SpaceEval Annotation Guidelines

July 7, 2014

Contents

1 Introduction

This document provides annotation guidelines for the SpaceEval annotation task. The SpaceEvaltask involves annotating English language texts using a markup scheme drawing largely from ISO-Space, an emerging standard for the annotation of spatial and spatiotemporal information. The SpaceEvalannotation task subsumes previous ISO-Space annotation tasks and includes minor modifications to certain tag identifiers and attribute identifiers. These guidelines provide details on how to apply the specification to a text as well as how to use MAE, an annotation environment, for annotation.

This document is organized as follows: Sections 2 through 4 are divided into subsections covering the different extent tags. In each subsection, first, we describe what should be captured by each tag (the extent) and then we explain what additional information must be supplied (the attributes for that tag). Section 5 then explains how to annotate spatial relationships by constructing links between tagged extents. Section 6 discusses a number of attributes, which are common to many different tag types, that are used to capture quantification information. Section 7 outlines a phased approach for dividing the annotation task. Section 8 provides fully annotated examples and some discussions of the annotation decisions involved. In Section 9 we list some heuristics and other generally helpful tips. Finally, Section 10 provides directions on how to use MAE for SpaceEvalannotation.

Note: When example annotations are provided in this document, they are represented using a predicate argument structure for readability. That is, the tagged extent is marked in the example text with a color and an identifier such as [**extent**$_{id\#}$], and then the attributes assigned for that tag are listed separately in a form such as TAG_TYPE(`id=id#`, `attribute1`=VALUE, `attribute2`=VALUE, ...). When an attribute does not receive a value or the default value is assigned, it may not be included in the briefer examples in Sections 2 through 5.

2 Places, Paths, and Spatial Entities

2.1 Places

The PLACE tag is used to annotate geographic entities and regions like *lakes* or *mountains* as well as administrative entities like *towns* or *countries*. This tag is predominantly inherited from SpatialML, but, for the purposes of this Pilot Annotation, PLACE tags will be manually annotated as follows.

2.1.1 Place Extents

In general, places are referred to explicitly in language such that there is some textual extent that can be captured, i.e., consumed, with a PLACE markup tag. A place is generally a noun phrase (NP) and, in this Pilot Annotation, only the head nouns of NPs should be marked. Example (1) shows some extents that should be captured with the PLACE tag.

(1) a. [**Boston**$_{pl1}$] is north of [**New York**$_{pl2}$].

b. John entered the [**store**$_{pl3}$].

c. [**Boston, MA**$_{pl4}$] – It was reported today that...

d. [**Back Bay**$_{pl5}$] is also served by Amtrak...

e. My father flew to [**Managua**$_{pl6}$] with a silly looking bicycle.

2.1.2 Place Attributes

The majority of attributes for the PLACE tag (shown in Table 1) are inherited from SpatialML.

Attribute	Value
`id`	`pl1`, `pl2`, `pl3`, ...
`dimensionality`	POINT, LINE, AREA or VOLUME
`form`	NAM or NOM
`elevation`	a MEASURE ID
`mod`	a spatially relevant modifier
`dcl`	TRUE or FALSE (default: FALSE)
`countable`	TRUE or FALSE

Table 1: PLACE Tag Attributes

The `dimensionality` attribute specifies how many spatial dimensions the object occupies. POINT specifies a zero-dimensional object that occupies a singular point in space. LINE specifies that the object occupies one dimension of space. A value of AREA (a common value for the PLACE tag) specifies a two-dimensional region of space. VOLUME is used to specify an object which occupies all three spatial dimensions. In selecting a value for this attribute, annotators should consider the context of the PLACE tag since certain values place constraints on the sorts of qualitative relations that the PLACE may participate in. E.g., specifying that a PLACE has a `dimensionality` value of POINT means that it is spatially indivisible, and therefore can't have any other spatial objects inside it. Similarly, a value of AREA would preclude any volumetric object from being inside it.

The next attribute that annotators are tasked with filling is the `form` attribute. Annotators should mark whether the extent of a PLACE tag is a nominal form, as in *the mountain*, or a named location, as in *Mount Everest*. This is done using the `form` attribute with the values NOM and NAM, respectively.

The `elevation` attribute takes a MEASURE ID[1] if a location's elevation is established. E.g., *The city of Zacatecas, at 8,000 feet ...* directly introduces an elevation in feet above sea level. This is an optional attribute, thus if no explicit elevation is mentioned in the text, the annotator should not enter any value and leave this attribute unspecified.

The `mod` attribute is used to capture cases like *tall mountain*, *upper observation deck*, and *longest lake*, where *tall*, *upper*, and *longest* do not constrain the location of the entities they modify, but they do contribute spatial information. The following is a list of scalar adjectival terms in English that may act as spatial modifiers, however it is not intended to be an exhaustive list of modifiers, and it is possible that these terms may not always act as modifiers: *long, short, tall, low, high, deep, shallow, thin, narrow, wide, upper, lower, inner, outer, inland, near, far, northern.* Remember that it is only the head of an NP that should be marked as a PLACE, thus these modifier terms are indicated with the `mod` attribute as shown in Example (2a) below. Like the `elevation` attribute, if no spatially relevant modifier is mentioned, this attribute should be left unspecified.

The `dcl` attribute is used to specify the Document Creation Location (DCL) associated with a text. The DCL is a special location that the author must declare to be the location where the text was written. Such a location is often used in newswire text, though an author may indicate where a document is written in other types of texts. Any given text can have at most one DCL and the

[1]Refer to Section 4.2 for more details regarding the MEASURE tag

default value for this attribute is set to FALSE. If the annotator can determine that a location tag is the unique DCL for a text, the value of the `dcl` attribute for the first mention of that location should be set to TRUE and, from that point on, no other location tag in the text can be declared the DCL, even if the same location is referred to later in the text. Example (2b) shows an example where it would be appropriate to mark a PLACE tag as the DCL. Note that the second reference to *Boston* in this example is not the DCL.

The `countable` attribute is used to distinguish between countable (e.g., *continents, countries, cities, suburbs, towns, parks, lakes, islands, ...*) and uncountable (e.g., *highlands, foothills, waters, backcountry, countryside, ...*) location terms. Note that these are not exhaustive lists. In fact, in English, it is possible to coerce many countable terms to act like mass terms. Annotators should make their decisions for this attribute based on the context. Refer to Section 6 for a more complete explanation of the `countable`, `gquant` and `scopes` attributes.

Note: Sometimes it is necessary to create non-consuming PLACE tags for locations when they are referred to indirectly in the text. For further details regarding non-consuming tags refer to Section 2.4.

