Context Relations

What is a Context Relation?

There are a number of annotation schemes that we use to represent things that change over the course of a text. Context relations, on the other hand, are for capturing static relationships and states that do not change. A context relationship most often holds between two referents in a text, usually two characters. For example, consider the following sentence:

(1) Jack and Jill were siblings.

This sentence indicates that there is a relationship that never changes between Jack and Jill. In this guide, we will use straight underlines to indicate relation arguments, squiggly underlines to indicate relation signals, and subscript text in parenthesis to indicate argument roles.

Rarely a context relationship holds between a referent and an event, or two events. The first is rare because usually when there is a relationship between a referent and event it is one of the referent being an argument to the verb expressing the event, and so this is coded in a different annotation (namely, the semantic role annotation). The second is rare because relationships between events are usually temporal, and these are captured, as well, in a different annotation (namely, the Time Link scheme).

Basic Components of a Context Relation

When marking context relations, one needs to specify the signal (if any) and the arguments. For each argument, one specifies its filler and its role. This section will deal first with simple binary relations, which have exactly two arguments.

Signal

The signal is the word or phrase which indicates the relationship in question. Most context relations have a signal, but there are rare cases when no signal is present. In (1), the signal is the word "siblings." In the following example, we mark a context relation of set/member between the referents "birds" and "sparrow." This context relation is implied by our understanding of the text, but there is no explicit signal that indicates the relationship.

(2) Birds$_{(set)}$ were chirping. A sparrow$_{(member)}$ was the loudest.

Argument Filler

Every context relation has at least two *arguments*, each of which have both a *filler* and a *role*. The filler is a specific referring expression that is most closely associated with signal. If there is a signal that indicates the relation, then usually the arguments will be in some sort of syntactic relationship to the signal, meaning they will be, at the very least, in the same sentence. In some cases, such as (2), the referring expressions are in different sentences.

Argument Role

The argument role is a marking that tells us how each argument is related to the other arguments in the relation. We have the option of choosing roles from two different semantic lexicons: Wordnet and Propbank. Wordnet is an electronic dictionary that provides many hundreds of thousands of word meanings. Propbank is an encoding scheme that provides a list of verb *frames*, which describe, and give names to, the possible arguments to several thousand verbs. Wordnet roles, on average, tend to be more specific. Therefore we have **Rule #1: Prefer Wordnet roles to Propbank roles**.

There are times, though, when only a Propbank role properly expresses the relation in the text. This is usually when the context relation is a static instantiation of a verb, and there is no single meaning in Wordnet that captures the role that referent plays with respect to the verb and relation. A good example is the preposition "in" that expressions presence in a location. Consider the following example:

(3) In a certain kingdom in a certain land there lived a king.

What is the relation between the kingdom and the land? The kingdom is located in the land. Yet, there is no single word in Wordnet that expresses "thing located." Therefore we fall back to Propbank and use the arguments to the verb "locate" which are "thing located" and "location."

In general we want the context relations to mimic the actual expression of relationships in the text, this overrides our preference for Wordnet over Propbank role markings. In (4), then, we prefer to choose the Propbank roles associated with the verb "love" rather than any possible Wordnet markings. **Rule #2: If a relation is expressed with a verb, prefer Propbank role markings.**

(4) John loves Mary.

When there are multiple role possibilities to choose from that seem to fit the relation, observe the following conventions: **Rule #3: For semantically equivalent roles, choose the one whose lexical expression is more common.**

Rule #4: Whenever possible, do not mix Propbank and Wordnet roles, or roles from different Propbank frames, in the same relation.

Redundant Relations

In many texts the same phrase will be repeated time and again, such as "his sword" for a particular sword, or "her mother and father" for a particular set of parents. Although tedious, we will mark all of these, even if they are redundant with previously-marked relationships.

What not to Mark

Rule #5: Do not to mark temporal relationships, relationships between a verb and its arguments, or physical relationships that change over the course of the text. Temporal relationships, such as those between two events, two times, or between an event and a time, are capture in the TimeML suite of annotations and are not covered here. Verbal arguments are captured in the semantic role annotation and should not be marked as context relations.

