

Frame Semantics

Jean Mark Gawron

January 31, 2008

For there exists a great chasm between those, on the one side, who relate everything to a single central vision, one system more or less coherent or articulate, in terms of which they understand, think and feel — a single, universal, organizing principle in terms of which alone all that they are and say has significance — and, on the other side, those who pursue many ends, often unrelated and even contradictory, connected, if at all, only in some de facto way, for some psychological or physiological cause, related by no moral or aesthetic principle.

Berlin (1997:436),
cited by Minsky (1975)

1 Introduction

Two properties of word meanings contribute mightily to the difficulty of providing a systematic account.

One is the **openness** of word meanings. The variety of word meanings is the variety of human experience. Consider defining words such as *Tuesday*, *barber*, *alimony*, *seminal*, *amputate*, and *brittle*. One needs to make reference to diverse practices, processes, and objects in the social and physical world: repeatable calendar events, grooming and hair, marriage and divorce, discourse about concepts and theories, and events of breaking. Before this seemingly endless diversity, semanticists have in the past stopped short, excluding it from the semantic enterprise, and attempting to draw a line between a small linguistically significant set of primitive concepts and the openness of the lexicon.

The other problem is the closely related problem of the **richness** of word meanings. Words are hard to define, not so much because they invoke

fine content specific distinctions, but because they invoke vast amounts of background information. The concept of buying presupposes the complex social fact of a commercial transaction. The concept of alimony presupposes the complex social fact of divorce, which in turn presupposes the complex social fact of marriage. Richness, too, has inspired semanticists simply to stop, to draw a line, saying exact definitions of concepts do not matter for theoretical purposes.

This boundary-drawing strategy, providing a response if not an answer to the problems of richness and openness, deserves some comment. As linguistic semanticists, the story goes, our job is to account for systematic, structurally significant properties of meaning. This includes:

- (1) a. the kinds of syntactic constructions lexical meanings are compatible with.
 - 1. the kinds of participants that become subjects and objects
 - 2. regular semantic patterns of oblique markings and valence alternations
- b. Regular patterns of inference licensed by category, syntactic construction or closed class lexical item.

The idea is to carve off that part of semantics necessary for knowing and using the syntactic patterns of the language. To do this sort of work, we do not need to pay attention to every conceptually possible distinction. Instead we need a small set of semantic primitives that make the distinctions that linguistically matter; what is left over can be dealt with using some open class of predicates or features whose internal details are not of concern. Jackendoff (1990) is a good example of this kind of approach. The generative semantics program, especially as outlined in Lakoff (1972), is another. Dowty (1979) has many of the same features, but in places expresses doubts that the program can be completely carried out. The kind of analysis I have in mind can be exemplified through Dowty's generative-semantics-like analysis of causatives like *break.tr* (transitive *break*):

- (2) a. John broke the glass.
- b. DO(John, CAUSE(BECOME(broken'(glass'))))

Here the predicates in capitals (DO, CAUSE, BECOME) are from the inventory of linguistically significant primitives, and the lower case predicates (*broken'*, *glass'*) are from the open class predicates whose internal structure does not matter. At most we need to know that one expresses a state (*broken'*) and the other a kind (*glass'*). The details beyond that are linguistically

insignificant. Of course there are differences in truth-conditions between states like *broken*' and *dead*', but these have only minor selectional effects on the causative inchoatives created from them (*break.tr* = DO ... CAUSE BECOME broken' and *kill* = DO ... CAUSE BECOME dead'). I will refer to this view of lexical semantics as the **classical view**.

In this paper I wish to consider a view of semantics in general and lexical semantics in particular that is quite at odds with this classical picture: *frame semantics* (Fillmore 1975, Fillmore 1978, Fillmore 1977b, Fillmore 1982, Fillmore 1985). Someone wishing to contest the classical picture has two options: first, contend that the wrong kinds of questions are being asked; second, argue that the program as outlined is not very well-suited to attaining its goals. As we shall see, both kinds of reply provide motivations for frame semantics.

1.1 Motivations

The version of frame semantics I will present here is largely the brainchild of Charles J. Fillmore. Although frame semantics has sprouted off in a number of directions and been applied to a number of problems, I will limit the present discussion in two ways: First I will confine myself largely to fleshing out the Fillmorean picture; second, I will confine myself mostly to questions of the lexicon, lexicography, and the lexicon-syntax interface, leaving for other work questions of discourse and text understanding to which frames are also relevant.

Although Fillmore has had many interesting things to say about the kinds of problems listed in (1) in early and late works on Case Grammar, the primary motivations given in Fillmore (1982, 1985) focus on frame semantics as a contribution to a theory of *text understanding*. Consider for example, the very different scenes evoked by the following pair of sentences:¹

- (3) a. I can't wait to be on the ground again.
b. I can't wait to be on land again.

Sentence (3a) evokes a speaker who is in the air (on a plane), sentence (3b) a speaker who is at sea (on a ship). This contrast is tied to some difference between the words *land* and *ground*, yet, on the face of it, *land* and *ground* denote very similar things. Fillmore would say *land* is understood within a conceptual frame of sea travel, and within that frame it is opposed to *sea*, while *ground* is understood within a conceptual frame of air travel, and

¹This example is a modified version of an example Fillmore (1985) discusses.

within that frame, it is opposed to *air*. Thus we can explain something that is very difficult to explain in terms what the words in the sentence denote by investigating the conceptual background against which the relevant word senses are defined. That conceptual background is what Fillmore calls a frame.

Frames are conceptual structures that provide context for elements of interpretation; their primary role in an account of text understanding is to explain how our text interpretations can (validly) leap far beyond what the text literally says. Frames can be introduced into interpretation in a variety of ways. They may be directly tied to word senses as in the example of *land* and *ground*, or they may be introduced by patterns among the facts the text establishes. To use another example of Fillmore's (1985), 232:

(4) We never open our presents until morning.

This sentence *evokes* the Christmas frame by describing a situation that matches salient facts of Christmas practice, even though no word in it is specific to Christmas. If in fact the Christmas frame is the right one, that evocation makes a significant contribution to the understanding of the surrounding text.

Frames are motivated not just by words, then, but by stereotypes about customs, practices, institutions, and games. Moreover, the kinds of cognitive structures Fillmore has in mind have been proposed by a variety of researchers for a variety of purposes. Fillmore has adopted the terminology of AI researcher Minsky (1975) in calling them frames, but *schemata* in psychology (Bartlett 1932, Rumelhart 1975) are getting at something very similar, as are *scripts* (Schank and Abelson 1977), *cognitive models* (Lakoff 1983), *experiential gestalts* (Lakoff and Johnson 1980), *the base* (as opposed to the profile) (Langacker 1984), and Fillmore's own notion of *scene* (Fillmore 1976, Fillmore 1977a).

As an approach to word meanings specifically, the starting point for frame semantics is that the lexical semantics "problems" of openness and richness are connected. Openness depends on richness. Openness does not mean lack of structure. In fact, it presupposes structure. Most concepts are interpretable or understandable or definable only against the background of other concepts. Many backgrounds are rich enough to define a cluster of concepts, in particular, a cluster of words. These backgrounds are the frames. Thus because words are networked together through their shared backgrounds, frames can provide an organizing principle for the openness of the lexicon.

Consider one of the examples already discussed, discussed in Fillmore (1982). The concept of alimony depends on the concept of divorce. The concept of divorce in turn depends on the concept of marriage. The dependency is definitional. Unless you define what a marriage is, you can't define what a divorce is. Unless you define what a divorce is, you can't define what alimony is. Thus there is a very real sense in which the dependencies we are describing move us toward *simpler* concepts. Notice, however, that the dependency is leading in a different direction than an analysis that decomposes meanings into a small set of primitives like CAUSE and BECOME. Instead of leading to concepts of increasing generality and abstractness, we are being led to define the situations or circumstances which provide the necessary *background* for the concepts we are describing. Marriages and divorces are equally concrete, but the institution of marriage provides the necessary background for the institution of divorce.

Or consider the complex subject of Tuesdays (Fillmore 1985). We live in a world of cyclic events. Seasons come and go and then return. This leads to a cyclic calendar which divides time up into repeating intervals, which are divided up further. Years are divided into months, which are divided into weeks, which are divided into days, which have cyclic names. Each week has a Sunday, a Monday, a Tuesday, and so on. Defining Tuesday entails defining the notion of a cyclic calendar. Knowing the word *Tuesday* may not entail knowing the word *Sunday*, but it does entail understanding at least the concept of a week and a day and their relation, and that each week has exactly one Tuesday.

Or consider the words *momentum* and *speed*. Obviously both in some sense conceptually presuppose motion. Here we might say that motion is an abstract concept underlying the concepts of speed and momentum, but notice the relationship between the particular and abstract concepts is not one we might call an *instance of* relation or a *subtype of* relation. Speed is not a subtype of motion, nor is momentum. Rather the concepts of concepts of speed and momentum can only arise once we assume motion as a background. And given that, the particulars of how they arise are rather complicated.