(2) a. I camped one night ... in the small [**town**$_{pl1}$] of [**Monte Frio**$_{pl2}$].
PLACE (`id`=`pl1`, `form`=NOM, `mod`="small")
PLACE (`id`=`pl2`, `form`=NAM)

b. [**Boston**$_{pl3}$], [**MA**$_{pl4}$] – It was reported today in [**Boston**$_{pl5}$] that...
PLACE (`id`=`pl3`, `form`=NAM, `dcl`=TRUE)
PLACE (`id`=`pl4`, `form`=NAM, `dcl`=FALSE)
PLACE (`id`=`pl5`, `form`=NAM, `dcl`=FALSE)

c. The new skyscraper at [**111 Huntington Avenue**$_{pl6}$] was completed in 2002...
PLACE (`id`=`pl6`, `form`=NAM)

d. [**The Plaza de Armas**$_{pl7}$] is [**one**$_{pl8}$] of the most visited [**places**$_{pl9}$] in [**Cochabamba**$_{pl10}$]
PLACE (`id`=`pl7`, `form`=NAM)
PLACE (`id`=`pl8`, `form`=NOM)
PLACE (`id`=`pl9`, `form`=NOM)
PLACE (`id`=`pl10`, `form`=NAM)

e. There are two [**Dunkin' Donuts**$_{pl11}$] in the [**area**$_{pl12}$].
PLACE (`id`=`pl11`, `form`=NAM, `countable`=TRUE)
PLACE (`id`=`pl12`, `form`=NOM, `countable`=TRUE)

f. I was impressed that both in the [**lowlands**$_{pl13}$] and the [**highlands**$_{pl14}$], the [**countryside**$_{pl15}$] is entirely cleared for cattle grazing.
PLACE (`id`=`pl13`, `form`=NOM, `countable`=FALSE)
PLACE (`id`=`pl14`, `form`=NOM, `countable`=FALSE)
PLACE (`id`=`pl15`, `form`=NOM, `countable`=FALSE)

2.2 Paths

The PATH tag captures locations where the focus is on the potential for traversal or the location functions as a boundary between regions. This includes common nouns like *road*, *river*, and *border*

as well as proper names like *Route 66* and *Kangamangus Highway*. The distinction between places and paths is not always clear-cut. Take for example the case of *river* in *follow the river* and *cross the river*. The first case clearly is a path, but one could argue that in the second case the traversal functionality is not accessed and therefore *river* could be annotated as a PLACE. While the inferences that can be made with both places and paths are actually identical, it is important for inter-annotator agreement to try to apply a diagnostic test to determine if a location should be tagged as a PLACE or a PATH.

One useful heuristic is the "*be-at/take*" test. That is, while you can *be at* a place or path location, you can only *take* or *follow* a path. In general, we are looking for consistency such that when an annotator comes across something like a *road*, which typically has the potential to be traversed or followed, the annotator need not consider whether that particular road is actually functioning more like a place in that particular context. For example, *roads, rivers, alleys, walls, stairways, shorelines, ridgelines* and *mountain ranges*, among other locations, should be consistently marked with the PATH tag since they all function as static paths or boundaries.

2.2.1 Path Extents

As with places, paths occur predominantly within noun phrases, and, consistent with the PLACE tag, only the head of the NP should be captured as the extent for the PATH tag. Note that the PATH tag is only used to capture static, non-stative paths. Stative paths and event-paths introduced by motion predicates are captured using the MOTION tag[2]. As an example, *trip* in *the **trip** from the parking lot to the beach* would not be tagged as a PATH, but *walkway* in *the **walkway** from the parking lot to the beach*, would.

2.2.2 Path Attributes

You will notice that the PATH tag attributes in Table 2 overlap with with the PLACE tag. The attributes specific to the PATH tag that annotators must consider, namely `beginID`, `endID`, and `midIDs`, are discussed in this section.

Attribute	Value
`id`	`p1`, `p2`, `p3`, ...
`type`	BODYOFWATER, MTS, ROAD
`beginID`	ID of a location/entity/event tag whose location is a bounding point for the PATH
`endID`	ID of a location/entity/event tag whose location is a bounding point for the PATH
`midIDs`	ID(s) list of midpoint locations
`dimensionality`	LINE, AREA or VOLUME
`form`	NAM or NOM
`elevation`	a MEASURE ID
`mod`	a spatially relevant modifier
`countable`	TRUE or FALSE

Table 2: PATH Tag Attributes

Excluding continuous loops, or rays with an individual bounding point, paths typically have discernible endpoints. However, the locations of a path's endpoints may not be explicit in the text. Example (3a) illustrates a PATH tag for which the endpoints happen to be explicit and Example (3b) shows a case where the endpoints are unspecified.

[2]Refer to Section 3.2 for more details regarding the MOTION tag.

(3) a. ... the [**railroad**$_{p1}$] between [**Boston**$_{pl1}$] and [**New York**$_{pl2}$] ...
PATH (id=p1, beginID=pl1, endID=pl2, form=NOM)

b. We descended into a long [**valley**$_{p2}$].
PATH (id=p2, form=NOM, mod="long")

In such a case as Example (3a), the annotator should include the IDs for the relevant end point locations as values for the appropriate attributes in the PATH tag. Some paths may even have an explicit midpoint as shown in Example (4a) below. PATH tags also have the optional mod attribute as in Examples (4b) and (4c).

(4) a. John took the [**road**$_{p1}$] through [**Boston**$_{pl1}$].
PATH (id=p1, midIDs=pl1, form=NOM)

b. The car took the upper [**ramp**$_{p1}$].
PATH (id=p2, form=NOM, mod="upper")

c. The police assured me that the [**road**$_{p1}$] ahead was safe during the day, but that is also what we had heard of the [**road**$_{p2}$] behind.
PATH (id=p3, form=NOM, mod="ahead")
PATH (id=p4, form=NOM, mod="behind")

The form attribute indicates whether the PATH is a nominal form as in *road* or a named path as in *Massachusetts Avenue*. The remaining attributes are the same as those for the PLACE tag. Additionally, like PLACE tags, it may be necessary in certain cases to create non-consuming PATH tags (i.e., PATH tags that are note associated with any explicit extent in the text). Refer to Section 2.4 for more details regarding non-consuming tags.

