Chapter 6

Semantic Relations and Types

The frames that we create, and thus the Frame Elements and Lexical Units associated with them, are intended to be situated in semantic space by means of frame-to-frame relations and semantic types. The relations we use include **Inheritance**, **Subframe**, **Causative_of**, **Inchoative_of**, and **Using**. There are several benefits to extensive use of these relations:

- Improving the comprehensibility of frames: The intended meaning of a more complex frame can be clarified by relating it to an existing, easily-understood frame.
- Robustness: In some cases, it would clearly be possible for other researchers (or ourselves) to have made different frame divisions than the ones we have made. Having relations to semantically similar frames allows frames (and thus their Lexical Units) to be associated despite being separated.

A number of other benefits are outlined in 6.3.

We have spent considerable time recently improving the semantic relations encoded in our data, and also making the relations more accessible via the FrameGrapher (see `http://framenet.icsi.berkeley.edu/\~{}FrameGrapher/`). We have also marked a large number of Frame Elements with semantic types indicating appropriate fillers (6.2.1). There remain a dwindling number of completely disconnected frames, all of which are semantically quite distinct from other frames that we already describe.

In addition, we have added two new types of relations, Precedes (replacing the meta-relations for Subframes) and Perspective_on, which replaces a subset of the Using relations with a more specific and informative relation. These should make our data more useful for all forms of computational processing. All of the frame relations referred to in this chapter *other than these new ones* can be visualized with the FrameGrapher on the main FrameNet website (`http://\~{}framenet.icsi.berkeley.edu`).

In the sections that follow, general descriptions of each of the Frame-to-frame relations (6.1) and semantic types (6.2) are given, and then a more formal description of the relations and the ways that they may be used for reasoning is discussed (6.3). Note that the following does not discuss the FE-to-FE relations, which have a closer relationship with valence and annotation than with semantically-defined relations. For these relations, see sections 3.2.2.1-3.2.2.3.

6.1 Frame-to-frame Relations

Each frame relation in the FrameNet data is a directed (asymmetric) relation between two frames, where one frame (the less dependent, or more abstract) can be called the Super_frame and another (the more dependent, or less abstract) can be called the Sub_frame. We give a more specific, informative name to the Super and Sub_frames for each relation, as shown in Table 6.1 below.

In general, each frame has one relation to some other frame, but there are occasional exceptions, as seen in the relations between Assistance and Intentionally_act, shown in Figure 6.1 below. Assistance inherits from Intentionally_act, with the HELPER bound to the AGENT of Intentionally_act, but Assistance also uses Intentionally_act since there is a second intentional action presupposed, namely that of the CO_AGENT, which is thus also bound to the AGENT role of Intentionally_act in a separate relation.

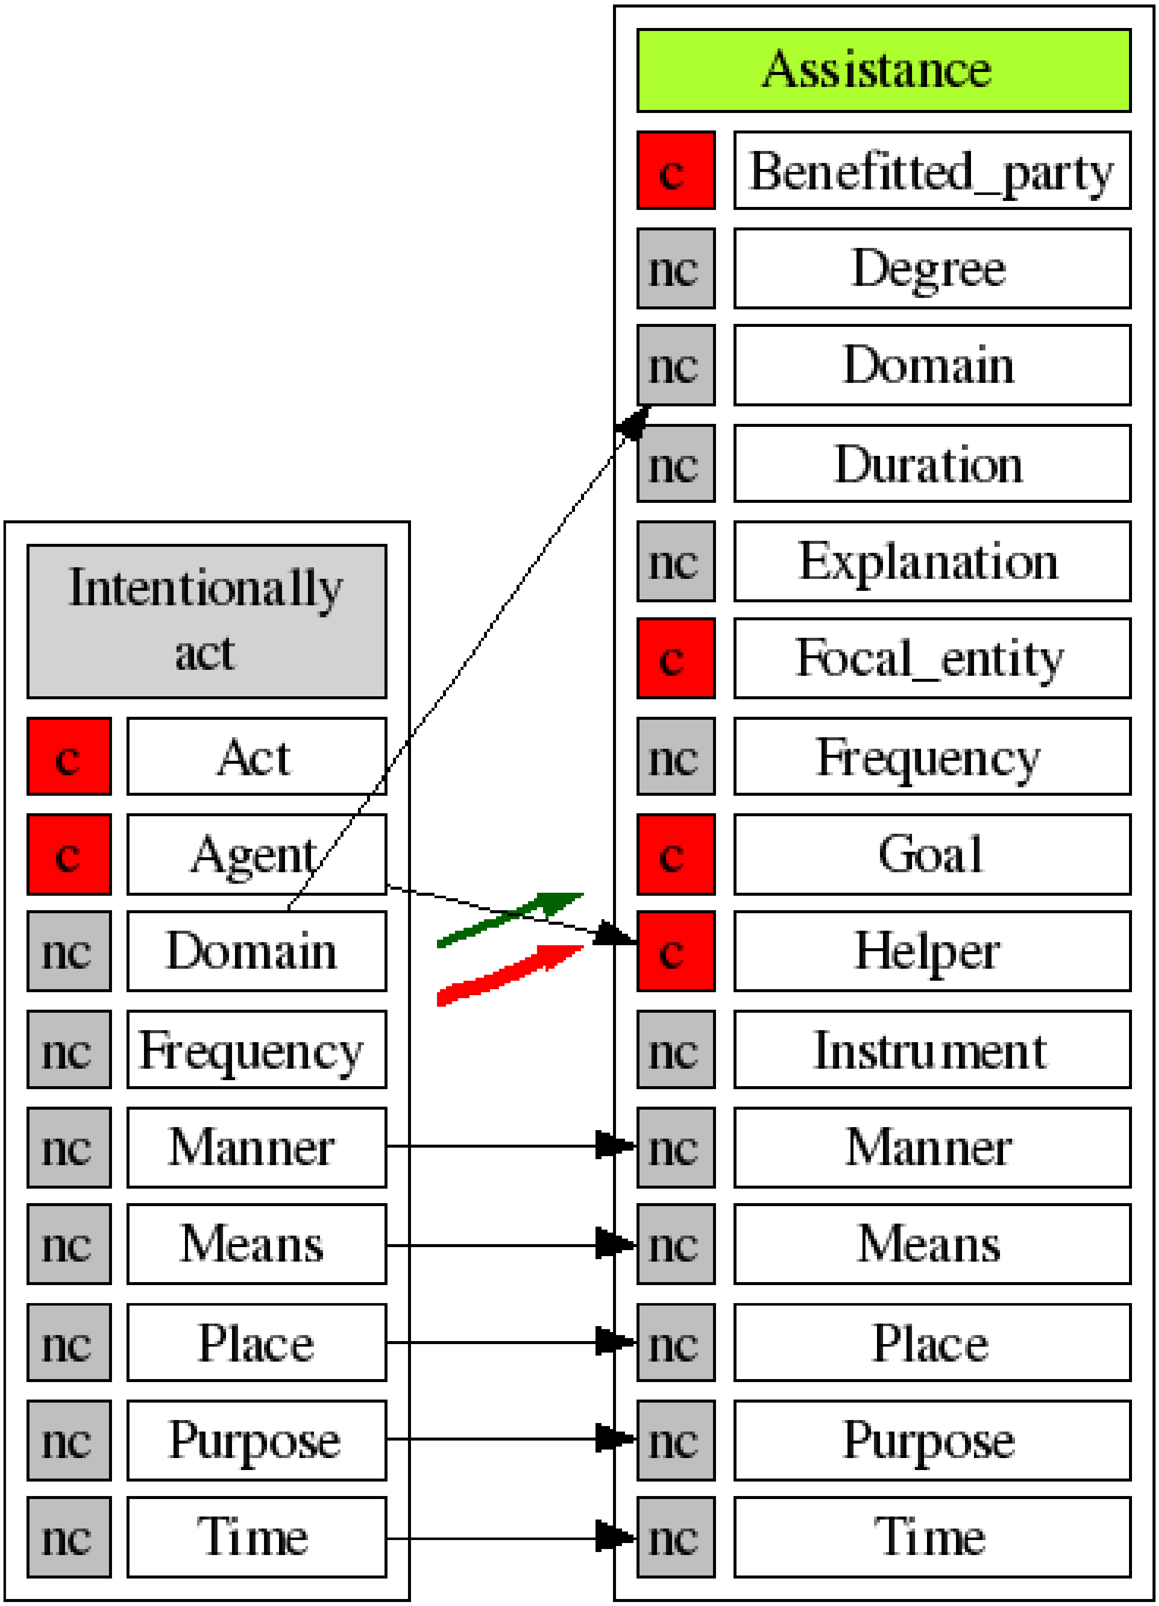

Figure 6.1: Assistance, with two relations to Intentional_act

Relation	Sub	Super
Inheritance	Child	Parent
Perspective_on	Perspectivized	Neutral
Subframe	Component	Complex
Precedes	Later	Earlier
Inchoative_of	Inchoative	State
Causative_of	Causative	Inchoative/State
Using	Child	Parent
See_also [1]	Referring Entry	Main Entry

Table 6.1: Types of Frame-frame Relations

6.1.1 Inheritance

This is the strongest relation between frames, corresponding to is-a in many ontologies.[2] With this relation, anything which is strictly true about the semantics of the Parent must correspond to an equally or more specific fact about the Child. This includes Frame Element membership of the frames (except for Extrathematic FEs), most Semantic Types, frame relations to other frames, relationships among the Frame Elements, and Semantic Types on the Frame Elements.

