

A Frame Element Group (FEG) is a list of the FEs from a given frame which occur in a phrase or sentence headed by a given word. Table 2 gives examples of such FEGs (including FEGs with only one member) paired with sentences whose constituents instantiate them. For purposes of this discussion, the frame elements are identified here using single letter abbreviations, and the structure of an FEG is shown as being merely a bracketed list. We recognize such a naming scheme is inadequate for a large annotation project, and certainly the representation of FEG structures will have to be more powerful. These, however, are minor problems with technical solutions. We focus below on other major issues we are confronting in interpreting the structure of frames as expressed by FEGs.

At the lexicographic level of description we could simply list the full set of FEGs for a given lexical unit. However, in many cases the FEG potential of a verb can be expressed in one or more simplifying formulas, by, for example, recognizing some FEs as optional. Thus, since we find both {H, B} ("The doctor cured my foot") and {H, B, T} ("The doctor cured my foot with a new treatment"), both sentences are using the verb *cure* in the same sense, we can represent both patterns in a single formula that treats the T element as an optional adjunct

[8]There are numerous suggestions, not reviewed here, on how to give full semantic representations (Jackendoff, 1994); (Sowa, 1984); (Schank, 1975), etc.

FEG (abbr.)	Frame Element Group	Example
{H,B,T}	HEALER, BODYPART, TREATMENT	The doctor treated my knee with heat.
{H,D}	HEALER, DISORDER	The doctor cured my disease.
{P}	PATIENT	The baby recovered.
{M,B}	MEDICINE, BODYPART	The ointment cured my foot.
{B}	BODYPART	His foot healed.
{W}	WOUND	The cut rapidly healed.

Table 2: Examples of Frame Element Groups (FEGs)

(expressed perhaps as {H, B, (T)}).

It will not be quite that automatic, however; further distinctions are needed. For example, while we can agree that the TREATMENT element in the previous examples was merely unmentioned, the omission of the DISEASE element in a sentence like "The doctor cured me" has a somewhat different status: there is clearly some DISEASE that the speaker has in mind, and its omission is licensed by the assumption that its nature is given in the context. That is, a possible "of" phrase was omitted from that sentence because its content had been previously mentioned or could otherwise be assumed to be known to both conversation participants. In the tagging of corpus lines, then, we will also indicate the status of "missing" elements to the extent that we can tell what that is. Such information will be presented in the representation of the FEG associated with the predicate. [9]

In contrast to cases where frame elements are "missing" (implied but unmentioned, optional, etc.), some examples require that we explicitly recognize (i.e. encode) multiple frame elements for a single constituent. Thus, the disorder may be identified in the description of the patient (e.g. *leper*, *diabetic*); we wish to annotate this constituent as P_d, which will be taken as indicating that the constituent satisfies the P role in the frame, but that it also secondarily instantiates a D role, since these nouns designate people who suffer specific diseases (leprosy, diabetes). It is important to recognize these cases, since the lexical semantics of verbs sometimes require that certain frame elements be instantiated or clearly recoverable from the context: corpus research on the verb *cure*, for example, shows that the DISORDER is regularly instantiated. Without explicit coding of the substructure of the PATIENT the sentence *He cured the leper* ({H,P_d}) would stand as a counter-example to this generalization.

There are cases where different but related senses of a predicate have distinct FEG possibilities. For example, the verb *heal* has two uses, one of which participates in a Causative/Inchoative valency alternation (Levin, 1993) and one which does not. In the use where it refers to the growth of new tissue over a wound, it can be found in both transitive and intransitive clauses: "The cut healed" ({W}) and "The ointment healed the cut" (the ointment facilitated the natural process of healing — {M, W}). But there is also a purely transitive use with a meaning very close to that of *cure*, with {H, D} or {M, D}, as in "The shaman healed my influenza" or "The waters healed my arthritis", and this use of *heal* usually implies something extra-medical or supernatural. In this usage, there is no corresponding intransitive "*My influenza/arthritis healed."

The verb sense distinctions we make may sometimes be less detailed than those appearing in most dictionaries, since, as many researchers have noted, dictionary sense distinctions are often overprecise and incorporate pragmatic and world knowledge that do not properly speaking inhere in the word itself. An excellent example of this kind of excessive distinction is pointed out in (Ruhl, 1989), p.7: one of the dictionary definitions of *break* is "to rupture the surface of and permit flowing out or effusing" as in *He broke an artery.* On the other hand, we would expect to capture by this process all the kinds of alternations that (Levin, 1993) has shown to be linked to semantic distinctions, some of them quite subtle.

The final versions of the lexical entries will encompass full semantic/syntactic *valence descriptions*, where the elements of each FEG associated with a verb sense will be linked to a specification of *sortal features*, indicating the "selectional" and syntactic properties of the constituents that can instantiate them.