2.3 Spatial Entities

Generally, anything that is spatially relevant but that does not fit into either the PLACE or PATH categories is considered to be a spatial entity. That is to say, a spatially relevant entity must be both located in real-space and participate in a link tag. In practice, moving objects and objects that have the potential to move are most commonly tagged as SPATIAL_ENTITY. In both Examples (5a) and (5b), *car* should be marked as a SPATIAL_ENTITY. In the first case, it is the mover in a motion event and, however in the second case, it behaves like a PLACE. Note, though, that it should still be annotated as a SPATIAL_ENTITY since cars function as a form of transportation, regardless if they are mentioned in a context where they don't happen to be in motion. Spatial objects, such as buildings, which do not function as a mode of transportation, should normally be annotated with the PLACE tag, though they may seem to function like more like spatial entities in certain contexts. E.g., *tower* is annotated as a SPATIAL_ENTITY in Example (5e), since it is participating as a mover in a motion event, rather than functioning like a PLACE.

(5) a. The [**car**$_{se1}$] drove down the street.

b. [**John**$_{se1}$] arrived at the [**car**$_{se2}$].

c. My [**father**$_{se1}$] and [**I**$_{se2}$] biked for two days.

d. Two [**men**$_{se1}$] with machetes and masks jumped out of the forest.

e. The [**tower**$_{se1}$] toppled when it was struck by lightning.

2.3.1 Spatial Entity Extents

The same extent rules apply for SPATIAL_ENTITY as do for PLACE and PATH. I.e., only the head of the NP should be captured as the extent.

2.3.2 Spatial Entity Attributes

Because tagging something as a SPATIAL_ENTITY is akin to treating it like a location, the SPATIAL_ENTITY tag shares its attributes with the PLACE tag type. The list of SPATIAL_ENTITY attributes is shown in Table 3.

Attribute	Value
`id`	`se1`, `se2`, `se3`, ...
`type`	FAC, VEHICLE, PERSON, DYNAMIC_EVENT, ...
`dimensionality`	POINT, LINE, AREA or VOLUME
`form`	NAM or NOM
`mod`	a spatially relevant modifier
`countable`	TRUE or FALSE

Table 3: SPATIAL_ENTITY Tag Attributes

The attributes for SPATIAL_ENTITY tags are basically the same as for those of PLACE tags. The `form` attribute should be used to specify whether the form of the textual extent of the SPATIAL_ENTITY is a proper name or a nominal. If a spatially relevant modifier is present, it should be entered as the value for the `mod` attribute. The `countable` attribute is used to distinguish between countable (e.g., *people, cars, ships, planes*) and uncountable (e.g., *water, oil, sand, concrete*) entities. To reiterate, countability is context-dependent.

(6) a. [**John**$_{se1}$] visited Boston.
SPATIAL_ENTITY (`id`=`se1`, `form`=NAM, `countable`=TRUE)

b. Two [**cars**$_{se2}$] are parked on the street.
SPATIAL_ENTITY (`id`=`se2`, `form`=NOM, `countable`=TRUE)

c. Two [**men**$_{se3}$] with machetes and masks jumped out of the forest.
SPATIAL_ENTITY (`id`=`se3`, `form`=NOM, `countable`=TRUE)

d. The [**dogs**$_{se4}$] in Costa Rica are more likely to chase [**cyclists**$_{se5}$]
SPATIAL_ENTITY (`id`=`se4`, `form`=NOM, `countable`=TRUE)
SPATIAL_ENTITY (`id`=`se5`, `form`=NOM, `countable`=TRUE)

e. So much [**oil**$_{se6}$] had been extracted from the ground...
SPATIAL_ENTITY (`id`=`se6`, `form`=NOM, `countable`=FALSE)

2.4 Special Section: Non-consuming Tags

Annotators may encounter situations where a spatially relevant entity is referenced indirectly. In such situations, it is possible to create so-called 'non-consuming' tags, whose tag IDs can then be filled as attributes for other tags or participate in links where appropriate.

Normally, for 'consuming' tags, there is some word or string in the text which is associated with the tag (called the tag's textual extent). Non-consuming tags are named as such because they have no associated extent in the text which is 'consumed'. That is, the extent of a non-consuming tag is a null or empty string.

Generally, non-consuming tags are not necessary to capture relevant spatial objects and relations. For this reason, non-consuming tags should be used sparingly. If an annotator is considering using a non-consuming tag, it may be worth reconsidering if there is anything spatially relevant being described at all, or whether there is an extent that was missed.

That said, the following are situations where the use of non-consuming tags is necessary:

1. ***Locations referenced by a* MEASURE.** When a relevant location is referenced indirectly by an elevation that will be captured as a MEASURE tag[3], a non-consuming PLACE tag can be used so that its PLACE ID may fill an attribute for other tags or links. In cases such as Example (7b) where the MEASURE does not indicate an elevation, it is necessary to create an MLINK[4] that links the non-consuming PLACE tag to some other entity.

   (7) a. John climbed to [**9,000 feet**$_{me1}$]. [$\varnothing_{pl1}$][5]
   PLACE (`id`=`pl1`, `elevation`=`me1`)

   b. We camped [**three miles**$_{me2}$] from the [**river**$_{p1}$]. [$\varnothing_{pl2}$]
   PLACE (`id`=`pl2`)
   MEASURE (`id`=`me2`, `value`=6, `unit`=miles)
   MLINK (`id`=`ml1`, `trajector`=`pl2`, `ground`=`p1`, `relType`=DISTANCE, `val`=`me2`)

2. ***Locations implied by 'cross' and 'across'.*** When the path traversed by an object 'crosses' a region, but there is no explicit PATH in the text, the use of non-consuming PLACE tags may be appropriate. This may occur in cases of CROSS class MOTION events. It also may be necessary in instances where an entity is located 'across' from another relative to some reference location. In Example (8a), the event-path, that is, the path traversed by *John*, is interpreted as entirely within the *town*, so the `source`, and `goal` for the MOVELINK[6] that would be triggered by the motion-event *walked* must be created by the annotator as non-consuming tags. The IDs of these non-consuming PLACE tags—`pl2` and `pl3`—can then participate in links with the PLACE tag for *town*—`pl1`. The QSLINK tags[7] `qsl1` and `qsl2` illustrate this. Additionally the non-consuming PLACE tag IDs are linked to the tag for *town* and each other via an OLINK to establish a 3-way relation such that, relative to `pl2`, `pl3` is across the *town*.