Properties of the Parent which are not strictly semantic in nature, such as not being evoked by lexical units (i.e. the Semantic Type Non-lexical_frame), being evoked by a particular set of Lexical Units, or having a See_also relation to another frame, are not inherited.[3]

Also, when there is a Core-set or an Excludes relation among Frame Elements of the Parent (see sections 3.2.2.1 and 3.2.2.3), these constitute disjunctive properties of the Parent. The Child frame may legitimately inherit only a subset of these disjunctive Frame Elements.

6.1.2 Perspective_on

This relation (new in Release 1.3) is a refinement of the more general Using relation (see 6.1.6). Perspective_on constrains related frames considerably more. The use of this relation indicates the presence of at least two different points-of-view that can be taken on the Neutral frame. For example, the Measure_scenario, in which an ENTITY's VALUE for some ATTRIBUTE is described, can be viewed either from the point-of-view of exact measurement (e.g. "Joey weighed 7 **pounds**.") or as a relative measure (e.g. "Joey was **heavy**."). The FEs in the two cases are quite different, so the words should not be included in the same frame (see 2), but they do make reference to the same scene. The Perspective_on relation allows us to refer directly to the scene (encoded by the Neutral Measure_scenario frame here) and connect the two. As in this case, the Neutral frame is normally Non-lexical and Non-perspectivalized. (See 6.2.2.)

A single Neutral frame generally has at least two Perspectivalized frames, but in some cases, words of the Neutral frame are consistent with multiple different points-of-view while the Perspectivalized frame is consistent with only one.[4] Whenever there is a state of affairs that is describable by a frame in a Perspective_on relation, all the other frames that are connected to it by the Perspective_on relation can also be used to describe the state of affairs.[5]

[2]In some few cases (enumerated in the release notes for data release 1.3), we have used the Inheritance relation in an apparently incorrect sense, where a Child frame has two types of realization, only one of which can straightforwardly be described as inheriting from the Parent. For example, the Manipulation frame has two possible points-of-view: one which focuses on the AGENT ("She **grabbed** the cigar.") and the other on the BODYPART_OF_AGENT ("Her hand **grabbed** the cigar."). Of these, only the AGENT-focused realization properly inherits from Intentionally_affect, as the BODYPART_OF_AGENT-focused realization has no AGENT, which is an obligatory FE of the Intentionally_affect frame. All the remaining errors in Inheritance are of the same kind: where the Child frame has an alternation of FEs related by some metonymy and one of these related FEs is not present in the Parent. Although not indicated in Release 1.3, in future releases these metonymies will be explicitly marked.

[3]These descriptions of a frame actually represent meta-information rather than true properties of the frame.

[4]Many of our frames, including many which are not explicitly involved in the Perspective_on relation, have more than one point of view inherently. In particular, any frames which have exclusion sets among their FEs have a separate point-of-view associated with each different choice of excluded FEs.

[5]Note that the fact that all of the frames in the set can describe a situation does not always mean that they are intersubstitutable.

Ex: These skateboards are easy to *sell.* ≠ ...easy to *buy.*

A word like *easy* is specifically sensitive to the point-of-view of its complement, picking out the actor in that scene as the EXPERIENCER. In a scene like the one involving *buy* and *sell*, there are multiple individuals that can be construed as the actor, and *buy* and *sell* lexically (and framally) profile these different potential actors. It is still the case that to whatever degree a buying event occurs, a selling event does also.

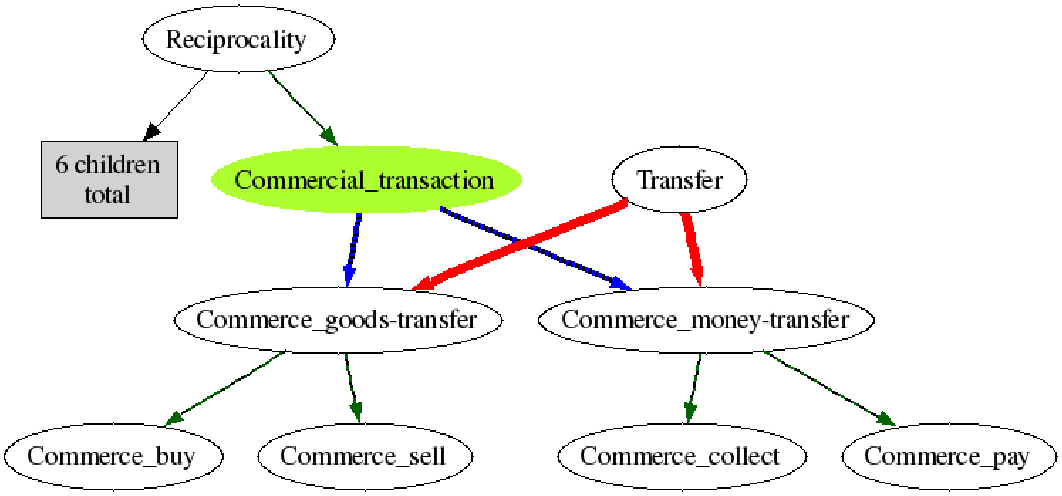

Figure 6.2: Commercial_Transaction

For example, the Commercial_transaction frame, diagrammed in Figure 6.2, specifies a complex schema involving an exchange of multiple Themes (the MONEY and GOODS) between the BUYER and SELLER, including also two subframes: Commerce_goods-transfer and Commerce_money-transfer. The Buying frame has a Perspective_on relationship with the Commerce_goods-transfer subframe in which the MONEY, GOODS, BUYER, and SELLER are identified.

Some other relations of this type include:

- Get_a_job and Hiring are perspectives of Begin_employment
- Rope_manipulation and Knot_creation are perspectives on Knot_creation_scenario

Further note that it is quite common for a frame to inherit from one frame and be a perspective on another. An example of this is again provided by the Commerce_buy frame, which inherits the Getting frame (not shown in the diagram 6.2) but is a perspective on the Commerce_goods_transfer frame. An act of buying is a sub-type of getting, which justifies the inheritance relation to Getting. Buying is an event that only occurs when the situation (the Commerce_goods_transfer) is also describable as selling. For that reason, it (and the Commerce_sell) is connected to the Commerce_goods_transfer frame via a Perspective_on relation.

6.1.3 SubFrames

Some frames are complex in that they refer to sequences of states and transitions, each of which can itself be separately described as a frame. The separate frames (called subframes) are related to the complex frames via the SubFrame relation. In such cases, frame elements of the complex frame may be identified (mapped) to the frame elements of the subparts, although not all frame elements of one need have any relation to the other. (In this respect, it contrasts with inheritance; see below.) Also, the ordering and other temporal relationships of the subframes can be specified using binary precedence relations. To illustrate, consider the complex Criminal_process frame, which is defined as below and whose frame relations are shown in Figure 6.3.

> A Suspect is arrested by an AUTHORITY on certain CHARGES, then is arraigned as a DEFENDANT. If at any time the DEFENDANT pleads guilty, then the DEFENDANT is sentenced, otherwise the DEFENDANT first goes to trial. If the FINDING after the trial is guilty, then the DEFENDANT is sentenced. In the end, the DEFENDANT is either released or is given a SENTENCE by a JUDGE at the sentencing.