Finally consider adjectives of temperature: *freezing*, *cold*, *cool*, *tepid*, *lukewarm*, *warm*, and *hot*. Scalar adjectives in general require some attribute, some class of objects that can possess the attribute, and some scale measuring the degree to which such objects possess the attribute. In this case we will say the temperature adjectives require the background concept of TEMPERATURE. But scales don't exist in Platonic isolation. Also requiring this concept, we have nouns like *temperature* and *thermometer*.

We thus have words and background concepts. We will call the background concept the *frame*. Now the idea of a frame begins to have some lexical semantic bite with the observation that a single concept may provide the background for a set of words. Thus the concept of MARRIAGE provides the background for words/suffixes/phrases such as *bride, groom, marriage, wedding, divorce, -in-law, elope, fiancée, best man, maid-of-honor, honeymoon, husband, and wife*.² The concept of CALENDAR CYCLE provides the frame for lexical items such as *week, month, year, season, Sunday, ..., Saturday, January, ..., December, day, night, morning, and afternoon*. Notice that a concept once defined may provide the background frame for further concepts. Thus, DIVORCE itself provides the background frame for lexical items such as *alimony, divorce, divorce court, divorce attorney, ex-husband, and ex-wife*.

Thus, Fillmore's move toward a theory of understanding has dramatic consequences for lexical semantics. Each frame organizes a vocabulary domain:

“Borrowing from the language of gestalt psychology we could say that the assumed background of knowledge and practices — the complex frame behind this vocabulary domain — stands as a common ground to the figure representable by any of the individual words.”

“[Words belonging to a frame] are lexical representatives of some single coherent schematization of experience or knowledge”

Fillmore (1985:223)

Now a premise of frame semantics is that the relation between lexical items and frames is open ended. Thus one way in which the openness of the lexicon manifests itself is in building concepts in unpredictable ways against the backdrop of other concepts. The concept of marriage seems to be universal or near-universal in human culture. The concept of alimony is not. No doubt concepts sometimes pop into the lexicon along with their defining frames (perhaps *satellite* is an example), but the usual case is to try to build them up out of some existing frame (Thus *horseless carriage* leading to *car* is the more usual model).

Summing up: openness does not mean structurelessness. Concepts and their related words have certain unidirectional backgrounding relations.

²More generally, a variety of kinship relations have the concept of marriage built in (*aunt* and *uncle*). I assume there is a cascade of further frames built up from the marriage frame. Here I confine myself to some directly related vocabulary.

(5)	Words	Frames
	<i>bride, groom, marriage, wedding, divorce, -in-law, elope, fiancée, best man, maid-of-honor, honeymoon, husband, wife</i>	MARRIAGE
	<i>alimony, divorce court, divorce attorney, ex-husband, and ex-wife</i>	DIVORCE
	<i>week, month, year, season, Sunday, ..., Saturday, January, ..., December, day, night, morning, afternoon</i>	CALENDAR CYCLE
	<i>momentum, speed</i>	MOTION
	<i>freezing, cold, cool, tepid, lukewarm, warm, hot, temperature, thermometer</i>	TEMPERATURE

Frames and their associated lexical sets are a way of describing these relations.

All of this obviously points in exactly the opposite direction from the classical view, a few salient primitives, a hard distinction between linguistic and encyclopedic, and a large uninvestigated class of open class predicates.

But from the other direction, support for the classical view has been eroding even among those whose concerns have primarily departed from the problems in (1) such as Levin (1993) or from classic lexical semantic problems like polysemy (Pusteyovsky 1995).

Consider the kind of problem Beth Levin discusses in her broad study of the lexical semantic classes of English verbs (Levin 1993). A theory that does not posit some kind of systematically significant difference between *broken*' and *dead*' can not account for the following contrast:

- (6) a. John broke the glass against the wall.
 b. # John killed the cockroach against the wall.

Nor can it account for the fact that verbs in some sense close in meaning to *break* (*shatter, smash, crack, flatten*) will follow pattern (a), while verbs in some sense close to *kill* will follow pattern (b) (*strangle, murder, smother, and drown*). The generalization at issue is (roughly) that state change or directed action verbs whose effect is commonly achieved by moving one object against another will allow pattern (a) when the object whose state is changed or potentially changed is direct object. Other examples are *hit, knock, rap, bang, and slam*. None of the kill-type verbs fit the bill.

Thus if valence patterns are part of what is to be explained, then a language like English, with its rich inventory of prepositions and situationally

specific constructions (see for example the pattern lists in Levin 1993), will require reference to a large inventory of concepts. It is difficult to see how a principled line between open class and closed class concepts can be drawn in carrying out this program. It is clear for example, that Levin's verbs of contact, which include the verbs like *hit* and *slap* discussed above, overlap significantly with the verbs list for the IMPACT frame in FrameNet, a large computational instantiation of the ideas of frame semantics (Fillmore and Atkins 1994, Baker et al. 1998, Fillmore and Atkins 1998, Baker and Fillmore 2001, Boas 2001, Boas 2002, Chang et al. 2002a, Chang et al. 2002b).³ Thus the problems of openness and richness arise whether one starts from text understanding or from syntax/semantics interface.

1.2 Basic Tools

As a theory of word meanings, then, frame semantics makes the assumption that there is always some background knowledge relative to which a word does some profiling/highlighting, and relative to which it is defined. Two ideas are central:

1. a background concept
2. a lexical set including all the words that utilize this conceptual background.

Two other important frame theoretic concepts are *frame elements* and *profiling*.

Thus far in introducing frames I have emphasized what might be called the modularity of knowledge. Our knowledge of the the world can usefully be divided up into concrete chunks. Equally important to the Fillmorian conception of frames is the integrating function of frames. That is, frames provide us with the means to integrate with other frames in context to produce coherent wholes. To illustrate this, it will be useful at once to introduce the notion of a *frame element* as implemented in FrameNet. A frame element is simply a regular participant, feature, or attribute of the kind of situation described by a frame. Thus, frame elements of the marriage frame will include the husband, wife, wedding ceremony, wedding date, best man and maid of honor, for example. Frame elements need not be obligatory; one may have a marriage without a best man; but they need to be regular recurring features.

³At last count the NSF Framenet project (Fillmore and Baker 2000) which is building a frame lexicon for English had over 800 frames for about 4500 words.

Thus, frames have slots, replaceable elements. This means that frames can be linked to other frames by sharing participants or even by being participants in other frames. They can be components of an interpretation.

In frame semantics, all word meanings are relativized to frames. But a word meaning doesn't activate an entire frame. Different words select different aspects of the background to *profile* (here we use the terminology introduced in Langacker 1984). Sometimes those different aspects are just mutually exclusive parts of the kinds of circumstances being described, like distinct participants, such as the husband and wife in the marriage frame. But sometimes word meanings differ not in what they profile, but in how they profile it. In such cases, I will say words differ in *perspective* (Fillmore 1977a). Fillmore's much-discussed commercial event example (Fillmore 1976) is a case in point:

- (7) a. John sold the book to Mary for \$100.
- b. Mary bought the book from John for \$100.
- c. Mary paid John \$100 for the book.
- d. John collected \$100 for the book from Mary.

Verbs like *buy*, *sell*, *pay*, and *collect* have as background the concept of a commercial transaction, an event in which a *buyer*, gives *money* to a *seller* in exchange for some *goods*. Now because the transaction is an exchange it can be thought of as containing what Fillmore calls two *subscenes*: a *goods_transfer*, in which the goods is transferred from the seller to the buyer, and a *money_transfer*, in which the money is transferred from the buyer to the seller. Here it is natural to say that English along with many languages has as a valence realization possibility that, in transfers of possession, the object being transferred from one possessor to another is realized as direct object. Thus verbs profiling the money transfer will make the money the direct object (*pay* and *collect*) and verbs profiling the goods transfer will make the goods the direct object (*buy* and *sell*). Then the difference between these verb pairs can be chalked up to what is profiled.

But what about the difference between *buy* and *sell*? Surely there is one. By hypothesis, both verbs profile a goods transfer, but in one case the buyer is subject and in another the seller is. Perhaps this is just an arbitrary choice the grammar allows. This is in some sense what the thematic role theory of Dowty (1991) says: Since (7a) and (7b) are mutually entailing there can be no semantic account of the choice of subject.

In frame semantics, however, we may attempt to describe the facts as follows: in the case of *buy* the buyer is viewed as (*perspectivalized* as) agent, in the case of *sell*, the seller is. There are two immediate advantages to this

description. First, of course, it allows us to preserve a principle assumed by a number of linguists, that cross-linguistically agents must be subjects. Second, under the assumption that agents are understood as controllers of instruments, it explains why the following two sentences cannot be understood as paraphrases:⁴

- (8) a. John bought the book from Mary with his last pay check.
 b. Mary sold the book to John with his last paycheck.