   (8) a. John walked across [**town**$_{pl1}$]. [$\varnothing_{pl2}$] [$\varnothing_{pl3}$]
   PLACE (`id`=`pl1`)
   PLACE (`id`=`pl2`)
   PLACE (`id`=`pl3`)
   QSLINK (`id`=`qsl1`, `relType`=IN, `trajector`=`pl2`, `ground`=`pl1`)
   QSLINK (`id`=`qsl2`, `relType`=IN, `trajector`=`pl3`, `ground`=`pl1`)
   OLINK (`id`=`ol1`, `relType`="ACROSS", `figure`=`pl3`, `ground`=`pl1`, `frame_type`=RELATIVE, `referencePt`=`pl2`)

[3] Refer to Section 4.2 for more details regarding the MEASURE tag
[4] Refer to Section 5.4 for more details regarding MLINK tags.
[5] The ∅ symbol is used here to indicate a non-consuming tag.
[6] Refer to Section 5.3 for further details regarding MOVELINK tags.
[7] Refer to Section 5.1 for further details regarding QSLINK tags.

b. The [**forest**$_{pl4}$] sits across the [**border**$_{p1}$]. [$\varnothing_{pl5}$]
PLACE (id=pl4)
PATH (id=p1)
PLACE (id=pl5)
OLINK (id=ol2, relType=ACROSS, figure=pl4, ground=p1,
frame_type=RELATIVE, referencePt=pl5)

3. ***Sets whose members are mentioned.*** When sets of objects are referenced in the text, sometimes it will be necessary to create non-consuming tags to reify those sets so that they may participate within links. For instance, in the case of a split antecedent, a non-consuming tag must be created to reify the antecedent set in order to create a split or sub-coreference relation. The process for handling coreference is discussed in more depth in Section 2.5.

2.5 Special Section: Metalinks

Often, in natural language texts, the same entities are referenced multiple times or a group of entities is introduced whose members maybe be referenced individually or as subsets. These relationships between sets and their members are not really spatial in nature, but it is important to keep track of them nonetheless. To do this, the METALINK tag is used.

The METALINK tag type is an all-purpose tag for relating objects in a non-spatial way. This link tag type's current intended purpose is for relating spatial entities such as the ones in Example (9) below. There are three relevant METALINK attributes that annotators must fill as shown in Table 4.

Note: Unlike other link tag types, the arguments of METALINK tags may occur within separate sentences from each other. That is to say, METALINK relations may be non-local such that the participant tags may occur anywhere within a text. In order to reduce the number of METALINK tags that must be created, it is preferable to link back to the earliest mention in the text rather than to the closest mention.

Attribute	Value
id	meta1, meta2, meta3, ...
objectID1	ID of a location/entity/event tag
objectID2	ID of the related location/entity/event tag
relType	COREFERENCE, SUBCOREFERENCE, SPLITCOREFERENCE

Table 4: METALINK Attributes

The key attribute for the METALINK tag type is relType. The value for this attribute describes the relationship between the entities identified in objectID1 and objectID2. When the entities refer to the same object, the COREFERENCE value should be used to specify that both mentions refer to the same entity. When an entity, or subset of entities, is mentioned, which is part of a larger set, then the subset will be linked with the subsuming set via a METALINK with a relType value of SUBCOREFERENCE. When an anaphor has a split antecedent, then a non-consuming tag must created to represent the antecedent set. Then the individual parts of the antecedent set should be linked with SUBCOREFERENCE METALINK tags to the non-consuming tag, which represents the set. Finally, this non-consuming tag will be linked via a SPLITCOREFERENCE METALINK to the mention of the subsuming set.

In the case of a relType with a value of COREFERENCE, the order of the arguments is not significant, since the relation is symmetric. However, for a relType value of SUBCOREFERENCE or

SPLITCOREFERENCE, the order is important. Notice that for SUBCOREFERENCE, the smaller set (or individual) should be stored in `objectID1` and the subsuming set is stored in `objectID2`. In the case of SPLITCOREFERENCE, the mention in the text will be `objectID1` and the non-consuming tag that is created to represent the split antecedent set will be `objectID2`. Example (9) includes several cases of METALINK tag annotations for reference.

(9) a. Two [**cars**$_{se1}$] are on the street. [**One**$_{se2}$] of [**them**$_{se3}$] turns left.
SPATIAL_ENTITY (`id`=se1, `form`=NOM, countable=TRUE)
SPATIAL_ENTITY (`id`=se2, `form`=NOM, countable=TRUE)
SPATIAL_ENTITY (`id`=se3, `form`=NOM, countable=TRUE)
METALINK (`id`=meta1, `objectID1`=se3, `objectID2`=se1, `relType`=COREFERENCE)
METALINK (`id`=meta2, `objectID1`=se2, `objectID2`=se1, `relType`=SUBCOREFERENCE)

b. [**John Smith**$_{se4}$] left Boston. [**John**$_{se5}$] arrived in New York.
SPATIAL_ENTITY (`id`=se4, `form`=NAM, `countable`=TRUE)
SPATIAL_ENTITY (`id`=se5, `form`=NAM, `countable`=TRUE)
METALINK (`meta3`, `objectID1`=se5, `objectID2`=se4, `relType`=COREFERENCE)

c. [**John**$_{se6}$] and [**Mary**$_{se7}$] [$\varnothing_{se8}$] met at the store. [**They**$_{se9}$] went shopping.
SPATIAL_ENTITY (`id`=se6, `form`=NAM, `countable`=TRUE)
SPATIAL_ENTITY (`id`=se7, `form`=NAM, countable=TRUE)
SPATIAL_ENTITY (`id`=se8, `form`=NOM, `countable`=TRUE)
SPATIAL_ENTITY (`id`=se9, `form`=NOM, `countable`=TRUE)
METALINK (`id`=meta4, `objectID1`=se6, `objectID2`=se8, `relType`=SUBCOREFERENCE)
METALINK (`id`=meta5, `objectID1`=se7, `objectID2`=se8, `relType`=SUBCOREFERENCE)
METALINK (`id`=meta6, `objectID1`=se9, `objectID2`=se8, `relType`=SPLITCOREFERENCE)

d. [**The Prudential Tower**$_{pl1}$], also known as [**The Prudential Building**$_{pl2}$] or, colloquially, [**The Pru**$_{pl3}$]...
PLACE (`id`=pl1, `form`=NAM, countable=TRUE)
PLACE (`id`=pl2, `form`=NAM, countable=TRUE)
PLACE (`id`=pl3, `form`=NAM, countable=TRUE)
METALINK (`id`=meta7, `objectID1`=pl2, `objectID2`=pl1, `relType`=COREFERENCE)
METALINK (`id`=meta9, `objectID1`=pl3, `objectID2`=pl1, `relType`=COREFERENCE)