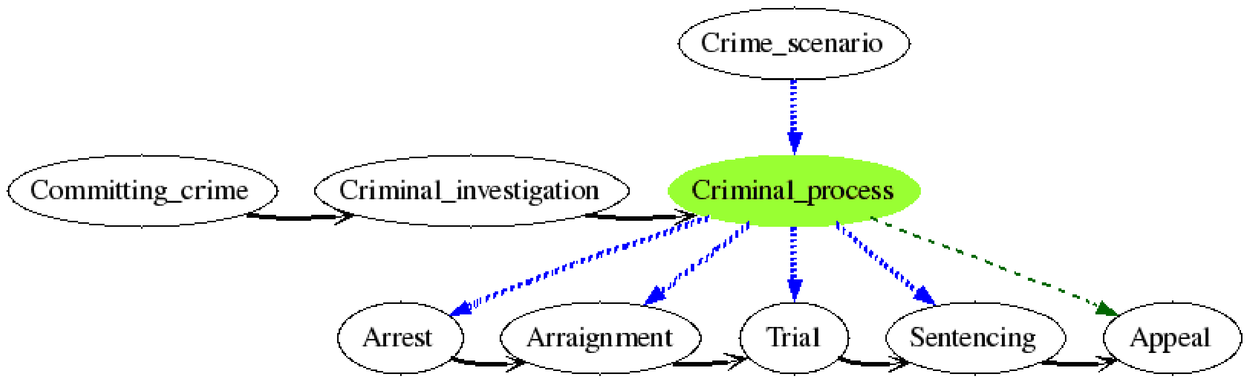

Figure 6.3: Subframes of the Criminal_process frame

For each step in the process, there is a separate frame in the database, including Arrest, Arraignment, Trial, Sentencing, and so on. Each of these frames is related to the Criminal_process frame via the SubFrame relation in the database. Moreover, subframes (of the same complex frame) are related to each other through their ordering. (See Narayanan et al. (2002) and Scheffczyk et al. (2006).) Note that the FEs of the subframes need not map to those of other subframes. So, in this same group of frames, the Arrest frame includes an ARRESTING_AUTHORITY which is not included in the Trial frame.

Notice that a given subframe may itself be a complex frame. For example, the Trial frame is a subframe of the Criminal_process frame, and has its own rich structure, some of which can be decomposed into simpler frames that are related to each other. A Trial is made up of court appearances, and involves opening arguments, presentation of evidence and testimony, and closing arguments. The system of subframe links is also quite complex. At present, the subframe relation is not indicated in every relevant case.

6.1.4 Precedes

This relation occurs only between two Component frames of a single Complex frame, i.e. as extra information associated with a set of Subframe relations. It specifies the sequence of states and events that are definitional for a certain state-of-affairs. Most Subframe relations will naturally have precedence relations, as can be seen in the foregoing diagram of the subframes of the Criminal_process frame (6.3), in which the Precedes relations are indicated by the black lateral arrows.

This is the only frame relation for which cycles are allowed. For example, in the subframes of the Sleep_wake_cycle frame (shown in Figure 6.4), Being_awake precedes Fall_asleep, which precedes Sleep, which precedes Waking_up or Getting_up, which in turn precedes the first frame, Being_awake.

6.1.5 Causative_of and Inchoative_of

We record the especially close and fairly systematic non-inheritance relationships between stative frames and the inchoative and causative frames which refer to them using the frame-to-frame relations **Causative_of** and **Inchoative_of**. Consider the following frames: Position_on_a_scale, Change_position_on_a_scale, and

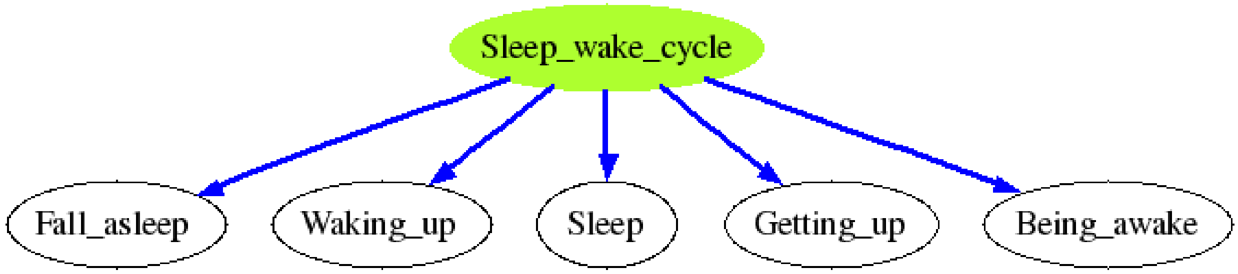

Figure 6.4: Subframes of the Sleep_wake_cycle frame

Cause_change_of_scalar_position. The following sentences illustrate the kind distinction captured in these frames:

(1) [Cause_change_of_scalar_position] (raise.v) Billie Blount **raised** taxes on farmers 18 times in 2002!

(2) [Change_position_on_a_scale] (rise.v) During the Elizabethan age, there was an increased emphasis on genealogy in the heralds' work as the gentry class **rose** in importance.

(3) [Position_on_a_scale] (high.a) Most fish from lakes is too **high** in mercury.

Frames which participate in these relations as Causatives should inherit from the Transitive_action frame, Inchoatives should inherit from Event, and States from State or Gradable_attribute.[6] Many of these inheritance relations have not yet been made.

6.1.6 Using

Often a particular frame makes reference in a very general kind of way to the structure of a more abstract, schematic frame. Since the creation of the more specific Perspective_on relation (see 6.1.2), the Using relation is used almost exclusively for cases in which a part of the scene evoked by the Child refers to the Parent frame. For example, Volubility uses the Communication frame, since Volubility describes a quantification of communication events.

It is possible for a frame to use more than one frame. An example of this situation is the Judgment_communication frame, shown in Figure 6.5 below. It uses both the Judgment frame and the Statement frame. The Judgment_communication frame does not inherit Judgment because it is not a simple subtype of a purely cognitive state. Judgment_communication does not inherit Statement either since it distributes the content of the Statement frame's MESSAGE frame element over two frame elements, EVALUEE and REASON.

6.1.7 See_also

In cases where there are groups of frames which are similar and should be carefully differentiated, compared, and contrasted, each of the frames in question has a SeeAlso relation with a representative member of the group. In the frame definition of the representative member, there will be a comparison which will contrast the frames to make clear the intended boundaries between them. For example, since the Scrutiny and Seeking frames are similar, there is a SeeAlso relation from Scrutiny to Seeking, and text in the Seeking frame that explains the difference.

This relationship does not imply any particular relationship between the Frame Elements or subframes of the frames involved.

[6]For many frames inheriting from Gradable_attribute, Inchoative or Causative frames corresponding to them should exist but have not yet been created. E.g. for the Age frame, which has LUs like *old.a*, there should be corresponding frames for the verb *age*, namely a Causative frame we might call *Cause_change_of_age and an Inchoative frame *Change_of_age.

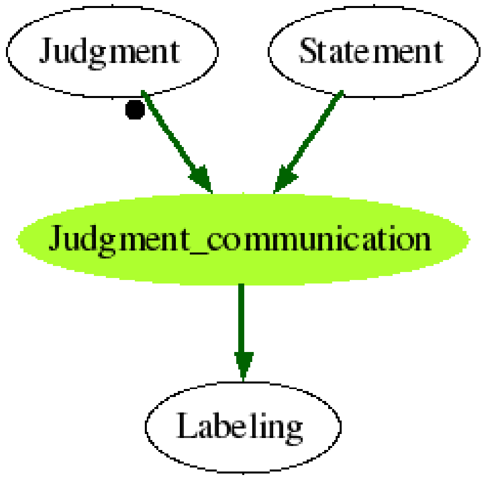

Figure 6.5: Judgment_communication's Using relations

6.2 Semantic Type

The general use of semantic types in the FrameNet project is to record information that is not representable in our frame and frame element hierarchies. In this section there is a detailed description of each major category of the semantic type hierarchy, which is broadly split by function. In specific, the functions we currently employ semantic types for are:

- Indicating the basic typing of fillers of frame elements, e.g. "Sentient" for the COGNIZER FE. These basic types are constrained by the frame hierarchy (see section 6.3.5.1), but not predictable from it since frame elements which are arbitrarily far away according to the frame hierarchy, such as the EXPERIENCER of Perception_body and the PERPETRATOR of the Piracy frame, are often marked as the same semantic type (in this case, Sentient). This kind of semantic type is designed primarily to aid frame parsing and automatic FE recognition.

- Useful, functional marking on frames, such as the type "Non-lexical" on frames which are present purely to participate in Inheritance, Subframe, or Using relations with other frames. This kind of property is actually a meta-description, not a fact about the semantics of the frame at all, and thus independent of the hierarchy. In fact, a frame which in English is Non-lexical might well have associated LUs in another language.

- Marking important dimensions of semantic variation among the lexical units in a frame that are not related to the kind of semantic combinatorial possibilities that we use for making frame distinctions (see Chapter 2). For instance, in the Judgment frame the difference between LUs such as *praise.v* and *criticize.v* in terms of the negative versus positive evaluation of the EVALUEE is marked with the semantic types Positive_judgment and Negative_judgment, respectively.