Physical relationships between referents (i.e., those expressing relations in space, *not* the romantic kind) are tricky, in that both Wordnet and Propbank contain no comprehensive set of roles that can be used to express the infinite variety of possible physical configurations. Fortunately, most physical relations expressed in texts are transient and so do not satisfy our first requirement of being static. Generic "in" relationships can be marked using the "locate" Propbank frame, as shown previously. If there are other examples that fall outside these possibilities, note them and bring them up to your adjudicator.

Finally, be careful to mark just the relationship between the two referring expression, and not any part of the semantics of the phrases in which it the relation might be embedded. Consider the sentence:

(5) He gave his blessing.

Look at the relationship between "his" and "blessing." One's initial temptation is to create a context relation with roles "giver" and "thing given." However, this is taking on some of semantics of the verb "gave" in which "his blessing" is embedded. When taken in isolation, the phrase "his blessing" expresses a generalized ownership that doesn't implying a giving action. For example, we embed the phrase "his blessing" in larger phrases with no give, e.g.,

(6) "His blessing is worthless to me."
(7) "I say 'ha!' to his blessing."

The relationship stays the same, a generic ownership, even though the embedding changes. Therefore the correct relationship between "his" and "blessing" is one of owner/possession.

Complex Context Relations

Until now we have restricted ourselves to simple binary context relations, those with only two arguments. To save time and space, the annotation scheme allows *N-ary* relations to be marked, that is, relations with more than two arguments. Consider the simple case of a 3-way relationship:

(8) Joe$_{(brother)}$, Jim$_{(brother)}$, and John$_{(brother)}$ were brothers.

Here we have three arguments, "Joe," "Jim", and "John." If we were to express the brother/brother relationship in terms of binary relationships, we would need three different relationships: Joe-Jim, Joe-John, and Jim-John. Perhaps this does not sound like many, but with five brothers, we would need twenty, and with seven brothers (not uncommon in some tales) we would need forty-two relations! N-ary relations provide an easy way around this. We merely add multiple arguments to the relation.

Relation Head

The most important thing to remember with a simple N-ary relation is that it is interpreted as meaning that single pairing of the arguments in the relation is a valid relationship. This works for sibling relationships, as in (8), but what about the following set/member relationship:

(9) Jack$_{(member)}$ went up the hill first, Jill$_{(member)}$ second. They$_{(set)}$ fetched a pail of water.

This simple N-ary relationship expresses the following three pair-wise relationships:

1. Jack is a member of the set “they”.
2. Jill is a member of the set “they.”
3. Jack is a member of the member “Jill.”
4. Jill is a member of the member “Jack.”

The first two look fine, that is what we intended to mark. But the third is incorrect. Jack is not a member of Jill no more than Jill is a member of Jack. How do we express this? To do so, we mark “they” as the *head* of the relationship. Marking an argument as a head means that each other non-head argument is in a binary relationship with the head, and non-heads are **not** in binary relationships with each other. You may also mark multiple heads, which means the natural thing: each head/non-head pairing is a valid relationship, but no head/head or non-head/non-head relationship is. This allows to concisely express the four-way relationship in the following example:

(10) Robert$_{\text{(parent-head)}}$ and Mary$_{\text{(parent-head)}}$ were the parents of Jack$_{\text{(child)}}$ and Jill$_{\text{(child)}}$.

Which means:

1. Robert is the parent of child Jack.
2. Robert is the parent of child Jill.
3. Mary is the parent of child Jack.
4. Mary is the parent of child Jill.

But not:

5. Robert is the parent of parent Mary. (and conversely)
6. Jack is the child of child Jill. (and conversely)

Applies-to-Constituents Flag

Sometimes you do not have an individual referring expression to mark as a participant in a relationship. Suppose, you have the following sentence, where this is the only mention of the parents:

(11) The parents of Jack were killed.