Note: Note that while it is not incorrect to create a link between [**The Prudential Building**$_{pl2}$] and [**The Pru**$_{pl3}$], since they are both linked to [**The Prudential Tower**$_{pl1}$], which is the earliest mention, it would be redundant. Annotators should avoid creating redundant links by linking only to the earliest mention.

e. [**I**$_{se10}$] biked with [**Gregg**$_{se11}$] and [**Brooks**$_{se12}$] [$\varnothing_{se13}$] for one more day. [**We**$_{se14}$] climbed over the mountains.
SPATIAL_ENTITY (`id`=se10, `form`=NOM, `countable`=TRUE)
SPATIAL_ENTITY (`id`=se11, `form`=NAM, `countable`=TRUE)
SPATIAL_ENTITY (`id`=se12, `form`=NAM, `countable`=TRUE)
SPATIAL_ENTITY (`id`=se13, `form`=NOM, `countable`=TRUE)
SPATIAL_ENTITY (`id`=se14, `form`=NOM, `countable`=TRUE)
METALINK (`id`=meta8, `objectID1`=se10, `objectID2`=se13, `relType`=SUBCOREFERENCE)
METALINK (`id`=meta9, `objectID1`=se11, `objectID2`=se13, `relType`=SUBCOREFERENCE)

METALINK (id=meta10, objectID1=se12, objectID2=se13, relType=SUBCOREFERENCE)
METALINK (id=meta10, objectID1=se14, objectID2=se13, relType=SPLITCOREFERENCE)

3 Non-Motion Events and Motion Events

3.1 Non-Motion Events

For the purposes of this task, non-motion events are taken to be a sub-species of event. The term "event" is borrowed directly from TimeML. It has the following definition:

> *Event* is used as a cover term for situations that *happen*, *occur*, *hold*, or *take place*. Events can be punctual (Example (10)) or last for a period of time (Example (11)).

(10) a. Ferdinand Magellan, a Portuguese explorer, first **reached** the islands in search of spices.

b. A fresh flow of lava, gas and debris **erupted** there Saturday.

(11) a. 11,024 people, including local Aeta aborigines, **were evacuated** to 18 disaster relief centers.

b. "We're **expecting** a major eruption," he said in a telephone interview early today.

c. Tropical Depression Seven **formed** Wednesday in the far eastern Atlantic.

In this task, non-motion events are tagged with the NONMOTION_EVENT tag in order to capture TimeML events that are spatially relevant in that ***they do not involve movement, but they are directly related to another spatial element by way of a link tag***.

3.1.1 Event Extents

The following information on events is taken directly from the TimeML Annotation Guidelines. This information is meant to help identify non-motion events in general, but, remember that for this task, only events that are spatially interesting should be tagged. For this task, annotate an NONMOTION_EVENT if all of the following hold: the NONMOTION_EVENT is directly, spatially related to a PLACE, PATH, SPATIAL_ENTITY or other NONMOTION_EVENT or MOTION, and the NONMOTION_EVENT advances the narrative. E.g., *party* in *the party lasted for hours* wouldn't be annotated as an NONMOTION_EVENT (at least not without further context), but in *the party on the roof lasted for hours* it would be. Additionally, for the purposes of this task, emotional or experiential states should not be tagged as an NONMOTION_EVENT since such eventualities do not exist in "real space" (i.e., emotions and other experiences should be considered non-spatial even if they are associated with an "experiencer" that is a spatial entity).

TimeML events are always single tokens (words) in the text. For example, in the phrase *may not leave*, only **leave** is actually tagged as an event while the modality and polarity markers are represented as attributes for that event. So, we say that the *head* of the expression is what is actually marked up. In the examples that follow, the actual extent of the event is in **boldface** and the rest of the expression, whether in be a verb phrase, noun phrase, or some other construction, is surrounded by square brackets.

As always, it is a good idea to be thinking ahead to the next part of the annotation as you mark event extents. Remember that events must participate in some kind of TimeML link. The discussion and examples below should help you identify event extents in most cases, but, if you are unsure about a particular event candidate, think about how it would be anchored to a time or ordered relative to other events in the annotation.

Events denoted by VERBS**.** We consider that all verbal predicates express an event and hence will be marked up as such, including those which denote states. In the following examples, the event extent to be marked appears in bold.

(12) a. The citizens of Loudon county [**embrace**] religious and human freedom.

b. Amir, 27, [is **serving**] a life sentence for the November 1995 assassination of Rabin at a Tel Aviv peace rally.

c. President Clinton [**says**] he and Blair [will **stand**] together not just on Iraq but also on arresting the terrorists.

d. We [are **going**] [to **maintain**] our forces in the region for the foreseeable future.

e. "Some of these bands of kidnappers [are **made** up] of foreigners," Toledo [**said**].

f. Only Saudi Arabia [**has**] more oil reserves.

g. With Kuwait now [**annexed**], Iraq [has **control**] of 20 percent of the world's oil reserves.

h. Israel [has been **scrambling**] [to **buy**] more masks abroad, after a shortage of several hundred thousand gas masks.

For simple VPs (Example (13a)), the event tag covers just the verbal head, as stipulated by the general rule. As in several of the above examples, when a complex VP (Example (13b)) is present (i.e. the verb is accompanied by auxilliaries and related particles), the event extent is again only the head of the VP. The same is true for phrasal verbs (Example (13c)) as well as idioms (Example (13d)). Additional examples of each of these constructions are shown below:

(13) a. A fresh flow of lava, gas, and debris [**erupted**] there Sunday.

b. Amir [may have been **trying**] [to **impress**] Har-Shefi.

c. Additional distribution centers would be [**set** up] next week.

d. Even more hard drives [**kick** the bucket].

Events denoted by NOUNS**.** Not all nouns are considered TimeML events. A nominal event must be able to appear in at least two of the following contexts:

- NOUN lasted for several seconds/minutes/days/years/...
- NOUN was very fast/immediate/...
- NOUN took/takes/will take place in TEMPORAL EXPRESSION
- NOUN began/continued/ended in TEMPORAL EXPRESSION

Event-denoting nouns will be marked up according to the following guidelines:

1. ***Prenominal modifiers.*** Event-donoting nouns acting as prenominal modifiers, like those underlined below, are NEVER annotated as events.

   (14) a. Many of the same reactions occur in a [panic attack].
   b. The likely shape of the [leadership contest] emerged yesterday as Labour began an inquest on its [election defeat].