The most interesting function of semantic types for human users is the third one, recording important semantic differences between lexical units that recur within several frames.

For example, "Positive_judgement" and "Negative_judgement" semantic types, indicating the speaker's attitude toward a situation, can be applied to lexical units across a range of frames. (Note that the term *speaker* may either refer to a frame element such as SPEAKER or COGNIZER, or to the author of the utterance containing the lexical unit.) We capture the distinction by marking LUs like the aforementioned Judgment verb *praise*, the Experiencer_subj verb *like*, and the Frugality adjective *generous* as "Positive_judgement" and, by contrast, marking "Negative_judgement" on some other words from the same frames, e.g. *criticize* and *hate*, and *stingy*.

For more description and further examples, see the following subsections.

6.2.1 Ontological Types

These semantic types are used to classify the denotation of lexical units, frames, and frame elements. On *lexical units*, they classify the denotation of LUs, usually cross-cutting their classification by frames.[7] When, on the other hand, such a distinction corresponds to a frame distinction, the *frame* is labeled with the type, which then signifies that every LU of the frame could be labeled with an equally or more specific Type. For *FEs*, semantic types actually categorize the sort of *filler* that is expected in the FE, rather than classifying what kind of role the FE is.[8]

Most of these types correspond directly to synset nodes of WordNet, and can be mapped onto ontologies, e.g. Cyc. The types are related to each other via subtype relations, which are logically equivalent to Inheritance or is-a. However, the FrameNet semantic type hierarchy is not guaranteed (or even likely) to correspond exactly to that of WordNet, Cyc, or any other resource.

One example of a semantic type which has been used for frames, FEs, and LUs is *Container*. This corresponds most closely to the WordNet synset node *container* and the Cyc Container Collection node.[9] The Container frame, unsurprisingly, contains LUs which refer to containers like jar.n, box.n, etc., and thus both the frame and the LUs in it can be labeled with the semantic type *Container*. The SOURCE of the Emptying frame is also labeled with the ST *Container*. This means that head-nouns of this FE are likely to be the sorts of words that belong as LUs in the Container frame. Pronouns or other kinds of nouns may also occur as fillers of the FE, in which case they are construed as containers, as in the following:

(4) She **emptied** *the lid* back into the bottle.

In a neutral context, we wouldn't expect a lid to itself be a container. Not every lid is at all readily describable as a *Container*, e.g. the lid of a tin can. Regardless of whether a filler of this role is inherently a *Container*, its occurrence as a filler of the role means that it is being *used* as a *Container* in the current context, and has the properties necessary for that.[10]

6.2.2 Framal_type

These types, as mentioned above, can only be applied to frames and are **not** to be interpreted as applying to LUs of a frame or any connected frames, including Inheritance daughters (which normally should match any semantic types that are placed on their parent frames).

6.2.2.1 Non-lexical_frame

Such frames have no lexical units and are present purely to connect two (or more) frames semantically. One example is the Post_getting frame, connected to Getting via a Precedes relation and connected to Possession via an Inheritance relation. This allows us to succinctly encode the fact that the state following "getting X" (the Getting frame) is "having X" (the Possession frame).

6.2.2.2 Non-perspectivalized_frame

This semantic type is used for frames which have a great diversity of lexical units, all of which share a kind of scene as a background. Such frames do not have a consistent set of FEs for the targets, a consistent time assigned to the events or participants, or (most especially) a consistent point-of-view between targets. An example of this type of frame is the Performers_and_roles frame, which contains such diverse LUs as *co-star.v*, *feature.v*, and *as.prep*. Like the Biframal_LU types, this semantic type is intended as a time-saving measure. All such frames could be split up into smaller frames with a consistent perspective, but these frames would contain very few LUs. (See 6.2.3.3 on Biframal_LUs.)

[7] For example, the semantic type Body_of_water is applied to some LUs in the Biological_area frame, e.g. *bog.n* and to LUs in the Natural_features frame, e.g. *bay.n*.

[8] We have implemented the typing on roles themselves (such as whether they are complements or adjuncts) using the separate Coreness status feature (see Section 3.2.1).

[9] These are not quite equivalent to the FN definition. Cyc, for example defines Container Collection as referring to "tangible objects whose primary function is to contain something" (`http://opencyc1.cyc.com:3602/cgi-bin/cyccgi/cg?cb-start`) rather than the broader FrameNet definition which allows any entity which is construed as containing something.

[10] In some cases, a filler occurs in a particular role despite the fact that it has features which directly contradict the semantic type assigned to a role. For example, in "She filled her days with meaningless tasks", *day.n* is a non-physical entity and thus clashes with the type *Container* which inherits from *Physical_entity*. All such cases will need to be interpreted metaphorically or metonymically by resources beyond the scope of FrameNet.

6.2.3 Lexical_type

There are a number of labels applied to LUs or frames which do not refer directly to the kind of entity denoted by an LU or the semantics of a frame. These demand, rather, some more complicated interpretation of the LU or frame that they are attached to. How entities with such types are to be interpreted is different for each of the types that are described in this section (and in the following section on Framal types), and in some cases the label even indicates that an LU should be interpreted in a different frame from the one that it is attached to (see Guest_LU types below 6.2.3.3), so LUs with these types, if interpreted as subtypes of the semantics of the frame according to the normal FrameNet logic (see 6.3.4.1), can be very misleading. (This means that for some purposes, it would be more sensible to leave out items with these types attached than to include them without further comment.) A description of each of the types follows.

6.2.3.1 Transparent_nouns

LUs marked as **transparent nouns** have an unusual sort of semantics since, unlike most nouns, their primary function is to give some kind of description of another noun, as seen in the examples with *kind*, *number*, and *side* below:

(5) I **read** that kind of book in college.
(6) Monotremes **lay** a number of eggs per clutch.
(7) Put it **on** the left side of the shelf.

Specifically, transparent nouns describe either 1.) information which is not typically inherent in the nouns they describe, such as quantity, grouping, and shape, or 2.) information which can be left implicit in the interpretation of nouns when they denote an instance of an entity, such as a.) instancing itself (e.g. “this bat” = “an instance of a bat”), b.) subtyping (e.g. “this bat” = “a type of bat”, since an instance is a maximally specific type, and thus automatically a subtype of its class, or c.) picking out subparts (e.g. “with a knife” = “with the edge/point of a knife”), as an **active zone** subpart of an entity is often meant even when the whole entity is mentioned.

The descriptive aspects of meaning profiled by transparent nouns are usually of little consequence for determining acceptability as fillers of predicate roles; rather more important is the semantics of the entity being described. This can be seen in the examples below, in which the above sentences are summarized/paraphrased by omitting the transparent noun, leaving only the described entity:

(8) I **read** those books in college.
(9) Monotremes **lay** many eggs per clutch.
(10) Put it **on** the shelf, towards the left.

We judge that recognition of these LUs is vital for correct cataloging of FE fillers, summarization, paraphrase, etc. since they violate the normal rule that the syntactic head of a phrase is the semantic head. For most purposes, transparent nouns can simply be omitted from the analysis of a sentence.

However, note that transparent nouns are not always transparent. In the right semantic contexts, they are so to speak ”opaque” and dominate the frame of their dependent (as most targets do); in such cases the meaning of the so-called transparent noun itself rather than that of its dependent is selected by a predicator. In particular, this is the case with governing predicates that evoke an open proposition or question, e.g. determine, find out, measure, etc.

(11) Archaeologists have been determining the **number/shape/part/type** of tools used by the Maya.

Here none of the transparent nouns can be used with quite the same meaning as the that of the described entity alone:

(12) Archaeologists have been determining the tools used by the Maya.

In addition, transparent nouns convey important, if optional, information like quantity and configuration. For tasks concerned with these dimensions of meaning, transparent nouns should be treated like any other target.

6.2.3.2 Bound_LU

Some LUs of a frame cannot be used unless they co-occur with the right other words. The limiting case, where a precise group of words must all co-occur in order to evoke a frame, is handled in FrameNet by the use of multi-word lexical units (3.1). But there are many cases where an LU may evoke the frame when occurring with any of a number of different words, or even when occurring with a class of words. Such LUs are to be marked with the *Bound_LU* semantic type.

There are two subtypes of Bound LU that we recognize in FrameNet: Supports and Bound_dependents.

Support This subtype of bound lexical unit is applied to LUs that only evoke a frame when they are used as Supports of certain dependents, as exemplified by *give.v* in sentences like the following:

(13) Receiving the notification so late almost **gave** me a heart attack.

In this sentence, *give* is used approximately synonymously with *cause*:

(14) Receiving the notification so late almost **caused** (me to have) a heart attack.