The parents are not named individually in the text. The temptation is to mark “parents” in the parent role and “Jack” in child role. But this literally means:

1. The parents is the parent of Jack.

This is incorrect. Each *individual* parent is the parent of Jack; the group of the two parents is not a parent. To express these relationships, there is the *applies-to-constituents* flag, which indicates that it is the constituents of the group that are actually participating in the relation, not the group itself. Therefore we mark it as:

(12) The parents$_{\text{(parent-constituents)}}$ of Jack$_{\text{(child)}}$ were killed.

This discussion of constituents leads us to another convention: **Rule #6: Where possible, prefer marking individual members of a group as participants in a relation rather than a set of which they are a part.** For example, in (10), if Robert and Mary are ever referred to as “they” in the text, then the

phrase “Robert and Mary” will be marked as a separate referent, and you will have the option of marking either “Robert and Mary” as a head, with member flag, or “Robert” and “Mary” individually, each as heads. You should prefer the latter.

Argument Order

When marking context relations, the order of the arguments matters. **Rule #7: Whenever possible, mimic the syntactic order of the expressions.** In (4) we should make sure that John is the first argument and Mary the second. Arguments should be ordered first by when they begin in the text (their lefthand side). If they both begin at the same place, they should be ordered by how long they are (their righthand side). So, for the relation:

(1) Her father and mother.

The argument order for the child/parent relationship should be “her” then “father” and then “mother.” For the set/member relationship the order should be “father,” “father and mother,” and then “mother.”

Rules

1	Prefer Wordnet roles to Propbank roles.
2	If the relation is expressed with a verb, prefer Propbank roles.
3	Do not to mark temporal relationships, relationships between a verb and its arguments, or physical relationships that change over the course of the text.
3	For semantically equivalent roles, choose the one whose lexical expression is more common.
4	Whenever possible, do not mix Propbank and Wordnet roles, or roles from different Propbank frames, in the same relation
5	Do not to mark temporal relationships, relationships between a verb and its arguments, or physical relationships that change over the course of the text.
6	Prefer marking individual members of a group as participants in a relation rather than a set of which they are a part.
7	Whenever possible, mimic the syntactic order of the expressions.

Conventions

As noted previously, you should prefer Wordnet roles generally, and use Propbank roles only when there is no appropriate Wordnet entry or the relationship is expressed with a verb. Even with these conventions, there can be quite a few options for how you mark the roles. Fortunately, there are a number of common context relationships types for which we have established conventional role markings. These are listed below.

Relation Type	Preferred Roles	Examples
Membership	set (SID-07996689-N) member (SID-13810615-N)	father and mother set: father and mother member: father member: mother
Possession	owner (SID-10389398-N) possession (SID-00032613-N)	his car owner: his possession: car
Part-Whole	whole (SID-00003553-N) part (SID-13809207-N)	the shaft of the sword whole: sword part: shaft
Material	whole (SID-00003553-N) material (SID-14580897-N)	a river of blood whole: river material: blood
Body Part	body (SID-05216365-N) body_part (SID-05220461-N)	the dog's neck body: dog body_part: neck
Living your life somewhere	inhabitant (SID-09620078-N) place (SID-08642037-N)	John lived in Kiev inhabitant: John place: Kiev
Living someplace	inhabitant (SID-09620078-N) home (SID-08559508-N)	a bear's den inhabitant: bear home: den
Living someplace one owns	owner-occupier (SID-10389865-N) home (SID-08559508-N)	the man's house owner-occupier: man home: house
Marriage	husband (SID-10193967-N) wife (SID-10780632-N)	his wife husband: his wife: wife
Location	located (SID-02126430-A) location (SID-00027167-N)	a stove in the field located: stove location: field
Support	supported (SID-02350729-A) support (SID-04360501-N)	a house on stilts supported: house support: stilts
Surrounding	Thing surrounding (sourround.01-ARG1) Thing surrounded (surround.01-ARG2)	a wall around a house thing surrounding: wall thing surrounded: house