   The same policy applies for present participle forms as in:

   (15) a. the waiting room
   b. knitting needle
   c. shooting match
   d. drinking problem

2. ***Sortal states.*** Sortal states are generally expressed by:

   - Agentive nominals; that is, nouns referring to the agents of certain activities or actions. In most cases they correspond to deverbal nouns (e.g., *employer, winner, customer*), but not always (e.g., *passenger, pedestrian*).
   - Nouns denoting professions, roles, or positions (e.g., *doctor, plumber, CTO, scholar*).
   - Rigid designators; that is, terms referring to the same entity in all possible worlds; e.g., *the 3rd president of America*.

   These elements will be annotated as events ONLY when functioning as the head of a *predicative complement*. A predicative complement expresses a predication and is the complement of a verb belonging to one of the classes listed below, among others. In the examples that follow, the predicative complement is in square brackets and the sortal state in bold face.

   - Copulative predicates (e.g., *to be, seem*, etc.)–underlined below.
     (16) Currently she is [**President** of the Macedonian Teachers Association in Victoria].
   - Inchoative predicates (e.g., *become*). They express the coming to existence of a situation.
     (17) In 1821 Simn Bolvar became [the first **president** of the Republic of Bolivia].
   - Aspectual predicates (e.g., *begin, continue, end, finish*, etc.).
     (18) He continued as [**director** of research at the Bell Telephone. Laboratories established in 1925].
   - Change of state predicates, that is, any predicate expressing the coming to be or ending of the position expressed by the sortal state, such as *retire, appoint, elect, resign*, etc.
     (19) In 1998 she retired as [a **professor** emeritus of mathematics].
   - Predicates of evaluation and description (e.g., *consider, describe, depict, evaluate*, etc.).
     (20) Daimler is considered [the first **inventor** to have invented a practical internal-combustion engine].

3. ***All other event-denoting nouns.*** will ALWAYS be marked as events. For example:

   (21) a. An **embargo** on Iraq could plunge the U.S. into a **depression** and the rest of the world into an economic **crisis**.
   b. The economic **chokehold** appears to be working.
   c. Through the Pope, Cuba can begin a more productive **relationship** with the world.
   d. The men explained that that was an **attack** by masked individuals.
   e. Here again, it was the democratic government that quickly became a **threat** to the United States.
   f. Moscow depicted the **situation** as a **conflict** between "the forces of democracy and progress against those of reaction."

4. ***Multiword nomial events.*** The event tag extends only over the head noun, disregarding any determiners, specifiers, complements, or modifiers, as in the following examples:

   (22) a. [The financial **assistance** from the World Bank and the International Monetary Fund] are not helping.
   b. [The Pope's **visit**] will persuade a great many more Cubans to break loose of the Cuban government.

   The same policy is used for named events such as *the Vietnam **War*** and *the Industrial **Revolution***. Finally, as was seen in the examples for sortal states, only the head noun of the construction is tagged as the event.

Events denoted by ADJECTIVES**.** Adjectives generally express a property or attribute of an entity, and as such, they denote an event of a stative nature. Adjectives can appear in attributive or predicative position. *Attributive adjectives* function as premodifiers of the noun:

- **furious** reaction
- **unbearable** pain
- **fair** trial
- **beautiful** garden

On the other hand, *predicative adjectives* act as the predicative complement of a verb belonging to one of the types listed below, among others. In the examples, the main verb is underlined and the predicative adjective is in bold face.

- Copulative predicates (e.g., *be, seem*, etc.).

  (23) The students <u>seemed</u> **exhausted** after three weeks of classes.

- Inchoative predicates (e.g., *become, turn into*). They express the coming to existence of a situation.

  (24) The Chinese dissident said he left China because his life <u>became</u> **unbearable** there.

- Aspectual predicates (e.g., *begin, continue, finish, terminate*, etc.).

(25) Families kept **hopeful** and many did see the return of their loved ones.

- Causative predicates (e.g., *cause, make*, etc.).

  (26) Dan Hollander, skater and entertainer, really made the audience **happy**.

- Change of state predicates in general.
- Predicates of perception (e.g., *look, hear*, etc.).

  (27) Ellen DeGeneres and Portia de Rossi looked **ecstatic** as they married in an intimate ceremony on Saturday.

- Predicates of evaluation and description (e.g., *consider, describe, present*, etc.).

  (28) He is often characterized as **eccentric**.

When annotating adjectives, the following guidelines apply:

1. ***Attributive adjectives.*** NO adjective in attributive position will be marked up as an event.
2. ***Predicative adjectives.*** We will only annotate as events those predicative adjectives that express a non-persistent property of the entity denoted by the noun.

There are indeed many properties of entities that can be considered as non-persistent. People, for instance, can change their nationality and cars can be painted into a color different that their original one. According to these considerations, the adjectives *Senegalese* in Example (29a), and *red* in Example (29b) should be marked up as events.

(29) a. Most of the people aboard were **Senegalese**, including many schoolchildren.

b. The defendants car was **red**.

However, we will **only** annotate those adjectives satisfying at least one of the following conditions:

1. The property they denote is clearly fluid, non-persistent, such as *red* in Example (30a) but not in Example (30b).

   (30) a. Nycks face turned **red** with shame and anger.

   b. The defendants car was **red**.

2. The property is presented as temporally bound to a particular point or period of time as in Example (31).

   (31) France was**under-developed** in the eighteenth century, and Germany at the beginning of the nineteenth.

3. The property is presented as the opinion, knowledge, or belief of somebody, or as a matter under discussion; e.g., *Senegalese* in Example (32a) but not in Example (32b).

   (32) a. The government claims that some abductees are themselves **Senegalese**.

   b. Most of the people aboard were **Senegalese**, including many schoolchildren.

Some cases may be hard to evaluate. **As a general rule, in case of doubt, do NOT annotate adjectives as events**.

Events denoted by PPs**.** PPs denoting events will be annotated ONLY when functioning as predicative complements. This involves PPs that are complement of verbs belonging to the types listed below, among others. The PP in each example is in square brackets while the actual extent of the event is in boldface.

- Copulative predicates (e.g., *be, seem*, etc.).

  (33) No woman has been [**in** charge] of a mission until now.

- Inchoative predicates (e.g., *become*). They express the coming to existence of a situation.

  (34) As a boy, he was an excellent horseman and also became [**on** good terms] with a tribe of Sioux Indians.

- Aspectual predicates (e.g., *begin, continue, finish, end*, etc.).

  (35) The US economic and political embargo has kept Cuba [**in** abox].

- Causative predicates (e.g., *cause, force, put*).

  (36) She says this puts the very existence of women [**at** the hands of their husbands].