This use of *give*, then, should go in the Causation frame, but it is readily apparent that this meaning is not possible in many (if not most) contexts:

(15) ??Releasing these old files nearly **gave** a disaster.
(16) Releasing these old files nearly **caused** a disaster.
(17) *What **gave** these events?
(18) What **caused** these events?

As implied above, *give* occurs in this meaning only when it is to be annotated as a Support Verb (see 3.2.7.1). (Note that there are other support uses of *give* which are not equivalent to *cause*, e.g. *give a laugh, give a speech.*)

In principle, though not yet in fact, all supports with any semantics not included in the target noun, adjective, or preposition that they are associated with, such as causative supports, inchoative supports, and point-of-view-shift supports, should be annotated separately in a causative, inchoative, or point-of-view specifying frame.

Bound_dependent In addition to fixed expressions and support constructions, there is a small class of other LUs which only occur as dependents of a limited set of governors. This semantic type has not been included on LUs in the current data release, but it is intended to cover semi-productive LUs such as *attention.n* in the Posture frame (in phrases like *(stand) at/to attention*), *bind.n* in the Predicament frame (in phrases like *in/out of a bind*), and idiosyncratic degree modifiers from the yet-to-be-made *Degree frame (for LUs like *very.adv*), e.g. *dirt* in *dirt poor*, *crystal* in *crystal clear*, *sopping* in *sopping wet*, etc. These LUs can be added to appropriate frames and given a Bound_dependent type, from which the user should infer that they are only usable when they have particular governors. The appropriate governors must be gleaned or generalized from the usage in the annotation data.

6.2.3.3 Biframal_LU

LUs of a frame normally have a semantics which is a subtype of the semantics of the frame. However, when LUs are marked with a Biframal_LU type, they denote something related, but not equal to the semantics of the frame. In all cases where an LU is marked as biframal, we could have made a separate frame to more narrowly characterize the meaning of the LU. These types have been defined as a time and resource-saving measure, to avoid having to create these separate frames for tiny groups of LUs; for example, *(un)intelligible.a*, despite the close relationship with the Grasp frame, would otherwise require a separate frame.

Another way of looking at this is that the normal relation of the semantics of the LU to the semantics of the frame is Inheritance, while the relation of a biframal LU's semantics to the frame's semantics is generally Using. Except for the Guest_LU type, each of the following subtypes define what frame the biframal LU inherits from. (See below for Guest_LU.)

We tend to avoid using these semantic types, since they complicate the interpretation of our data, and in many cases there is a problematic mismatch in the FEs allowed for the biframal LU and the other LUs of the frame.[11]

[11] E.g., *gunman.n* is included as an Agentive_noun LU in the Bearing_arms frame, but this frame does not include all of the

Guest_LU This type is applied to LUs which are only tangentially used in the host frame, and whose interpretation is still largely dependent on their membership in some other frame, as demonstrated in Example 6.2.3.3.

(19) Statement: " You'll never catch up," she **grinned**.
(20) Gesture: She **nodded** him through the door .
(21) Caused_motion: He **sneezed** the handkerchief off the table.

Thus, for example, *grin.v* can be included in the Statement frame, but any full understanding of this sentence requires an understanding of the Making_faces frame.

Agentive_noun This type is used to mark LUs which denote the AGENT[12] of the frame in which they occur, as in *murderer.n* in the Killing frame. They can be thought of as virtually inheriting from People and using the frame they are listed in; the PERSON FE of the People frame is bound to the same FE that is bound to the AGENT FE of the host frame.

Participating_entity This type marks LUs which denote the PATIENT[13] of the frame in which they occur, e.g. *possession_((entity)).n* in the Possession frame, *victim.n* in the Crime frame, and *knowledge.n* in the Awareness frame. These LUs can be thought of as inheriting from the Entity frame, and as using the host frame, with the PATIENT (or equivalent) FE of the host frame bound to the same FE as is bound to the ENTITY FE of the Entity frame.

Tendency_grading_LU [14]

This type marks LUs which denote the propensity (of something) to be a PATIENT (or similar FE, see 6.2.3.3) in the host frame; e.g. *verifiable* in the Verification frame is marked as a tendency-grading LU since it concerns something's tendency/ability to be verified. The semantics of a tendency-grading LU can be understood as a virtual frame inheriting from Inclination and using the host frame.

The most common subtype of these LUs are abilitative passives, usually constructed morphologically from a verb (X) + *(a)ble*, which denote the propensity of an entity to be Xed. Normal, productively formed examples include *findable* (Locating), *usable* (Using), *likable* (Experiencer_subj), and (with negative prefix) *unfixable* (Resolve_problem). There are also semantically indistinguishable examples which are not productively formed from verbs, e.g. *intelligible* (Grasp). These LUs are listed in frames that inherit from the causal Transitive_action frame.

In addition, there are some LUs which do not directly make reference to a causal frame, but rather refer to an inchoative frame, e.g. *fragile.a* (Fragment), and *mortal.a* (Death).

6.2.3.4 Affect_describing: Positive- and Negative_judgement

Many LUs describe the positive or negative opinion of an implicit or explicit JUDGE, without actually necessarily being of similar type otherwise, e.g. *like.v* in the Experiencer_subj frame, and *stingy.a* in the Stinginess frame. This type labels such LUs so that they can be compared despite their dissimilarities. In addition to the positive or negative dimension, there are two main types, distinguished by *who* is identified as having a high or low regard for something: 1.) Emotion-related targets, where the EXPERIENCER (or descendant of EXPERIENCER) is explicitly identified as the opinion-holder and the opinion concerns an overt STIMULUS or EVALUEE, e.g. *approve.v*, and 2.) other targets where the user of the word in the discourse is the opinion-holder and the opinion concerns the target predication as a whole, e.g. *thrifty.a*. All LUs in frames which inherit from or use the Emotions frame, especially the Judgment family of frames, fit in type

FEs appropriate for annotating people, like ETHNICITY, so that "Somali" in "Somali gunmen" is not annotatable; this problem is even more obvious with a word like *breakable.a* in the Render_non-functional frame, since it requires a DEGREE FE which is not available for other members of the frame. In such cases, we have the unenviable decision between adding a frame element which is not useful for most targets of the frame, and ignoring an important FE; in practice, we follow the latter approach and do not include FEs which are only possible for biframal LUs.

[12] For the purposes of this section, AGENT is taken to refer to any frame element which is connected to an AGENT FE via inheritance, even if this daughter is not called an AGENT; e.g. the COOK Agentive FE in the Apply_heat frame. Note that if an LU simultaneously denotes an FE inheriting from AGENT and any other FE (e.g. *braggart.n* in the Bragging frame, which also denotes the EVALUEE of Judgment) then Agentive_noun is used.

[13] For the purposes of this section, PATIENT is taken to refer to any frame element which is connected to a PATIENT FE via inheritance, even if this daughter is not called a PATIENT.

[14] This type is renamed and broadened from the Abilitative_passive type in previous releases.

1, with the COGNIZER, SPEAKER, or EXPERIENCER as the opinion-holder, whereas all other targets express the opinion of the individual using the word.[15]

In words of the second type, where it is the language-user whose view is expressed, the full semantics of the LU could be thought of in terms of a virtual frame which inherits from the host frame and uses the Regard frame, with the COGNIZER specified as the language-user, the EVALUEE identified as the whole predicate, and the JUDGMENT identified as positive for Positive_judgement and negative for Negative_judgement. Thus the target *bony.a* of the Body_description_holistic frame in 6.2.3.4 below indicates not only that someone was thin (which is at least partly specified by the meaning of the frame) but also the person declaring this state considers this undesirable.[16]

(22) Ignacio said she was looking **bony**.

6.2.3.5 LU_with_FE_specified

These types are employed on LUs to indicate that the LU inherently specifies some information about an FE. Although the types of information that LUs can specify about FEs is considerably broader (including most saliently filler-types for an FE which are particular to an LU, as in *tie.v* in the Attaching frame), we have so far used only two basic kinds, which specify DEGREE and sensory modality respectively.

Degree_specified_LU (Negative and End_of_scale) These types are used for LUs whose framal semantics provides a DEGREE FE. All such LUs inherently describe a particular ATTRIBUTE of an ENTITY as deviating in a particular direction from the norm. (This norm is usually the expected value of the ATTRIBUTE for the type of ENTITY described, but may be a norm for some more abstract supertype of the ENTITY. See the Gradable_attributes frame.) LUs should pre-specify the direction and difference from the norm; Negative indicates the direction and End_of_scale indicates maximal difference.