Only the first sentence allows the reading of the instrumental *with his last paycheck* on which the pay check provides the funds enabling the purchase. In general the semantic consequence of calling a participant an agent is both a set of entailment facts and a set of constraints on the interpretation of adjuncts like instrumentals, benefactives, and purpose clauses:

- (9) a. John bought the house from Sue for Mary. [allows reading on which Mary is ultimate owner, disallows the reading on which Mary is seller and Sue is seller's agent]
 b. Sue sold the house to John for Mary. [allows reading on which Mary is seller and Sue is seller's agent; disallows reading on which Mary is ultimate owner.]
 c. John bought the house from Sue to evade taxes/as a tax dodge. [tax benefit is John's]
 d. Sue sold the house to John to evade taxes/as a tax dodge. [tax benefit is Sue's]

But what does it mean to say that a verb takes a perspective which “views” a particular participant as an agent? The facts are, after all, that both the buyer and the seller are agents; they have all the entailment properties that characterize what we typically call agents; and this, Dowty's theory of thematic roles tells us, is why verbs like *buy* and *sell* can co-exist. I will have more to say on this point in Section 4; for the moment I will confine myself to the following general observation on what frame semantics allows: What is profiled and what is left out by a verb meaning is not determined by the entailment facts of its frame. Complex variations are possible. The

⁴One might try to say that the instrument control property here is a property of subjects not agents, but this would make instrumentals rather unique. Unlike other subject control adverbials, passivization has no effect on the control possibilities of instruments:

- (i) # John was sold the book with his last paycheck. (just as odd as (8b))
 (ii) The missionaries were willingly sacrificed. ['missionaries willing' interpretation available]

frames themselves do not even determine a particular subscene structure to constrain the subsets of participants which are chosen as arguments. For example, as Fillmore observes, the COMMERCIAL TRANSACTION frame is associated with verbs that have no natural way of realizing the seller:

(10) John spent \$100 on that book.

Nothing in the valence marking of the verb *spend* suggests that what is being profiled here is a possession transfer; neither the double object construction, nor *from* nor *to* is possible for marking a core COMMERCIAL TRANSACTION participant. Rather the pattern in use here seems to be available for one might call *resource consumption* verbs like *waste*, *lose*, *use (up)*, and *blow*. In this profiling, there is no room for a seller and none is allowed. Given that such rich variation in what is abstracted out of a frame is allowed, the idea that the agenthood of a participant might be part of what's included or left out does not seem so far-fetched. As I will argue in Section 4, the inclusion of events into the semantics can help us make semantic sense of what abstractions like this might mean.

Given this machinery, one obvious way in which frames can be of service in lexical semantic description is in the account of polysemy. Different senses will in general involve relativization to different frames. As a very simple example, consider the noun *hit* as used in baseball (a certain kind of event in which a batter is ruled safe after striking the ball into the field of play) or as used in fencing (a legal point-scoring strike on one's opponent). These senses need to be defined with respect to two distinct complex frames, a baseball frame and a fencing frame. The examples involve what might be called technical language, but from the frame semantics point of view, technical language is not really such a special case. The great specificity of technical senses is due simply to the complexity of the frames involved.

Next consider a non-technical example, the use of *spend* in the following sentence:

(11) John spent 10 minutes fixing his watch.

How are we to describe the relationship of the use of *spend* in this example, which basically describes a watch fixing event, with that in (10), which describes a COMMERCIAL TRANSACTION? One way is to say that one sense involves the COMMERCIAL TRANSACTION, and another involves a frame we might call ACTION DURATION which relates actions to their duration, a frame that would also be invoked by durative uses of *for*. A counter-proposal is that there is one sense here, which involves an actor using up a resource.

But such a proposal runs up against the problem that *spend* really has rather odd disjunctive selection restrictions:

(12) John spent 30 packs of cigarettes that afternoon.

Sentence (12) is odd except perhaps in a context (such as a prison or boarding school) where cigarette packs have become a fungible medium of exchange; what it cannot mean is that John simply used up the cigarettes (by smoking them, for example). The point is that a single general resource consumption meaning ought to freely allow resources other than time and money, so a single resource consumption sense does not correctly describe the readings available for (12); however, a sense invoking a COMMERCIAL TRANSACTION frame constrained to very specific circumstances does. Note also, that the fact that 30 packs of cigarettes can be the *money* participant in the right context is naturally accommodated. The right constraint on the money participant is not that it be cash (for which Visa and Mastercard can be thankful), but that it be a fungible medium of exchange.

A COMMERCIAL TRANSACTION is a very specific kind of scene, leading to a very specific sense for a word like *spend*. How specific can frames be?

Consider the word *clammy*. Farlex the online dictionary defines this as *disagreeably moist, sticky, and cold to the touch*; this definition makes reference to 3 scales, moistness, stickiness, and coldness, and possibly a 4th (disagreeableness). Is this a 3 (4?) frame word? Very likely, this is an example of a one-word frame that needs 4 other frames as background for its definitions. The key point is that *clammy* still behaves like a scalar adjective with one scale, allowing, for example, comparatives like *John's hand was clammier than Sue's*. That in turn presupposes some single well-ordering of clamminess. What these observations establish is that the background required for *clammy* is quite complex, and does not appear to be shared by any other words. I assume then that this is a plausible example of a one-word frame.

The answers to our question then is: Frames can be very specific indeed. For example, it makes sense to have a specialization of the HEIGHT frame which includes just the polar adjectives, as well as a specialization of that which includes just *short*. Similarly, one can have specializations of the CALENDAR CYCLE frame that include just the days of the week, the the months, the day parts (morning, afternoon, evening, night) and the seasons.

Summarizing:

1. Frames are motivated primarily by issues of understanding and converge with various schema-like conceptions advanced by cognitive psy-

chologists, AI researchers, and cognitive linguists. They are experientially coherent backgrounds with variable components that allow us to organize families of concepts.

2. The concept of frames has far reaching consequences when applied to lexical semantics, because a single frame can provide the organizing background for a set of words. Thus frames can provide an organizing principle for a rich open lexicon. FrameNet is an embodiment of these ideas.
3. In proposing an account of lexical semantics rich enough for a theory of understanding, frame semantics converges with other lexical semantic research which has been bringing to bear a richer set of concepts on problems of the syntax semantics interface.

Having sketched the basic idea, I want in the next two sections to briefly contrast the notion frame with two other ideas that have played a major role in semantics, the idea of a relation, as incorporated via set theory and predicate logic into semantics, and the idea of a lexical field.

2 Comparing

In this section I compare the idea of frames with two other concepts of major importance in theories of lexical semantics, relations and lexical fields. The comparison offers the opportunity to develop some other key ideas of frame semantics, including profiling and saliency.

2.1 Frames versus relations: Profiling and saliency

Words (most verbs, some nouns, arguably all degreeable adjectives) describe relations in the world. Love and hate are relations between animate experiencers and objects. The verb *believe* describes a relation between an animate experiencer and a proposition. These are commonplace views among philosophers of language, semanticists, and syntacticians, and they have provided the basis for much fruitful work. Where do frames fit in?

For Fillmore, frames describe the factual basis for relations. In this sense they are “pre-”relational. To illustrate, Fillmore (1985) cites Mill’s (1847) discussion of the words *father* and *son*. Although there is a single history of events which establishes both the father- and the son- relation, the words *father* and *son* pick out different entities in the world. In Mill’s terminology, the words *denote* different things, but *connote* a single thing,

the shared history. This history, which Mill calls the *fundamentum relationis* (the foundation of the relation), determines that the two relations bear a fixed structural relation to each other. It is the idea of a determinate structure of relations that Fillmore likens to the idea of a frame.

Thus, a frame defines not a single relation but, minimally, a structure of relations.

This conception allows for a natural description not just of pairs of words like *father* and *son*, but also of single words which do not in fact settle on a particular relation. Consider the verb *risk*, discussed in Fillmore and Atkins (1998), which seems to allow a range of participants into a single grammatical “slot”. For example,

- (13) Joan risked $\left\{ \begin{array}{l} \text{a.} \text{ } \text{censure.} \\ \text{b.} \text{ } \text{her car.} \\ \text{c.} \text{ } \text{a trip down the advanced ski slope.} \end{array} \right.$

The RISK frame has at least 3 distinct participants, (a) the bad thing that may happen, (b) the valued thing that may be lost, and (c) the activity that may cause the bad thing to happen. All can be realized in the direct object position, as (13) shows. Since, there are three distinct relations here, a theory that identifies lexical meanings with relations needs to say there are 3 meanings as well. Frame semantics would describe this as one frame allowing 3 distinct profilings. It is the structure of the frame together with the profiling options the language makes available which makes the 3 alternatives possible.

Other verbs with a similar indeterminacy of participant are *copy*, *mix*, and *hit*:

- (14) a. Sue copied her costume (from a film poster).
 b. Sue copied the film poster.
 c. John mixed the soup.
 d. John mixed the paste into the soup.
 e. John mixed the paste and the flour.
 f. Fred hit the fence with the stick.
 g. Fred hit the stick against the fence.