5 Conclusion

We have suggested a theoretical basis and a working methodology for coming up with an appropriate set of semantic tags for the semantic frame elements, and believe that such frames may constitute a sort of "basic level" of lexical semantic description. As

[9] Where feasible, because of our interest in sortal features of arguments, we will identify the nature of the missing element from the context. A similar issue arises in cases of anaphora; we may or may not resolve the anaphora's referent in the annotations, depending on practical considerations of time and effort involved.

such they would be an appropriate starting-point for both a broad-coverage semantic lexicon and for the semantic tagging of corpora.

We have also pointed out the importance of incorporating the notions of inheritance and other substructuring conventions in tagsets to reduce the size and complexity of the descriptions and to capture generalizations over natural classes.

We recognize several shortcomings with our approach which we hope to be able to address in the future.

First, it is clear that the size of the descriptions will increase rapidly as the annotation proceeds and we will need to find some explicit means of abbreviating representations, of collapsing FEGs in a principled way, and of relating frames together (both within and across semantic fields). This is both a practical and theoretical problem. We have shown a few clear examples in which the judicious use of the notion of inheritance, along the general lines of the ACQUILEX Project (Briscoe et al., 1993), should permit the concise representation of the lexical knowledge required to give a useful and relatively complete description of a word's semantic range. If the valence description (the FEG together with links to grammatical functions) associated with individual words is attached to each valence-bearing lexical token in a corpus, then if the corpus is parsed according to the same criteria by which the linking has been stated, we can avoid the problem of actually tagging the phrases that instantiate frame elements (and hence avoid the problem of multiple tagging for constituents that figure in more than one frame in the same sentence), because the constituents that play specific semantic roles in the sentence can be computed from the parse. The ability to accomplish something like that is desirable, but it is not something to which we are presently committed.

We intend first to focus on *prototypical* or *core* uses of the words. However, our preliminary research indicates that it would be difficult, and undesirable, to exclude metaphorical uses, if only because the metaphorical uses can often shed light on the structure of the core uses. However, we are limiting our attention to a limited number of semantic domains, and metaphorical extensions from the words in our wordlist that go far beyond our semantic fields will probably have to be set aside.

Finally, we should make a few remarks on the scope of our intended effort. We plan to create a "starter lexicon" containing some 5,000 lexical items indexed to examples of their use. With each entry we shall associate token frequencies with the various FEGs for each word sense, in order to assist NLP programs in picking likely interpretations. Initially the frequencies would be generated using our hand-tagged corpus examples; eventually we hope to be able to train on the hand-tagged examples and ultimately automate (at least partially) the tagging of instances, at least for preliminary word sense disambiguation, to be reviewed by a researcher. The automatic categorization of the arguments would use such information as WordNet synonyms and hypernyms (cf.(Resnik, 1993)), machine-readable thesauri, etc.,

References

Ted Briscoe, Valeria De Paiva, and Ann Copestake, editors. 1993. *Inheritance, Defaults and the Lexicon.* Studies in Natural Language Processing. Cambridge University Press, Cambridge, England.

Charles J. Fillmore and B.T.S. Atkins. 1992. Towards a frame-based lexicon: the semantics of risk and its neighbors. In A. Lehrer and E. F. Kittay, editors, *Frames, Fields and Contrasts*, pages 75–102. Lawrence Erlbaum Associates, Hillsdale, NJ.

Charles J. Fillmore and B.T.S. Atkins. 1994. Starting where the dictionaries stop: the challenge for computational lexicography. In B.T.S. Atkins and A. Zampolli, editors, *Computational Approaches to the Lexicon.* Oxford University Press, New York.

Charles J. Fillmore. 1968. The case for case. In *Universals in linguistic theory*, pages 1–90. Holt, Rinehart and Winston, New York.

Charles J. Fillmore. 1977a. The need for a frame semantics within linguistics. *Statistical Methods in Linguistics*, pages 5–29.

Charles J. Fillmore. 1977b. Scenes-and-frames semantics. In Antonio Zampolli, editor, *Linguistics Structures Processing*, volume 59 of *Fundamental Studies in Computer Science*, pages 55–82. North-Holland Publishing.

Charles J. Fillmore. 1982. Frame semantics. In *Linguistics in the morning calm*, pages 111–137. Hanshin Publishing Co., Seoul, South Korea.

Ray S. Jackendoff. 1994. *Patterns in the mind: language and human nature.* Basic Books, New York.

Beth Levin. 1993. *English Verb Classes and Alternations: A Preliminary Investigation.* University of Chicago Press, Chicago.

Philip Resnik. 1993. *Selection and Information: A Class-Based Approach to Lexical relationships.* University of Pennsylvania dissertation.

Charles Ruhl. 1989. *On monosemy : a study in linguisitic semantics*. Albany, N.Y.: State University of New York Press.

Roger C. Schank. 1975. *Conceptual information processing.* North-Holland., New York.

John F. Sowa. 1984. *Conceptual structures: information processing in mind and machine.* Addison-Wesley systems programming series. Addison-Wesley, Reading, Mass.

Michael Sperberg-McQueen and Lou Burnard. (eds.) 1994. *Guidelines for electronic text encoding and interchange (TEI P3).* ACH, ACL, ALLC, Chicago.