LUs which have the type Negative specify that the direction of deviation is in the negative direction, i.e. either less in quantity or desirability (or at any rate in the opposite direction from the unmarked description given in the frame definition). These LUs thus form antonyms to unmarked LUs in the same frames. Some select cases are presented in Table 6.2. Note that many LUs that should receive this type have not been marked in the FrameNet data.

Frame	Negative LU	Positive LU
Ambient_temperature	cold.a	hot.a
Experiencer_subject	hate.v	love.v
Mental_property	stupid.a	smart.a

Table 6.2: Antonymy in select frames

End_of_scale marks LUs that characterize an ENTITY as having a maximal value for an ATTRIBUTE, e.g. *fabulous* in the Desirability frame. Such LUs are not usable with normal DEGREE modifiers like *very*.

Sensory_related_LU LUs marked with these types inherently reference some type of sensory experience, either directly denoting such an experience, e.g. *see.v*, *look.v*, or other LUs in the Perception_experience or Perception_active frames, or they use such a frame, and thus imply the occurrence of a perceptual event, e.g. *tasty*. In any case, the important function of the subtypes of this type (viz. Visual_modality, Tactile_modality, etc.) is to identify the subtype of sensory experience which is inherent to an LU. Some examples of the use of the Visual_modality type:

(23) *see.v* in Perception_experience
(24) *glance.v* in Perception_active
(25) *shiny.a* in Location_of_light
(26) *ugly.a* in Aesthetics

[15] In some few cases, such as *fob/pawn off.v*, it seems that the negative judgement (of the language-user on the THEME in this case) is also ascribed to another participant in the frame (in this case the DONOR); this fact is not currently representable in the FrameNet data.

[16] In this example, since the word is embedded as an FE in the context of a speech verb, a full analysis might conclude that the explicit SPEAKER is also the opinion-holder, but this type of inference lies strictly in the domain of mental-spaces analysis rather and outside frame analysis proper.

6.3 How to use relations

This section outlines some of the major uses for the relations and semantic types. Section 6.3.1 is primarily devoted to describing the utility of the relations for humans using FrameNet as a reference, Section 6.3.2 is of use to both humans and automatic programs, whereas the other sections are of more interest to developers intending to use FrameNet for computational purposes. Virtually any computational use of the FrameNet relation information relies on and presupposes semantic parsing of texts–a process not discussed here, but well covered in several publications (e.g., Erk and Padó (2006) and Litkowski (2004)).

6.3.1 Looking up related words

One obvious benefit of defining relations between frames is that it enables the look-up of words that are related. For example, we can look at the word *murder.v* which occurs in the Killing frame. Obviously, there's some level of similarity with the other members of the frame, since they all involve something causing an entity to die. This sort of relationship might be captured in a good thesaurus under an entry for *kill.v*. However, if we follow the Causative_of relation from the Killing frame to the Death frame, we find a whole new group of words, also about death, e.g. *die* and *pass away*, but in this case focusing on the change of state of the PROTAGONIST. This kind of relation is neither synonymy nor antonymy, and therefore not to be found in a thesaurus.

6.3.2 Comparing to other systems of semantic annotation

In Fillmore's earlier work Fillmore (1968, 1977), a case was made for the universality of certain types of semantic roles, a concept which was further developed and is now enshrined as the theta role system of many syntactic/semantic formalisms. These roles include such labels as Agent, Instrumental, and Objective (roughly corresponding to Agent, Instrument, and Patient in other formulations). However, as the description of the semantics of Lexical Units has progressed, it has become apparent that the theta-role and original case-role account covers only a subset of the full set of roles.

We now take it that theta roles should be mapped to FEs in high-level, abstract frames like Transitive_action, which has FEs like AGENT and PATIENT. The relevance or irrelevance that these labels have for the roles of more specific predicates like *break.v* (in the Cause_to_fragment frame) or *resemble.v* (in the Similarity frame) is explained explicitly by the inheritance or non-inheritance of the AGENT and PATIENT frame elements in the relevant frames.

There are inherent problems to reducing our role-set of Frame Elements to the considerably smaller (and inarguably more computationally tractable) set of theta roles often used. One of these is deciding on the initial set of theta roles – a well-known problem in the theta-role literature. The Frame Elements that we define, however, are more immediately verifiable. Presumably, then, any theta-role system proposed to cover all predicates should allow us to specify, in a simple way, which FEs should be mapped to which theta-roles. Of the theta-role systems known to us, none allows any simple mapping to high-level FEs (and thence to the FEs that inherit from them) without covering some FEs multiple times and/or leaving some FEs uncovered, unless there is an unsatisfactory catch-all theta-role.

6.3.3 Paraphrase and translation

In many ways, paraphrasing is at the core of what we intend FrameNet to facilitate. A properly powerful ability to paraphrase enables many of the other goals of semantic NLP, including Question Answering, Summarization, and Translation. Question Answering can be thought of as looking in a corpus to find a paraphrase, but with real information filled in for the questioned FE. Summarization is equivalent to paraphrase of a text, but with the strategic omission of information from FEs and targets. Translation is paraphrasing with the limitation that all the resulting paraphrase must be in the target language.[17]

One of the basic insights behind FrameNet is that grouping words according to the scenes that they evoke, regardless of whether they are synonyms, antonyms, or some other relation to each other, groups words that are useful for paraphrasing. In particular, since FrameNet lists words together despite part-of-speech differences (unlike WordNet), paraphrases involving an interchange of noun, verb, adjective, or preposition are (in principle) discoverable with the FrameNet data. (See Sect. 6.3.3.1.)

[17] This requires FrameNet-style data for both source and target language; this currently limits such efforts to English, German, Japanese, and Spanish.

6.3.3.1 Differences in LUs: Semantic Types and WN

Despite the fact that FrameNet is built to facilitate paraphrase, more specific relationships between the words usually need to be ascertained to recognize or generate actual instances of paraphrases. What kind of paraphrase results is dependent on whether a proposed paraphrasing target is an antonym, exact synonym, subtype, or supertype of the initial target to be paraphrased.[18] We have some of this more specific information (see the Semantic type "Negative", 6.2.3.5), but for the most part, we have envisioned the information that we record as complementary to WordNet.

As WordNet has a reasonable coverage of synonymy, antonymy, and subtyping ("is-a") relations between words, we have tried to avoid reinventing the wheel by duplicating this information in FrameNet. However, for tasks that require information from both sources, integrating the information would depend on mapping FrameNet LUs to words in WordNet synsets in some way. This has not been consistently accomplished anywhere, although work has been done by several researchers.[19]

Once a proposed list of paraphrases for a starting LU is established (however buggy), then targets can be paraphrased with other words which are synonyms, supertypes, subtypes, or (with the addition of negation) antonyms in the same frame; this excludes, then, targets which are "sisters" or "cousins", i.e. which are subtypes of some supertype of the original target.

Working without a WN-FN mapping, antonyms can at least provisionally be established if some members of the frame have the Semantic type Negative. Such targets are antonyms of unmarked targets in the same frame. In addition, targets with FE incorporation should be considered subtypes of targets in the frame without incorporated FEs. In the absence of any other information, all other targets could be considered synonyms, although this will be erroneous in many cases.

Degree differences, especially antonymy Paraphrasing an initial target with a target that is an antonym requires adding in negation, as can be seen in 27 and 28 with antonymic targets from the Possession frame:

(27) Initial sentence: She may **lack** the money for a more nutritious meal.

(28) Paraphrase: She may *not* **have** the money for a more nutritious meal.

Supertypes and synonyms When paraphrasing a more specific initial target with a more general target from the frame, no adjustment is generally necessary, just as with a synonym, as seen in examples 29 and 30 from the Self_motion frame:

(29) Initial sentence: I enjoyed watching the giraffes **saunter** by.

(30) Paraphrase: I enjoyed watching the giraffes **walk** by.

Supertype paraphrase works across frames as well; compare 31 from the Motion frame, which Self_motion inherits from:

(31) Paraphrase: I enjoyed watching the giraffes **go** by.

Interchanging part-of-speech of targets The simplest type of interchange of part-of-speech of targets does not require any change in the embedding context. This is most often the case when a target combines with a support to "simulate" a different part-of-speech. For example, a noun + support verb can often paraphrase a verb target (example from the Bragging frame):

[18] Some pairs of targets within a frame will not be fit paraphrases at all. This occurs when the targets have differing semantic relations (synonymy, antonymy, hyponymy, etc.) to a more generic concept within the frame; for example, the Subject_stimulus frame has such relative incomparables as *funny.a*, *shocking.a*, *sad.a*, *encouraging*, etc. These have no sensible paraphrase relationship with each other, except for the fact that they are all subtypes of emotional description focusing on the EXPERIENCER. Something similar can be said for frames like Biological_area with words like *forest.n* and *bog.n*.