In each of these cases the natural frame semantics account would be to say the frame remains constant while the profilings or perspective changes. Thus, under a frame semantics approach, verbal valence alternations are to be expected, and the possibility of such valence alternations provides motivation for the idea of a background frame with a range of participants and a range of profiling options.

At the same time the idea of profiling raises certain theoretical challenges. The following examples are due to David Dowty:

- (15) a. The truck collided with the lamp post.
b. # The truck and the lamp post collided.

Now we would like to say the sentences in (15) differ in just the way the sentences in (14) did: Different profiling options have been taken. Yet these are not the only differences. Sentence (15b) is odd because, in contrast to (15a), it seems to attribute movement to a lamp post. Using a different terminology (with “scenes” substituting for “frames”) but assuming essentially the same idea of profiling developed elsewhere, Fillmore 1976, 1977b introduces the idea of a Saliency Hierarchy, including features such as definiteness, motion, animacy, and state change. For a participant to manifest any of these features promotes or even secures its being *perspectivalized*, where being perspectivalized as a *participant* specifically means becoming subject or object.⁵ Let us assume that if a frame entails movement, then, one of the profiled participants must be moving (unless outranked by some more salient participant); it follows that the only interpretive option for the single collective participant of (15b) is that the collection be moving. In contrast, in (15a), only the truck is profiled, so only the truck is entailed to be moving. Movement is optional, however, for the participant expressed in an oblique.

Now on a theory in which senses are relations, the verbs in (15) must have different senses. This is because the sentences in (15a) and (15b) have different numbers of arguments and different truth-conditions. Frame semantics allows another option. We can say the same verb sense is used in both cases. The differences in interpretation arise because of general constraints on perspectivalization.

2.2 Frames versus lexical fields

Because frames define lexical sets, it is useful to contrast the concept of frames with an earlier body of lexical semantic work which takes as central the identification of lexical sets. This work develops the idea of **lexical fields** (Weisgerber 1962, Coseriu 1967, Trier 1971, Geckeler 1971). Lexical fields define sets of lexical items in mutually defining relations, in other words, lexical semantic paradigms. The classic example of a lexical field is the

⁵Dowty (1991) develops a feature based theory of subject and object selection with different assumptions, but many of the same features.

set of German labels used for evaluating student performance (Weisgerber 1962:99):

(16) *sehr gut, gut, genügend* and *mangelhaft*

The terms are mutually defining because the significance of a single evaluation obviously depends on knowing the entire set and the relations of the terms in the set. Thus *gut* means one thing in a school system with the 4 possibilities in (16) and quite another if the possibilities are:

(17) *sehr gut, gut, befriedigend ausreichend, mangelhaft* and *ungenügend*

Fillmore also cites the example of the tourist industry use of the term *first class* in their categorization of hotels; to many travelers, *first class* sounds pretty good; in fact, the top ranked class of hotels is *luxury* and *first class* is fourth from the top. The misunderstanding here seems exactly like a case of applying the wrong frame in the process of understanding.

Domains in which lexical fields have provided fruitful analyses include color, temperature, furniture and artifacts, kinship relations, intelligence, livestock, and terrain features (Fillmore 1985:227).

The general hypothesis of lexical field theory is that the lexicon can be carved up into a number of (sometimes overlapping) lexical sets, each of which functions as a closed system. To this extent, there is agreement with the conception of frames, and in fact, the lexical sets associated with frames can include lexemes in paradigmatic, mutually defining relations. For example, we identified the TEMPERATURE frame in Section 1, and this includes the lexical field of temperature words like *cold*, *cool*, *lukewarm*, *warm*, and *hot*.

However, the idea of a frame is distinct from the idea of a lexical field. To start with, the idea of a one-word lexical field is incoherent: How can a word have a function in a field in which there is nothing for it to be opposed to? However, there is no inherent difficulty with the idea of a one-word frame. We have already encountered the example of *clammy*. Fillmore (1985) cites the example of *hypotenuse*, which requires for its background the concept of a right triangle. There appear to be no other English lexical items specific to right triangles (the term *leg* in the relevant sense seems to apply to triangle sides in general); and that is neither surprising nor problematic. The notion mutually defining is not necessary for lexical frame sets because words in frames are defined in contrast to or in terms of the frame alone. The frame, not its lexical instantiations, provides the background necessary to identify a semantic function. The primitive notion is not *defined in opposition to* but *profiled from the background of*.

Pushing this idea one step further, the idea of a no-word frame is also coherent. The idea of a frame is in no way limited to lexical meaning. Thus if the word *hypotenuse* did not exist in English, the right-triangle frame would still be called on to provide background for discourses about trigonometry and the Pythagorean theorem.

A second way in which frames differ from lexical fields is that, even when there is more than one word, there is no requirement that words in the set function in paradigmatic opposition to one another. Thus the TEMPERATURE frame cited above also contains the noun *temperature*, just as the HEIGHT frame containing polar adjectives like *tall* and *short* will contain the noun *height*.

Thirdly, because of the notion of mutual definition, lexical fields come with strict criteria of individuation. In contrast, as we saw in Section 1, frames of arbitrary specificity make sense. Thus, we have very general frames of TEMPERATURE and HEIGHT. But we also have a set of specific frames that recover the traditional mutually defining sets that preoccupied lexical field theorists, a specialization of HEIGHT that includes just the polar adjectives, a specialization of TEMPERATURE that includes just the set *cold*, *cool*, *warm*, *hot*, and so on. This level of specificity in fact roughly describes the granularity of FrameNet.

3 Minskian frames

As described in Fillmore (1982), the term *frame* is borrowed from Marvin Minsky. It will be useful before tackling the issue of how perspectivalization works to take a closer look at this precursor.

In Minsky's original frames paper (Minsky 1975), frames were put forth as a solution to the problem of *scene interpretation* in vision. Minsky's proposal was in reaction to those who, like the Gestalt theorists (Koffka 1963), viewed scene perception as a single holistic process governed by principles similar to those at work in electric fields. Minsky thought scenes were assembled in independent chunks, constituent by constituent, in a series of steps involving interpretation and integration. To describe this process, a model factoring the visual field into a number of discrete chunks, each with its own model of change with its own discrete phases, was needed.

A frame was thus a dynamic model of some specific kind object with specific participants and parameters. The model had built-in expectations about ways in which the object could change, either in time or as a viewer's perspective on it changed, formalized as operations mapping old frame states

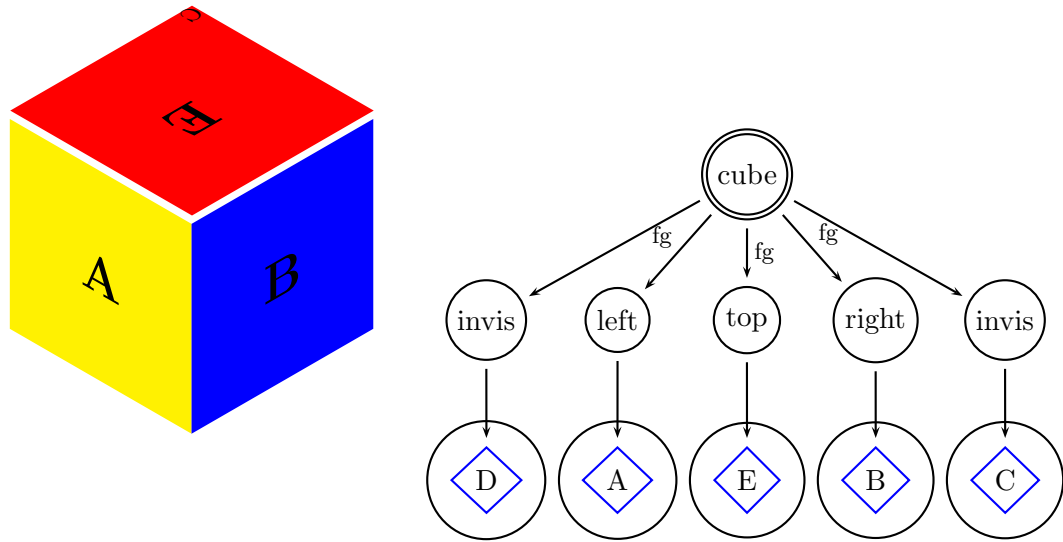


Figure 1: View of cube together with simplified cube frame representing that view. Links marked “fg” lead to foregrounded slots; slots marked “invis” are backgrounded. Faces D and C are out of view.

to new frame states. A frame also included a set of *participants* whose status changed under these operations; those moving into certain distinguished slots are *foregrounded*. Thus, for example, in the simplified version of Minsky’s *cube frame*, shown before and after a rotation in Figures 1 and 2, a frame state encodes a particular view of a cube and the participants are cube faces. One possible operation is a rotation of the cube, defined to place new faces in certain view-slots, and move old faces out and possibly out of view. The faces that end up in view are the foregrounded participants of the resulting frame state. Thus the cube frame offers the tools for representing particular views or *perspectives* on a cube, together with the operations that may connect them in time.