Note that when an event is expressed by means of a PP, **only the head preposition is annotated as the event**. Note that not all PPs which denote events will have its preposition head tagged as such, but the verb, noun, or adjective head of the complement of that preposition instead, whenever this is the element conveying the event. We mark up the head preposition only *in case the verb, noun, or adjective head within the PP does not denote the event itself.* The following are examples where the element to be tagged as an event is the head of the preposition complement (boldface), and not the preposition (underlined). Example (37a) illustrates the case for verbs, Example (37b) for nouns, and Example (37c) for adjectives.

(37) a. He glared at Conroy [without **seeing** him] and charged back.

b. The programme began [with an **interview** with someone from that TV series].

c. He is often characterized [as **eccentric**].

Events denoted by OTHER ELEMENTS**.** Events can also be referred to by other elements, most typically locative adverbs such as *here* and *there*. They will be marked up only when functioning as a predicative complement.

(38) a. We're **there** to stay for a fairly lengthy period.

b. We are **here** because what happens on this island will also have an impact on the United States.

Complex event constructions. It is often the case that multiple events are mentioned in the same construction. In some cases, both of these are verbs, but it is also possible for NPs, APs, and PPs to be involved. We now discuss some of the most common of these constructions.

1. **Aspectual Constructions.** These consist of an aspectual verb (e.g., *begin, stop, end, keep*) or noun (*beginning, ending*), and an event-denoting complement, which can be expressed by either a VP (Examples (39a) and (39b)) or an NP (Example (39c)). BOTH the aspectual predicate and its complement will be tagged as independent events. In the sentences below, the aspectual predicate is in bold face, its complement head in italics, and the two elements to annotate are delimited by square brackets. The sentences in Example (39) illustrate cases in which the aspectual predicate is a verb, whereas the sentences in Example (40) exemplify cases involving aspectual nouns.

   (39) a. US did not [**stop**] [*interfering*] in other countries' policies.
   b. They probably would have [**began**] [*responding*] to President Reagan's 600 ships plan with new construction.
   c. In 1939, Bohannon [**began**] the [*construction*] of Hillsdale.

   (40) a. The banks must wait at least 30 days before [**closing**] the [*purchase*].
   b. The [**outbreak**] of holy [*war*] could bring thousands of Americans home in coffins.
   c. Mr. Bush and his aides were leaning toward a military [**conclusion**] of the [*crisis*].

2. **Inchoative Constructions.** Inchoative constructions in English express the coming to existence of a situation. They generally involve the presence of verbs like *become* and *get*, in addition to their complement, which denotes the resulting situation or process. BOTH the inchoative predicate (in bold face) and the complement expressing the resulting situation (in italics) will be annotated as events:

   (41) a. They aren't being allowed to leave and could [**become**] [*hostages*].
   b. The President Ilham Aliyev [**got**] [*acquainted*] with reconstruction works in Vahid garden and National Park.
   c. The public clamor was so great that they [**got**] [*scared*] and a substitute was adopted appointing a committee to investigate the property.

3. **Causative Constructions.** Causal constructions involve one of the following causative predicates, or similar ones, in their causative senses: *cause, stem from, lead to, breed, engender, hatch, induce, occasion, produce, bring about, produce, secure.* Two different constructions can be distinguished here:

   (a) **EVENT**$_{e_1}$ ***cause***$_{e_2}$ **EVENT**$_{e_3}$**.** The causal expression (e_2), its logical subject (e_1) and its event complement (e_3) are ALL tagged as independent events–indicated with square brackets in the example below.

   (42) The [**rains**$_{e_1}$] [**caused**$_{e_2}$] the [**flooding**$_{e_3}$].

   (b) **ENTITY** ***cause***$_{e_1}$ **EVENT**$_{e_2}$**.** BOTH the causal expression (e_1) and its event complement (e_2) are tagged as independent events.

   (43) John [**caused**$_{e_1}$] the [**fire**$_{e_2}$].

4. **Light Verb Constructions.** These involve a verb of very light semantic content (e.g., *make, get, do, have, take, put, set, let*) and a nominal event acting as its selected complement. In these situations, BOTH the verbal and nominal elements are tagged as events. Below, the light verb is in bold face whereas the nominal is in italics.

(44) a. Several pro-Iraq [*demonstrations*] have [**taken**] place in the last week.

b. They will definitely [**take**] into [*consideration*] our readiness.

5. **Copulative Constructions.** Copulative constructions are VPs headed by verbs like *be* or *seem*, and which have an NP (Example (45a)), AP (Example (45b)), or PP (Example (45c)) as complement. In these constructions, BOTH the verbal predicate and the predicative complement will be marked up. All of the involved elements (verb, NPs, APs, or PPs) will be annotated according to the rules specified in the previous sections. In the examples above, the copulative predicate is in bold face, its complement head in italics, and the elements to annotate are delimited with square brackets.

(45) a. An eminent Indian origin woman [**is**] the new [*head*] of the British Medical Association.

b. If, in spite of everything, we will not [**be**] [*ready*], we will ask the United States to delay the operation.

c. Zarei [**was**] [*in*] charge of a program to clean cities from corruption.

The primary reason for annotating both the copula and its complement in TimeML is so that later processing of the document can assign a particular attribute or property to a particular entity. For example, the *woman* in Example (45a) is assigned the property of being the *head of the British Medical Association*. While TimeML itself says nothing about this connection, it is useful to have both events tagged for making this connection at a later time.[8]

6. **Constructions with Functional Nouns.** Examples of functional nouns are: *temperature, size, weight, population, intensity*, etc. They take an individual as argument (denoting, e.g., a person, physical location, group of individuals, etc.) and return a specific value on an appropriate scale, which can be numeric or not (e.g., *high, low; big, small; hot, cold*; etc.). Functional nouns can appear in three different constructions. In all of them, the funtional noun will be marked up as event, together with the main verb of its main clause. Some examples are provided below for different types of constructions, where all the markable expressions are in square brackets. In addition, the functional noun is in bold face and the main verbal predicate of its clause, in italics.

(a) **NOUN**$_{functional}$ ***is*** **X.**

(46) a. The current USA [**population**] [*is*] above 300 million.

b. An appraisal of the house indicated its market [**value**] [*is*] $150000.

(b) **INDIVIDUAL *has* a NOUN**$_{functional}$ **of X.**

(47) a. Catalonia [*has*] a [**population**] of around 7 million people.

b. This noise originated from cosmic radiation and [*had*] a [**temperature**] of 3 K.