[19] Martha Palmer and Andy Dolbey are working on the most thoroughgoing effort to make a broad alignment of WordNet, VerbNet, and FrameNet (Martha Palmer, personal communication), but this project has not yet produced publications.

Mapping in the other direction, Aljoscha Burchardt and others have created the "WN detour to FrameNet", which connects WN synsets to FN frames (`http://www.coli.uni-saarland.de/albu/papers/gnws05_burchardt_erk_frank-final.pdf` and `http://www.coli.uni-saarland.de/albu/cgi-bin/FN-Detour.cgi`).

Several other efforts have used a knowledge-engineering approach to put FN data together with WN (Shi and Mihalcea (2005)) or VerbNet and PropBank (Giuglea and Moschitti (2004)).

(32) Initial sentence: Eustace **boasted** uncomfortably.
(33) Paraphrase: Eustace *made* an uncomfortable **boast**.

In 32 and 33, *boast.v* and Support Verb + *boast.n* equivalently fill the finite-verb slot of the sentence.

They may also equivalently fill exactly the same valence pattern of a governing verb, as seen in 34 and 35. In this case, both noun and verb fill the VP-to realization of the GOAL FE of the verb *try.v* in the Attempt frame:

(34) Initial sentence: She *tried* to **boast** about her years of teaching experience.
(35) Paraphrase: She *tried* to make some **boast** about her years of teaching experience.

A noun + support preposition can paraphrase an adjective or adverb target:

(36) Initial sentence: He sat *in* some **discomfiture** as she explained.
(37) Paraphrase: He sat somewhat **discomfited** as she explained.

Other kinds of interchange of noun, adjective, verb, or preposition require changing the Phrase Type of the phrase that they occur in. This is only possible when the governing context has two different valence patterns with different phrase types for the FE filled by the target (see 6.3.3.2):

(38) Initial sentence: We all **want** to **succeed**.
(39) Paraphrase: We all **want success**.

Here, the paraphrasing of *succeed.v* with *success.n* is only possible because the EVENT FE of the target *want.v* in the Desiring frame can be realized as either an infinitival complement or an NP object.

6.3.3.2 Paraphrasing FEs with different PTs

One strategy of paraphrase is to exchange the valence pattern for the realization of FEs for a target; a typical case is the Dative Alternation:

(40) Initial sentence: **Hand** me that spatula.
(41) Paraphrase: **Hand** that spatula to me.

In this example (from the Giving frame), the RECIPIENT FE is filled either with an NP or a PP(to). Since the alternation of NP and PP merely involves the presence or absence of the marker (i.e. *to*), it is easily possible to paraphrase these mechanically.

In principle, for most possible pairings of PTs, there are cases when a particular FE will vary between the two PTs across a frame. All such cases provide potential paraphrases across phrase types.

For example, the Phrase Type "PP-ing(about)" in the sentence "She thought about returning" can be paraphrased in the sentence "She considered returning" with the PT "VPing", or in the sentence "She considered a return" with the PT "NP", or in "Her contemplation of a return" as "PP(of)", etc. These interchanges can be selected from the phrase types listed for a particular FE realization in the valence tables of the lexical entries. Thus, this type of paraphrase generation is relatively independent of the target choice (discussed in 6.3.3.1), limited only after the fact by what phrase types are available for the chosen target.

Many of the interchanges cannot be achieved by simply varying a marker, e.g. changing an NP to a VP or vice-versa. This more complicated type of interchange could only be accomplished by paraphrasing the targets inside the phrase, e.g. by paraphrasing the head noun of the NP with a verb from the same frame to make a VP.

The straightforward interchanges of PT can be produced mechanically by adding, subtracting, or changing a marking word (like a preposition) or changing the morphology of the head of the phrase (e.g. from present participle to infinitive). The paraphrasability of the unparenthesized PTs below in Figure 6.6 is merely a matter of changing the morphological marking on the head and/or changing the syntactic markers (such as prepositions, quotation marks, or the infinitive-marker to). For many other patterns (indicated in parenthesis below), there are very common constructions (such as copularization of adjectives) which allow an interchange of PTs.

- NP Poss PP (N) (A)
- VPing VPto VPbrst VPed VPfin PPing (Swhether Sinterrog PPinterrog)
- Sinterrog Swhether PPinterrog QUO
- Sing Sto Sforto Sbrst Sfin QUO
- AJP (AVP Sub) (Srel VPtorel VPed)
- QUANT (Srel)
- Sabs (Sub)

Figure 6.6: Simple phrase-type valence paraphrases

6.3.3.3 Paraphrasing by omission

Many FEs are optionally syntactically expressed, and many (less informative) paraphrases which omit them can be easily generated.

All peripheral FEs and non-target incorporated FEs are freely omissible, thus (omitting the Time FE of the Giving frame):

(42) Initial sentence: **Hand** me that spatula now.
(43) Paraphrase: **Hand** me that spatula.

In addition, FEs that ever occur with an LU as INI should be considered omissible.[20]

FEs that occur with an LU as DNI are omissible under essentially the same conditions as would license an anaphoric pronoun, deictic adverb, or the like for the same FE.

FEs that occur with CNI omission are omissible only in certain constructional contexts. The FrameNet data currently do very little to explicate what these contexts are, although they are generally well known syntactic phenomena.[21]

6.3.4 Inferencing

We intend that the Frame Relations, Frame Element mappings, and Lexical Units provide a significant amount of information for doing inferencing on natural language. We do not provide every kind of information that an inferencing engine might require (especially any kind of quantitative information), but the frame-to-frame relations should provide a good basis for further specifications (using X-Schemas, predicate logic, etc.) which are useful to existing systems, especially in combination with information from WordNet. To this end, an OWL version of these relations is included in the data release.

This section describes some of the kinds of inferences that are intended by the definitions we give to the various relations.[22]

The most basic summarization of the logic of FrameNet is that Frames describe classes of situations, the semantics of LUs are subclasses of the Frames, and non-Extrathematic FEs are classes that are arguments of the Frame classes.[23] An annotation set for a sentence generally describes an instance of the subclass associated with an LU as well as instances of each of its associated FE classes.

Inheritance has a special place among the frame relations. It is to be interpreted as a subtype or "is-a" relation between frame classes. This constrains the Child frame to have all of the properties of the Parent frame.

---

[20] In many cases, we have annotated the omitted FEs in generic sentences with INI. Such omissions would be more properly labeled CNI, on which see below.

[21] Correctly using CNI information for paraphrase would require specific methods for handling the constructions commonly referenced in the FrameNet data, i.e. Passives, Imperatives, Instructional imperatives, and Generics (including the common omission of external arguments with infinitives and participles).

[22] Note that the See_also relation has no formal semantic definition whatsoever, and thus is not intended for machine processing. It will not be mentioned further in this section.

[23] As mentioned in section 6.3.4.1, Extra-thematic FEs effectively evoke their own frame, separate from the frame of a target LU. The Extra-thematic FE itself is a Frame Element of this Extra-thematic frame. Other frame elements of the Extra-thematic frame should correspond to some subset of the Frame Elements of the target LU's frame, but exactly which subset of FEs is not identified in the current data release.

Although it is difficult to generalize across all of the relations, it is usually true that an instance of a Sub frame implies the existence of an instance of a Super frame. In the case of a Theft, which inherits from Committing_crime, an instance of Theft is obviously an instance of Committing_crime. For the Precedes relation, however, this is not exactly true .[24] For the Subframe relation, the existence of a Sub frame instance does imply the relevance of a Complex frame instance, but the reverse is not true. So, for example, if there is an instance of the Sentencing frame, then we can instantiate an instance of the Criminal_process frame. We can then also instantiate all of the subframes of Criminal_process which must precede Sentencing, including Trial, Arraignment, etc., but we cannot instantiate the Appeal subframe of Criminal_process which follows Sentencing.

6.3.4.1 Logical relations of FEs and LUs to Frames

The relations of FEs and LUs to frames are intertwined with our Frame relations in various ways. Some clarification and formalization of these concepts is necessary to simplify their logical description:

- The term "Frame Element" has two meanings: the relation itself, and the filler of the relation. When we describe the Coreness status of an FE (see Sect. 3.2.1), we are describing the relation; when we describe the Ontological type on an FE (see Sect. 6.2.1) we mean the type of the filler. Fillers are pronouns, proper names, or (more usually) common nouns that evoke entity or event frames. Entity reference, named entities, and anaphora are all outside the scope of the FrameNet project, but when FEs are filled by frame-evoking words, an interpretation engine should iteratively analyze these words in the same way as any other frame-evoking element.