Fillmore’s innovation, then, was to apply this Minskian idea in the domain of word meaning, importing not only the idea of chunked modular knowledge units, but also the idea of operations that take perspectives on such chunks. I used the terms *profiling* and *perpsectivalization* to describe such operations in Section 1. Although Fillmore himself does not attempt a formalization of these operations, I believe it is possible to clearly describe what is at issue using some ideas from event semantics (Davidson 1967, Davidson 1980, Parsons 1990), building on the event-based approach to frames in Gawron (1983).




Figure 2: Cube frame after counterclockwise rotation. Faces D and A are now foregrounded, B has moved out of view.

4 Events, Profiling, and Perspectivalization

To spell out the analogy between Minsky’s idea of perspectivalization and Fillmore’s, let us return to the case of the COMMERCIAL TRANSACTION frame discussed in Section 1. The following development takes up and extends the ideas of Gawron (1983).

A rather natural account of the interface between frames and compositional semantics becomes available if we make use of neo-Davidsonian event-semantics (Davidson 1967, Davidson 1980, Parsons 1990). On a neo-Davidsonian account, we have, as the schematic semantics for *John bought the book on sale*:

$$\exists e[\text{buy}'(e) \wedge \text{agent}(e) = j \wedge \text{patient}(e) = b \wedge \text{on-sale}(e, b)]$$

We call e in the above representation the **lexical eventuality**. We call

$$\text{agent}(e) = j$$

a **role function**. For simplicity, I will assume lexical roles are functional in the sense that they allow at most one filler per eventuality; but in principle non-functional roles (relational roles) are also possible.

I assume that Fillmorean frames can be related to events. That is, there is such a thing as a COMMERCIAL TRANSACTION event. Further, I assume that lexical predicates like *give* and *buy* are predicates true of events. These

lexical events cannot be directly identified with Fillmorean frame events. Rather the lexical events are profilings or perspectivalizations of Fillmorean frame events. Thus, for example, buying will be associated with two events, one, a **perspectivalizing** event that is directly related to syntactic realization, the other a COMMERCIAL TRANSACTION (the Fillmorean frame event). I will call this latter the **circumstance** event. Perspectivalizing events and circumstance events will be related by functions called perspectivalizing or profiling functions.

Borrowing the formal machinery of sorted logic,⁶ I will assume that all predicates and role relations are **sorted**; that is, it is a property of a predicates and relations that in all models, for any given argument position, there is a certain sort of individuals for which that argument position is **defined**. I will write sorts in boldface and predicates true of them in roman.

- (18) AGENT PATIENT : **agent_patient** \mapsto **truth-values**
agent : **agent_patient** \mapsto **animate**
patient : **agent_patient** \mapsto **entity**
source : **agent_patient** \mapsto (**entity**)
goal : **agent_patient** \mapsto (**entity**)

These declarations just say, in roughly standard mathematical notation that agent and patient are functions from one set to another. For example, the first declaration says that AGENT PATIENT is a function from the set (sort) to truth-values;. the second says *agent* is a function from the set (sort) of AGENT PATIENT events to animates; *patient* from the set of AGENT PATIENT events to the set of things (the domain of entities). The parentheses in the source and goal role definitions may be taken to mean that the role is optional (or the function is *partial*). Not every AGENT PATIENT event has a source or a goal, but some do.

I assume the declarations in (18) are sufficient to define a very simple kind of frame. The first declaration defines a predicate AGENT PATIENT that is true of events of that sort; the rest define a set of roles for that sort of event. Thus a minimal frame is just an event sort defined for a set of roles. I will call AGENT PATIENT an **argument frame** because syntactic arguments

⁶Sorts can be viewed as an extension of a system of types. The most relevant development of typed/sorted logic can be found in the literature on the logic of typed feature structures (Carpenter 1992, Smolka 1992). In Carpenter, an explicit *appropriateness* function is defined for a partially ordered system of types. See Rounds (1997) for an excellent survey. Throughout we will switch back and forth between constraints in a sorted logic and typed-feature structures, because the two formalisms are very closely connected. Copestake et al. (1988) apply typed feature structures to lexical semantics .

of a verb will need to directly link to the roles of argument frames (such as *agent* and *patient*). We can represent this set of axioms as an attribute-value structure:

$$(19) \left[\begin{array}{ll} \text{AGENT PATIENT} & \\ \text{agent} & \mathbf{animate} \\ \text{source} & \mathbf{entity} \\ \text{goal} & \mathbf{entity} \\ \text{patient} & \mathbf{entity} \end{array} \right]$$

I will use the attribute-value notation for its readability, but the reader should bear in mind that it is merely a shorthand. Throughout, the intention is that the actual axiomatic development of a lexical system is captured using axioms like those in (18), with the essential primitive notion being constraints on partial functions and relations from sorts to sorts,

We can just envision AGENT PATIENT events as very general event-types; that is, the kind of interpretation I have in mind for the agent and patient roles is the sort of *proto-role* interpretation Dowty (1991) uses. As we will see, the plasticity of their semantics follows from the fact that in order to use them, a lexical item must specify some circumstance frame in which participant roles are further specified.

Such declarations will be the basis for a treatment of semantic compatibility of modifiers and heads. When a role relation is predicated of an eventuality not in its appropriateness sort, the role relation is undefined:

$$(20) \quad \begin{array}{l} \text{a. } * \text{ John ate to school.} \\ \text{b. } \exists e[\text{eat}(e) \wedge \text{eater}(e) = j \wedge \text{to}(e) = \text{school}] \end{array}$$

Eating eventualities do not have a *to* role defined for them, so (20b) has no defined interpretation. Formally this is how argument frames will be linked to syntactic realization.

I will illustrate the connection between an argument frame like AGENT PATIENT with simple circumstance frames through the example of the POSSESSION TRANSFER frame (related to verbs like *give*, *get*, *take*, *receive*, *acquire*, *bequeath*, *loan*, and so on). Represented as an AVM, this is:

$$\left[\begin{array}{ll} \text{POSSESSION TRANSFER} & \\ \text{donor} & \mathbf{animate} \\ \text{possession} & \mathbf{entity} \\ \text{recipient} & \mathbf{animate} \end{array} \right]$$

Now both *give* and *acquire* will be defined in terms of the POSSESSION TRANSFER frame, but *give* and *acquire* differ in that with *give* the *donor* becomes subject and with *acquire* the *recipient* does. (Compare the difference between *buy* and *sell* discussed in in Section 1.2.)

We will account for this difference by saying that *give* and *acquire* have different mappings from the argument frame (AGENT PATIENT) to their shared circumstance frame (POSSESSION TRANSFER). This works as follows.

We define the relation between a circumstance and argument frame via a *perspectivalizing function*. Here are the axioms for what we will call the *acquisition* function, on which the recipient is agent:

- (21) (a) *acquisition* : **possession_transfer** → **agent_patient**
 (b) *agent* ∘ *acquisition* = *recipient*
 (c) *patient* ∘ *acquisition* = *possession*
 (d) *source* ∘ *acquisition* = *donor*

The first line defines *acquisition* as a total mapping from the sort **possession_transfer** to the sort **agent_patient**, that is as a mapping from POSSESSION TRANSFER eventualities to AGENT PATIENT eventualities; each POSSESSION TRANSFER is guaranteed to have an AGENT PATIENT eventuality associated with it. In the second line, the symbol ∘ stands for **function composition**; the composition of the *agent* function with the *acquisition* function (written *agent* ∘ *acquisition*) is the same function (extensionally) as the *recipient* relation. Thus the filler of the *recipient* role in a possession transfer must be the same as the filler of the *agent* role in the associated AGENT PATIENT eventuality. And so on, for the other axioms. Summing up these axioms AVM-style:

$$\left[\begin{array}{ll} \text{POSSESSION TRANSFER} & \\ \text{donor} & \boxed{1} \\ \text{recipient} & \boxed{2} \\ \text{possession} & \boxed{3} \end{array} \right] \xrightarrow{\textit{acquisition}} \left[\begin{array}{ll} \text{AGENT PATIENT} & \\ \text{agent} & \boxed{2} \\ \text{source} & \boxed{1} \\ \text{patient} & \boxed{3} \end{array} \right]$$

I will call the mapping that makes the *donor* agent *donation*. Written AVM style, it is:

$$\left[\begin{array}{ll} \text{POSSESSION TRANSFER} & \\ \text{donor} & \boxed{1} \\ \text{recipient} & \boxed{2} \\ \text{possession} & \boxed{3} \end{array} \right] \xrightarrow{\textit{donation}} \left[\begin{array}{ll} \text{AGENT_PATIENT} & \\ \text{agent} & \boxed{1} \\ \text{goal} & \boxed{2} \\ \text{patient} & \boxed{3} \end{array} \right]$$

With the *acquisition* and *donation* mappings defined, the predicates *give* and *acquire* can be defined as compositions with *donation* and *acquisition*:

$$\begin{aligned} \text{give}' &= \text{POSSESSION TRANSFER} \circ \text{donation}^{-1} \\ \text{acquire}' &= \text{POSSESSION TRANSFER} \circ \text{acquisition}^{-1} \end{aligned}$$

donation^{-1} is an inverse of donation^7 , a function from AGENT PATIENT eventualities to POSSESSION TRANSFERS. It is defined only for those AGENT PATIENT events related to POSSESSION TRANSFERS. Composing this with the POSSESSION TRANSFER predicate makes *give'* a predicate true of those AGENT PATIENT events related to possession transfers, whose agents are donors and whose patients are possessions. The treatment of *acquire* is parallel but uses the *acquisition* mappings.

Summarizing the assumptions:

- a. an argument frame AGENT PATIENT, with direct consequences for syntactic valence (agent becomes subject; patient direct object; obliques are restricted).
- b. a circumstance frame POSSESSION TRANSFER, which captures the circumstances of possession transfer.
- c. **perspectivalizing** functions *acquisition* and *donation* which map participants in the circumstances to argument structure.

This is the basic picture of perspectivalization. The picture becomes more interesting with a richer example.

In the discussion that follows, I assume a commercial transaction frame with at least the following frame elements:

$$(22) \quad \left[\begin{array}{ll} \text{COMMERCIAL TRANSACTION} & \\ \text{buyer} & \mathbf{animate} \\ \text{seller} & \mathbf{animate} \\ \text{money} & \mathbf{fungible} \\ \text{goods} & \mathbf{entity} \end{array} \right]$$

This is a declaration that various functions from event sorts to truth values and entity sorts exist, a rather austere model for the sort of rich backgrounding function we have assumed for frames. We will return to the richness of this model below.

Our picture of profiling and perspectivalization can be extended to the more complex cases of commercial transaction predicates with one more composition. For example, we may define *buy'* as follows:

⁷For the exact definition of \uparrow , see Gawron (2008).

$$(23) \quad \text{buy}' = \text{COMMERCIAL TRANSACTION} \circ (\text{acquisition} \circ \text{goods_transfer})^{-1}$$

What this says is that the relation *buy'* is built in a series of steps, out of 3 functions:

1. *acquisition*: the function from possession transfer events to AGENT_PATIENT events already introduced.
2. *goods_transfer*: a new function from commercial events to possession transfers:

$$\left[\begin{array}{l} \text{COMMERCIAL TRANSACTION} \\ \text{buyer} \quad \boxed{1} \\ \text{seller} \quad \boxed{2} \\ \text{money} \quad \boxed{3} \\ \text{goods} \quad \boxed{4} \end{array} \right] \xrightarrow{\text{goods_transfer} \sqsubseteq} \left[\begin{array}{l} \text{POSSESSION TRANSFER} \\ \text{recipient} \quad \boxed{1} \\ \text{donor} \quad \boxed{2} \\ \text{possession} \quad \boxed{4} \end{array} \right]$$

in these possession transfers, it is the goods which is being transferred, the buyer who is recipient, and the seller who is donor.

3. The composition of *goods_transfer* with *acquisition*

$$\text{acquisition} \circ \text{goods_transfer}$$

is a function from commercial transactions to associated agent_patient events

4. This function

$$(\text{acquisition} \circ \text{goods_transfer})^{-1}$$

is therefore a function from agent patient events to commercial transactions.

5. COMMERCIAL TRANSACTION: a sortal predicate true of commercial transactions.
6. *buy'* is therefore a predicate true of AGENT PATIENT events that are related in certain fixed ways to a POSSESSION TRANSFER and a COMMERCIAL TRANSACTION

The novelty in the definition above is the *goods_transfer* function. Conceptually, *goods_transfer* selects out the parts of the COMMERCIAL TRANSACTION event which the verb highlights. We will call this the **profiling function**. The function which determines subject object (the perspectivalizing function) is, as with *acquire*, *acquisition*. The perspectivalizing function is probably the most interesting object from a linguistic theoretical perspective, because it is a mapping to argument structure, but both the

perspectivalizing function and the profiling function are needed to define *buy*.

There are two independent reasons for positing a subscene role like *goods_transfer*:

a. It enriches our rather impoverished model of COMMERCIAL TRANSACTION. We started out in (22) with little more than the assumption that there were 4 sorted participants we were calling buyer, seller, money, and goods. Now with the assumption of the *goods_transfer* function, a possession transfer p is entailed (because the function is total). Paraphrasing what the definition of the function *goods_transfer* tells us in English we have:

- (a) The recipient in p is the buyer;
- (b) the donor in p is the seller;
- (c) the possession in p is the goods.

Thus *goods_transfer* can be viewed as part of an extended definition of the COMMERCIAL TRANSACTION frame. There will be other total functions that tell us more, for example, a *money_transfer* function that will be of use in the definition of verbs like *pay* and *collect*, which tells us the money goes from buyer to seller.

b. Both MONEY_TRANSFER and GOODS_TRANSFER are projections from commercial events to possession transfers; and possession transfer is a frame for which we have a pre-defined perspectivealization, independently motivated for other verbs like *acquire* and *get*. Thus by composing a commercial event subscene projection with a possession transfer argument projection we have an argument projection for commercial transactions. Thus the *good transfer* function simultaneously serves knowledge representation needs (a) and valence theory needs (b).


Saying that a buying event is the goods transfer of a COMMERCIAL TRANSACTION with the buyer perspectivealized as agent says all we need to say about buying. Everything else can be explained in terms of what we know about COMMERCIAL TRANSACTIONS, POSSESSION TRANSFERS, and agents. For example, being an AGENT PATIENT instance linked to a POSSESSION TRANSFER will license a certain argument structure and certain English valence possibilities, including the double object construction and the use of *for* phrases to refer to something exchanged.⁸

At this point, we have sketched enough of the picture of how profiling

⁸I leave out here a detailed treatment of the double object construction, obliques like *from NP* and *for NP* with *buy*. Interested readers should consult (Gawron 2008).

and perspectivalization work to complete the analogy to Minskyan frames. A Minskyan frame enables the integration of scene components in view with underlying objects by specifying, for example, how the faces of the cube in view relate to the cube as a whole. A Fillmorian perspective enables the integration of the realized elements of a text with an underlying text interpretation by specifying how syntactically realized frame components relate to frames as a whole. In both cases there are operations that mediate between rich representations and a constrained (perspectivalized) representation that belongs to an external representational system. Minskyan rotation operations mediate between 3D representations and the 2D representations of a scene. Ultimately the 2D representations are necessary because the human retina is a screen. Fillmorian profilings and perspectivalizations mediate between unlinearized representations in which there is no fixed individuation of participants and linearizable argument structure. Ultimately argument structure is linearizable because the syntax of human language forces us to individuate and linearize participants.

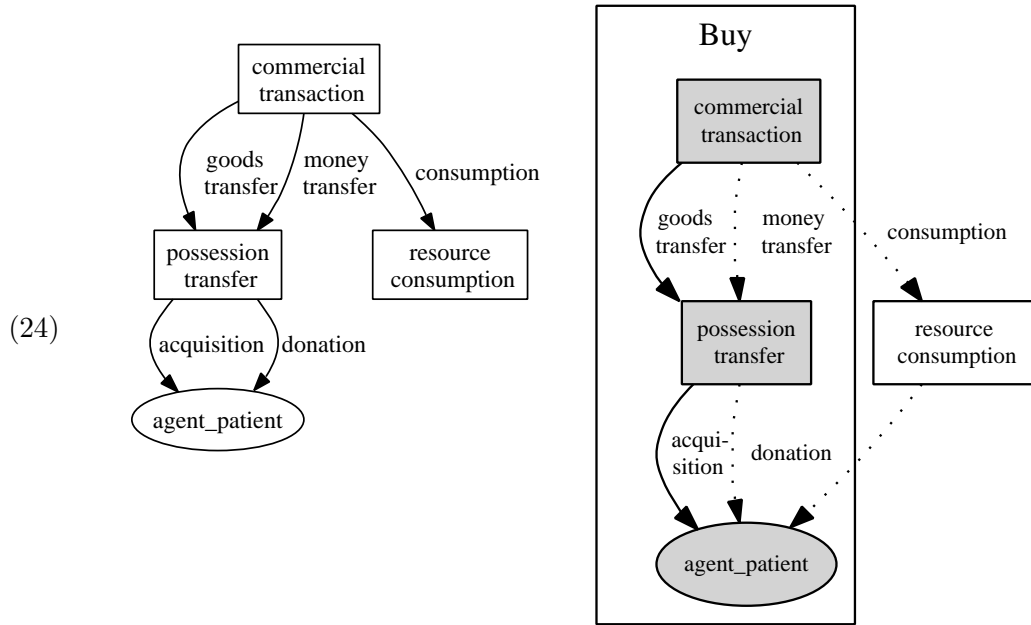
The power of these ideas can be seen when we start thinking about how things get left out in perspectivalization. Consider the case of *spend*.⁹



The verb *spend* views a commercial transaction as a RESOURCE CONSUMPTION, where RESOURCE CONSUMPTION is the frame used by verbs like *waste*, *lose*, *use (up)*, and *blow*. Notice *spend* and *sell* carve up COMMERCIAL TRANSACTIONS in different incompatible ways. The verb *spend* profiles a RESOURCE CONSUMPTION subscene which has the buyer, the money, and the goods; the seller is left out and cannot be realized, as discussed in Section 1. The verb *sell* profiles a subscene with the seller, the buyer and the goods. The two subscenes overlap in participants but choose distinct, incompatible event types, which leads to distinct realization possibilities in the syntactic frame.