- Extra-thematic FEs have a considerably different interpretation from all other FE types. Normal FEs (barring Coresets (see Sect. 3.2.2.1) or Excludes relations (see Sect. 3.2.2.3)) must always be logically present for the frame to make sense. Extra-thematic FEs, however, independently evoke a different frame from from the one they are listed in. The Extra-thematic FE itself fills one of the FEs of this frame, and the other FEs are filled by various frame elements of the original target word according to heuristics which must be separately specified for each Extra-thematic FE.

  For example, in 44, evoking the Ride_vehicles frame, the COTHEME FE evokes an additional instance of Motion whose THEME FE is filled by the COTHEME, and whose PATH and other FEs are co-identified with the instance of Ride_vehicles. Thus both "I" and "her" are described as moving to school in this example.

  (44) I **rode** to school [with her COTHEME] all the time.

- Except for the rare LUs with the Semantic Type Biframal_LU (see 6.2.3.3) or LUs in frames which are marked Non-perspectivalized (see 6.2.2.2), all LUs of a frame have a semantics which is best described as a subtype of the semantics of the frame. With these caveats, one may thus consider LU membership in a frame to be an identical relationship to Inheritance from a frame.

6.3.4.2 Inheritance

Not surprisingly, Inheritance, as our most formally defined relation, is an exact match for an ontologically defined relation, namely subtyping or "is-a", as mentioned above. Because these relations are so well-defined formally, there is very little else that needs to be said about Inheritance here.

6.3.5 Propagation of Information

6.3.5.1 Propagation of types on FEs

As discussed above, the Inheritance relation implies the correspondence of the Parent frame and its FEs to the Child frame and its FEs. This straightforwardly allows us to propagate the ontological semantic types (see Sect 6.2.1) on the Parent frame and its FEs downward to the Child frame and its FEs. However, what may not yet be apparent is the fact that the FE bindings in all frame-to-frame relations are equivalent to the Inheritance FE bindings, *regardless of the type of relation.* This means that normal semantic types on FEs

[24] The full interpretation of which frames in a group of subframes are implied to have happened requires the use of X-schema logic or an equivalent system for querying reachability and preceding states in relationally-defined state/event systems. See Narayanan (1999).

can be propagated down the hierarchy along every type of relation. Types are propagated explicitly in the current OWL representation of the FrameNet data.

Because we cannot anticipate all of the Semantic Types that will be useful for tagging FEs, it will certainly also be desirable to categorize the fillers of our FEs using WordNet (or a similar resource). Just as with the pre-marked semantic types applied by the FrameNet team, whatever information is gathered on fillers of an FE in one frame can be propagated to all FEs which are connected to it by any frame-to-frame relation.

6.3.5.2 Propagation of usage information for frame parsing

As shown by Mohit and Narayanan (2003), the frame hierarchy should help in the training of semantic parsers to label sentences with FrameNet frames and FEs. The basic principle is that each FE in an annotated sentence should give some evidence for all the connected FEs, both in terms of the types of fillers expected, and in terms of the expected syntactic realizations. Which of these two types of information for a connected frame can be gleaned from a given annotation set actually depends on which type of frame-to-frame relation we are faced with.

Information on fillers can in principle be gleaned from every type of frame relation (cf. section 6.3.5.1 above), and should be useful for identifying FEs of a target regardless of its part of speech. This means that the annotation of the THEME FE from Taking.seizure.n should help the annotation of the GOODS FE with Theft.steal.v.

Since Core FEs may have an idiosyncratic syntactic realization (see the Giving frame), the only way to have absolutely solid information about the realization of the FEs of an LU is to actually consult the patterns seen in the annotation data. However, in practice there is usually a fair amount of overlap in the way that semantically related predicates (especially those of the same syntactic category) realize their arguments. All of this indicates that a statistical parser may reasonably rely on the realizations of LUs semantically similar to the LU of interest as one input to its statistical decision.

6.4 Metaphor in FrameNet

The FrameNet treatment of metaphor is based on the well-known fact that metaphors differ in their novelty/pre-packaging (see Lakoff and Johnson (1980)). Although this pre-packaging is, of course, a matter of degree, FrameNet makes only a binary distinction along this continuum between "productive" and "lexicalized" metaphor, indicated by whether annotation is done with respect to the source domain of a metaphor (the literal frame) or with respect to the target domain (the frame that more directly encodes what the speaker was trying to say) respectively. We indicate productive metaphor by marking with the sentence-level tag "Metaphor" and annotating only in the source domain if:

- All synonyms of the the current target and related terms have a corresponding alternation between literal and metaphorical uses,[25]
- All FEs of the target domain are mapped to FEs of the source domain,[26]
- The sentence can only be understood by relatively consciously evoking the source domain. Although this criterion is especially vague, it has never been used, in practice, as a deciding factor since it correlates so well with the other criteria.

Thus the following examples are labeled as "productive" metaphor:

[25]Often there are semantic differences in instances of the target or related terms that explain their lack of participation in the metaphor, if we take the care to define semantics of the metaphor and the LUs sufficiently specifically. Thus "She **slid** into a coma" is a productive use of the CHANGE OF STATE IS MOTION metaphor, despite the fact that many of the other targets of the Motion frame (e.g. **roll**, **travel**, **weave**, etc.) cannot be used to describe becoming comatose, and some members (e.g. **circle**) cannot easily be used to describe any change of state. In any frame, there will be LUs with a diversity of semantics, and some specific details of a LU may well block participation in the metaphor, such as in the case of *circle*, which normally evokes a kind of cyclical motion that deemphasizes the changing position of the Theme. Because the change-of-location aspect of motion is deemphasized, *circle* is a poor fit for a metaphor that is entirely concerned with change.

[26]FEs introduced by the metaphor construction are obviously not included in the mapping back to the source domain. For example, in "Why have we chucked all our visions into the **political trashcan** and replaced idealists with policy wonks?", *trashcan*, a member (on the literal side) of the Containers frame, is a valid example of a specific, productive metaphor STATES OF IRRELEVANCY ARE WASTE CONTAINERS (an instance of the STATES ARE LOCATIONS metaphor), since *trashcan*, *garbage*, *dustbin*, etc. can all be used to convey an Undesirable_situation, even though the modifier *political* does not correspond to an FE of the Containers frame, since *political* represents the Target_domain FE of the metaphor construction.

(45) [Cause_motion] Once our parents are dead, we are **catapulted** into becoming the older generation ourselves.

(46) [Absorb_heat] There was a darkness in his eyes, anger **simmering** just beneath the surface.

These criteria are admittedly themselves vague, but hopefully listing them here helps to clarify what factors are important to consider. So, for example, in sentence 45, it is clear that:

1. *catapult* could be replaced by virtually any member of the Cause_motion frame that has sufficient "force" to convey the same (metaphorically) jarring transition,

2. the FEs Cause, Theme, and Goal are all mapped using the STATES ARE LOCATIONS metaphor onto Cause, Patient, and State, and

3. the expression is, impressionistically speaking, very vivid.

We see that the criteria are also met in sentence 46, although in this case, against our general policy, we (experimentally!) made the Emotion_heat frame to represent the particular blended semantics of the metaphor EMOTION IS HEAT.

If any of these criteria are not met (i.e. synonyms and semantically related words do not participate in the same metaphor, some FEs are independent of the source domain, or there is insufficiently vivid evocation of the source frame) we consider the metaphorical use to be a separate sense from the literal. So, for example, in sentence 47, there is a frame element that, in the target domain (involving banking) we could call a Bank FE. There is no correspondent to this conceptualization in any putative source domain, as can be seen in sentences 48 and 49. This is because, in a lexically specific way, the source domain's Goal FE is blended with a concept of a person that keeps and preserves, as seen in examples like sentence 50. In 51, we can also see the FE *Variable* which has no correspondent in literal usages like 52.

(47) [Bank_depositing] Bonlat claimed to have £2.7BN **deposited** *with the Bank of America*.

(48) [Alluvial_depositing] *The river **deposited** sand *with the delta*.

(49) [Placing] *She **deposited** her backpack *with the table*.

(50) Brenda has temporarily left her child *with this woman*.

(51) [Position_on_a_scale] It will be so **low** *in price* that no man making a good salary will be unable to own one.

(52) The great dirigible was so **low** that they could see every crease and contour from nose to fins.

In such cases, an LU should be created in a frame corresponding to the target domain of the metaphor, and, in principle, the metaphorical relation between the domains would be modeled as a Frame-to-Frame relation. In practice, however, FrameNet has not yet added any such links.

In neither type of annotation has FrameNet labeled the source and target domains simultaneously, since we deem this to be worthy of a whole research project by itself. For examples of what such a treatment would look like, see Burchardt et al. (2009, 216-9).