⁹This is a slight oversimplification for expository simplicity. Actually, the goods is only *sometimes* mapped to the resource requirer role. The difficulty is that the resource requirer does not always come into possession of the buyer, as in *John spent \$300 on his car*. See Gawron (2008) for a more complete account.

The picture that has emerged of commercial transactions can be schematized as follows:



The picture on the left summarizes the frame to frame mappings that we have discussed. It shows what we might call the **commercial transaction neighborhood**. The picture on the right shows that portion of the neighborhood that is activated by *buy*; the functions used in its definitions are linked by solid lines; the functions left out are in dashes; the boxed regions contain those frames that are used in the definition.

The analysis of *buy* thus involves 3 frames

- a. The circumstance frame COMMERCIAL TRANSACTION
- b. The profiling frame POSSESSION TRANSFER
- c. The argument frame AGENT_PATIENT

The difference between the analysis of *give* and *buy* is that for *give* the circumstance frame and the profiling frame are the same. Basically *give* describes simple circumstances all of which are included in what the verb profiles.¹⁰

¹⁰We make the following observation. Lexical predicates can only use profilings for which some argument mapping exists. Thus the profiling frame and the argument frame do the kind of work Goldberg (1995) attributes to lexical constructions like the double object construction. I conclude that the relation between the original Fillmorean conceptions of

We have called *goods_transfer* and *consumption* profiling functions. We might equally well have called them *subscene roles*, because they fit the technical definition of role we have assumed. Note that subscene roles don't attribute a fixed hierarchical structure to a frame the way DO ... CAUSE BECOME .. in Dowty's system attributes a fixed structure to causatives of inchoatives. As these examples show, a frame may have subscene roles which carve up its constituents in incompatible ways.

Now this may seem peculiar. Shouldn't roles be fixed structural elements of a relation? I submit that the answer is no. First there are event sorts. The roles associated with each event sort are parts of our theory of that sort, regularities that help us classify an event as of a certain type; thus they are functions from the events to certain constituent types. But nothing about these functions guarantees they will carve up each event of that type into non-overlapping or hierarchically structured parts. For example, sometimes they are partial functions, because a participant may be an optional constituents of an event type. COMMERCIAL TRANSACTIONS as we have conceived them always have buyers, sellers, money, and goods, but surely the concept of a sales slip needs to be understood in the context of the same kind of event. And not every commercial transaction need have one. A sales slip is thus an optional role filler.

More confusingly, sometimes distinct role functions may select overlapping constituents of events. This is particularly possible when independent individuation criteria are not decisive, as when the constituents are collectives, or shapeless globs of stuff, or abstract things such as events or event types. Thus we get cases like *collide*, *mix*, and *risk*, where different ways of profiling the frames give us distinct, incompatible sets of roles. We may choose to view the colliders as a single collective entity (X and Y collided), or as two (X collided with Y).¹¹ We may choose to separate a figure from

profiling and perspctivalization is closer than sometimes thought.

¹¹I think this is also the right way to look at valence alternations characteristic of so-called symmetric verbs like *meet*:

- (i) John and Mary met.
- (ii) John met (with) Mary.
- (iii) John and Mary walked.
- (ii) John walked with Mary.

The general challenge is to characterize the kind of role function *with* denotes, which must pick out a subpart of some underlying frame element *x* which the agent role also picks out a subpart of. The difference between the (i),(ii) cases with *meet* and the (iii),(iv) cases with *walk*, is that in the case of *meet*, that frame element *x* **must** be a collective (plural entity). In the case of *walk*, it can be. Nouns like *enemy* pose the same problem, allowing both *The two enemies faced each other* and *an enemy of the state*.

a ground in the mixing event (mixing X into Y), or or lump them together (mix X and Y), or just view the mixed substance as one (mix Z). Finally, risks involve an action and a potential bad consequence, and for that class of risking events where the bad consequence is loss of some X, there is a function that allows us just to select what is lost:

$$r = \text{lost-thing} \circ \text{bad-consequence}$$

the function r is certainly a role of RISK, in the sense we have defined roles, but it is an optional role (a partial function) because, assuming *lost-thing* is only defined for LOSS events, r is only defined for those instances of RISK whose *bad-consequence* is a loss.

Have we dispensed with relations in this picture? Formally, we have replaced relations with event predicates, each of which is defined through some composed set of mappings to a set of events that will be defined only for some fixed set of roles. Clearly, for every lexical predicate, there is a corresponding relation, namely one defined for exactly the same set of roles as the predicate. Thus in the end the description of the kind of lexical semantic entity which interfaces with the combinatorial semantics is not very different. However the problems has, I believe, been redefined in an interesting way. Traditionally, discussion of the lexical-semantic/syntax interface starts with a relation with a predefined set of roles. This is the picture for example, that motivates the formulation of Chomsky's (1981) Θ -Criterion. However, a major point of frame semantics is that, for many purposes, it is useful to to look at the pre-relational, the foundations out of which a relation's roles are abstracted. This is the domain of frames.

5 Lexicography

A word about the application of frames to lexicography is in order. Frames impose a certain classificational scheme on the lexicon. There have been lexical organization schemes in lexicography before. Roget's Thesaurus is probably the best known. The valence dictionary approach of Longman's is another. Frames differ significantly in that they are not first oriented either to the task of synonym-identification or syntactic frame identification. One expects to find synonyms and antonyms in the same frame, of course, and many examples of valence similarity, but neither trend will be a rule. As we saw in Section 1, near synonyms like *land* and *ground* may belong to different frames, and understanding those frames is critical to proper usage. As we saw in our investigations of profiling and perspective, differences of

both kinds may result in very different valence options for verbs from the same frame. The value of the frame idea for lexicography is that it seems the most promising idea if the goal is to organize words according to usage. This of course is a hypothesis. FrameNet (Fillmore and Baker 2000) is a test of that hypothesis. Accordingly, frame entries are not just lists of words; they are connected with rich sets of examples gleaned from the British National Corpus illustrating frame element realizations in a variety of syntactic contexts. Interested readers will find a tour of the web site far more persuasive than any discussion here. The website also contains pointers to other large scale description efforts in other languages, including Japanese, German, and Spanish.

6 Relation between frames

At least the following kinds of relations between frames have cropped up in our discussion:

- a. Specializing: For example, we have suggested there is a subframe of `CAL-
ENDAR CYCLE` which just includes the days of the week.
- b. Perspectivalizing: A relation between two frames A and B in which A is viewed as an instance B. More formally, there are functions from instances of A to instances B, maintaining certain role correspondences. Discussed in Section 4.
- c. Definitional: For example, the `MARRIAGE` frame is invoked to define the `DIVORCE` frame.

Now specialization is obviously an important relation between frames. Within FrameNet, frames are often specialized so as to identify lexical sets that share circumstantial background, category, and valence properties such as transitivity. For example, there is both an `IMPACT` and a `CAUSE IMPACT` frame distinguishing uses of verbs like *hit* from those that have agents from those that do not (*the rock hit window*). This is both lexicographically and linguistically useful. For other examples of the utility of specialization relations in frames and linguistic see FrameNet (Fillmore and Baker 2000), which provides a graph tool to visualize the frame hierarchy. See also Gawron (1983), for a various examples of using specialization relations to capture valence generalizations in English.

The kind of relation we have had the least to say about is definitional relations. This is obviously the most open-ended case, and can be seen as including the other two. For example, the subscene roles *goods_transfer* and *money transfer* which capture profiling relations between `COMMERCIAL`

TRANSACTIONS and POSSESSION TRANSFERS can be seen as part of the definition of COMMERCIAL TRANSACTIONS. The IMPACT frame can be seen as part of the definition of the CAUSE IMPACT frame, and so on. But these two possibilities only scratch the surface. Frames by describing practices and social categories define new sorts of things in the world and those sorts enter into new definitions. The DIVORCE frame, with respect to the MARRIAGE frame, is a case in point, with new categories like *alimony* emerging. It is not clear how much light formalizing such possibilities would shed, but it is clear that understanding the recursive structure of the framing relations is a central part of a theory of text understanding.

7 Conclusion

The syntax of human languages, with some limited but very interesting exceptions in sign language, forces us to linearize the participants of the events we are talking about. The logical notion of a relation, which preserves certain aspects of linearization, has at times appeared to offer an attractive account of what we grasp when we grasp meanings. But the data we have been looking at in this brief excursion into frame semantics has pointed another way. Lexical senses seem to be tied to the same kind schemata that organize our perceptions and interpretations of the social and physical world. In these schemata participants are neither linearized nor uniquely individuated, and the mapping into the linearized regime of syntax is constrained but underdetermined. Thus we frequently see individual words with options in what their exact participants are and how they are realized or closely related words with distinct realization possibilities for the same participants. Frames offer a model that is both specific enough and flexible enough to accommodate these facts, while offering the promise of a firm grounding for lexicographic description and an account of text understanding.

We conclude with a quote from Fillmore. Note the careful placement of the negation in the last sentence:

People need to categorize objects and events in their world. When we wish to study instances of categorization provided by the lexical items in their language, we can do this only by asking what functions such categorizations have in their lives. Some of the categorizations we find have only linguistic explanations: people do it that way because that's how their language evolved, and it could have evolved in a number of other ways. Others have, at least in part, explanations that depend in crucial ways

on such matters as how humans perceive the world around them, how people form categorizations in general, and what social institutions form the matrix of their daily activities. We need theories of word meanings that will not encourage us to lose sight of these realities.

Fillmore (1975)

References

- Baker, Collin, and Charles J Fillmore. 2001. Frame semantics for text understanding. In *Proceedings of WordNet and Other Lexical Resources Workshop*, Pittsburgh, PA. NAACL.
- Baker, Collin, Charles J Fillmore, and John B Lowe. 1998. The berkeley framenet project. In B. T. S. Atkins and A. Zampolli (Eds.), *Proceedings of the COLING-ACL*, Montreal, Canada.
- Bartlett, F. C. 1932. *Remembering: An Experimental and Social Study*. Cambridge: Cambridge University Press.
- Berlin, Isaiah. 1997. *The Proper Study of Mankind*. New York: Farrar, Straus, and Giroux.
- Boas, Hans C. 2001. Frame semantics as a framework for describing polysemy and syntactic structures of english and german motion verbs in contrastive computational lexicography. In P. Rayson, A. Wilson, T. McEnery, A. Hardie, and S. Khoja (Eds.), *Proceedings of the Corpus Linguistics 2001 conference*, Lancaster, UK: University. Centre for computer corpus research on language.
- Boas, Hans C. 2002. Bilingual framenet dictionaries for machine translation. In M. G. Rodriguez and C. P. S. Araujo (Eds.), *Proceedings of the Third International Conference on Language Resources and Evaluation*, Vol. IV, 1364–1371, Las Palmas, Spain.
- Carpenter, B. 1992. *The Logic of Typed Feature Structures*. Cambridge, MA: Cambridge University Press.
- Chang, Nancy, Srini Narayanan, and Miriam R. L. Petruck. 2002a. From frames to inference. In *Proceedings of the First International Workshop on Scalable Natural Language Understanding*, Heidelberg, Germany.

- Chang, Nancy, Srinu Narayanan, and Miriam R. L. Petruck. 2002b. Putting frames in perspective. In *Proceedings of the Nineteenth International Conference on Computational Linguistics*, Taipei, Taiwan.
- Chomsky, Noam. 1981. *Lectures on Government and Binding*. Dordrecht: Foris.
- Copestake, Ann, A Sanfilippo, Ted Briscoe, and V de Paiva. 1988. The acquilex lkb: An introduction. In T. Briscoe, A. Copestake, and V. de Paiva (Eds.), *Inheritance, Defaults, and the Lexicon*. Cambridge: Cambridge University Press.
- Coseriu, Eugenio. 1967. Lexicale solidaritäten. *Poetica* 1:293–303.
- Davidson, Donald. 1967. The logical form of action sentences. In N. Rescher (Ed.), *The Logic of Decision and Action*. University of Pittsburgh Press. Also in (Davidson 1980), chapter 6.
- Davidson, Donald. 1980. *Essays on Actions and Events*. Oxford University Press.
- Dowty, David. 1979. *Word Meaning and Montague Grammar*. Boston: Reidel.
- Dowty, David R. 1991. Thematic proto roles and argument selection. *Language* 67:547–619.
- Fillmore, Charles J. 1975. An alternative to checklist theories of meaning. In C. Cogen, H. Thompson, G. Thurgood, K. Whistler, and J. Wright (Eds.), *Proceedings of the First Annual Meeting of the Berkeley Linguistics Society*, 123–131, The University of California at Berkeley. Berkeley Linguistics Society.
- Fillmore, Charles J. 1976. Topics in lexical semantics. In R. Cole (Ed.), *Current Issues in Linguistic Theory*, 76–138. Indiana University Press.
- Fillmore, Charles J. 1977a. The case for case reopened. In P. Cole and J. Sadock (Eds.), *Syntax and Semantics: Grammatical Relations*, Vol. 8, 59–81. Academic Press Inc.
- Fillmore, Charles J. 1977b. Scenes-and-frames semantics: Linguistic structures processing. In A. Zampolli (Ed.), *Fundamental Studies in Computer Science*. Amsterdam: North Holland.

- Fillmore, Charles J. 1978. On the organization of semantic information in the lexicon. In *Papers from the Parasession on the Lexicon*, 148–173, The University of Chicago. Chicago Linguistics Society.
- Fillmore, Charles J. 1982. Frame semantics. In T. L. S. of Korea (Ed.), *Linguistics in the Morning Calm*. Seoul: Hanshin Publishing Co.
- Fillmore, Charles J. 1985. Frames and the semantics of understanding. *Quaderni di Semantica* 6(2).
- Fillmore, Charles J, and B T S Atkins. 1994. Starting where the dictionaries stop: The challenge for computational lexicography. In B. T. S. Atkins and A. Zampolli (Eds.), *Computational Approaches to the Lexicon*. Clarendon Press.
- Fillmore, Charles J, and B T S Atkins. 1998. Framenet and lexicographic relevance. In *Proceedings of the First International Conference on Language Resources and Evaluation*, Granada, Spain.
- Fillmore, Charles J, and Collin Baker. 2000. FrameNet web site. <http://www.icsi.berkeley.edu/~framenet>.
- Gawron, Jean Mark. 1983. *Lexical Representations and the Semantics of Complementation*. New York: Garland.
- Gawron, Jean Mark. 2008. Circumstances and perspective: The logic of argument structure. unpublished ms.
- Geckeler, Horst. 1971. *Strukturelle Semantik und WortfeldTheorie*. München: Willhelm Fink Verlag.
- Goldberg, Adele. 1995. *Constructions: A Construction Grammar Approach to Argument Structure*. Chicago: University of Chicago Press.
- Jackendoff, Ray S. 1990. *Semantic Structures*. Cambridge, MA: The MIT Press.
- Koffka, K. 1963. *Principles of Gestalt Psychology*. New York: Harcourt, Brace and World.
- Lakoff, George. 1972. Linguistics and natural logic. In D. Davidson and G. Harman (Eds.), *Semantics of Natural Language*, 545–665. Dordrecht: Reidel.

- Lakoff, George. 1983. Category: an essay in cognitive linguistics. In *Linguistics in the Morning Calm*, 139–194. Seoul: Hanshin.
- Lakoff, George, and Mark Johnson. 1980. *Metaphors We Live By*. Chicago: University of Chicago Press.
- Langacker, Ronald. 1984. Active zones. In *BLS 10*, 172–188. Berkeley Linguistics Society, University of California at Berkeley, Department of Linguistics.
- Levin, Beth. 1993. *English Verb Classes and Alternations*. Chicago: University of Chicago Press.
- Mill, John Stuart. 1847. *A System of Logic*. New York: Harper and Brothers.
- Minsky, Marvin. 1975. A framework for representing knowledge. In P. Winston (Ed.), *The Psychology of Computer Vision*. McGraw-Hill.
- Parsons, Terence. 1990. *Events in the Semantics of English*. Cambridge, MA: MIT Press.
- Pusteyovsky, James. 1995. *The Generative Lexicon*. Cambridge: MIT Press.
- Rounds, William C. 1997. Feature logics. In J. van Benthem and A. ter Meulen (Eds.), *Handbook of Logic and Language*. Cambridge, MA: Elsevier.
- Rumelhart, Donald. 1975. Notes on a schema for stories. In *Representations and Understanding*. New York: Longmans.
- Schank, Roger C, and R P Abelson. 1977. *Scripts, Plans, Goals, and Understanding*. Hillsdale, N. J.: Erlbaum.
- Smolka, G. 1992. Feature constraint logics for unification grammars. *Journal of Logic Programming* 12:51–87.
- Trier, Jost. 1971. *Ausätze und Vorträge zur Wortfeldtheorie*. The Hague: Mouton.
- Weisgerber, J. Leo. 1962. *Grundzüge der Inhaltbezogenen Grammatik*. Düsseldorf: Pädagogischer Verlag Schwann.