(c) **NOUN**$_{functional}$ ***changes/raises/drops/increases/...*** **(from X) (to Y).**

(48) In the nine months, [**net**] [*rose*] 4.3% to $525.8 million, from $504.2 million last year.

As with copulative constructions, constructions with functional nouns have no additional TimeML annotation, but both elements are annotated to allow for easy annotation at a later stage of the functional noun to its given value.

[8]For SpaceEval, do **not** tag copulas with the NONMOTION_EVENT tag.

3.1.2 Event Attributes

When you encounter a spatially relevant event, tagging it with the NONMOTION_EVENT tag will generate an NONMOTION_EVENT ID that can then be used to relate the event to other tag elements. Since events are really the responsibility of TimeML, any inherited attributes are not necessary to discuss here, and are not the responsibility of annotators for this task. The attributes relevant to this task are the optional `elevation` and `mod` attributes, and the `countable` attribute, which are treated the same as for the location tags discussed in Sections 2.1.2, 2.2.2 and 2.3.2.

Note: The `countable` attribute for NONMOTION_EVENT should only be specified for instances of nominal events—for verbal forms, the `countable` attribute should be left unspecified.

3.2 Motions

A motion is a species of event that involves movement. Note that every MOTION tag will participate in a relation with whatever participates in the motion-event. That is to say, in creating a MOTION tag the annotator is also committing to creating at least one MOVELINK for that MOTION. Motion-events receive special attention in this task since they are inherently spatial. Motion-events come in three varieties:

1. Bare-Manner Motion:
   e.g., *John walked.*
2. Path Motion:
   e.g., *John left home.*
3. Compound Motion:
   e.g., *John left home running.* or *John walked home.*

These different strategies for expressing motion in natural language are reflected in the attributes as described in Section 3.2.2.

3.2.1 Motion Extents

When identifying motion-events, follow the same extent rules for any TimeML event. Annotate a MOTION only if it passes the NONMOTION_EVENT tests and it involves a change of location in "real space". E.g., *followed* in *David followed the map* would not be annotated as a MOTION, but in *David followed the road* would it would be.

3.2.2 Motion Attributes

Table 5 shows the attributes for the MOTION tag. The `id` attribute is automatically generated, but the annotator should fill in values for the remaining attributes.

The `motion_type` attribute refers to the distinction mentioned earlier in this section. Manner-of-motion events (those with the `motion_type` value MANNER) are relatively rare in the corpus. In order to receive this value, there can be no indication of the source (starting location), goal (ending location), or mid-point locations of the event-path. PATH and COMPOUND motion-events are more common in the corpus.

MOTION tags of the PATH `motion_type` are those that have an explicit component of the path of motion evident in the text, but that have no indication of the manner in which the motion is performed. The sentences in Example (49) include only PATH type motion-events.

Attribute	Value
id	m1, m2, m3, ...
motion_type	MANNER, PATH, COMPOUND
motion_class	MOVE, MOVE_EXTERNAL, MOVE_INTERNAL, LEAVE, REACH, DETACH, HIT, FOLLOW, DEVIATE, CROSS
motion_sense	LITERAL, FICTIVE, INTRINSIC_CHANGE
mod	a spatially relevant modifier
countable	TRUE or FALSE

Table 5: MOTION Tag Attributes

(49) a. John [**left**$_{m1}$] the room.

b. John [**arrived**$_{m2}$] at the party.

c. John [**left**$_{m3}$].

d. John [**arrived**$_{m4}$].

e. Danielle was [**headed**$_{m5}$] west-northwest at near 17 mph (28 kph).

f. Projections show Danielle [**nearing**$_{m6}$] the Bermuda area by Sunday morning.

Notice that Examples (49c) and (49d) are considered PATH motions, though there are no explicit locations given as the `source` or the `goal`. This is because certain predicates are always interpreted as PATH motion-events even if the PATH information is implicit (e.g., LEAVE class motion-events require a `source` which is PATH information). When the `source`, `goal`, `midPoints`, or `landmark` locations are not made explicit, we naturally figure out what it should be using context. The same can be said for Example (49b) with the `goal` location. You can tell when you are dealing with such a predicate if you find yourself looking for missing information. That is, if you read the sentence *John left*, it is natural to wonder *left where? Leave* and *arrive* will be common PATH motion-events in the corpus, so you should consistently tag them as such, even when the `source`, `goal` etc. are not explicit (and if necessary, non-consuming tags may be created to represent the implicit locations).

MOTION tags of the MANNER `motion_type` capture what are known as "bare-manner" motion-events, which are a rarer type of motion-event in English. These are motion-events where no explicit path of motion is evident, but the manner of motion is indicated. The sentences in Example (50) exemplify MOTION tags of the MANNER `motion_type`.

(50) a. John [**ran**$_{m1}$] five miles yesterday.

b. John [**bikes**$_{m2}$] seriously.

c. The arrow [**flew**$_{m3}$] straight and true.

d. John [**took**$_{m4}$] the bus.

Note: Light verbs may express motion-events, though they do not carry any path or manner information by themselves. In Example (50d), for instance, *took* introduces a motion-event, but contributes no path or manner information by itself. In this case, the presence of the motion

signal[9], *the bus*, indicates manner, which would qualify this MOTION tag to be annotated with the MANNER `motion_type`.

The most common value for the `motion_type` attribute is COMPOUND. A COMPOUND motion-event has characteristics of both PATH and MANNER motions. Sometimes manner will be encoded in the verb itself while path information will appear in a motion signal. However, some motion verbs conflate path and manner without any separate motion signals. Bare-manner motion verbs can also coincide with a motion signal that encodes path information. And, in still another case, multiple motion signals may provide both path and manner information about a singular motion-event. The sentences in Example (51) provide examples of MOTION tags with the COMPOUND `motion_type`.

(51) a. John [**biked**$_{m1}$] from Virginia to Oregon.

b. John [**left**$_{m2}$] the concert on foot.

c. John [**went**$_{m3}$] through the tunnel on his bike.

d. John [**caught**$_{m4}$] a taxi home.

The values for the `motion_class` attribute are each associated with a spatial event structure that specifies the spatial relations between the arguments of the motion at different phases of the event. For example, a REACH motion such as *arrive* involves a pre-state in which the mover is not at the `goal` location and a post-state in which the mover is at the `goal` location. Table 6 lists event structures associated with the different `motion_class` values. To determine the appropriate `motion_class` value, annotators must identify which event structure the event-path resembles.

Note: Annotators should not necessarily annotate every instance of the same motion verb with the same `motion_class` value; the class of motion for a particular motion event may depend on the motion verb in conjunction with additional information about the event-path that is contributed by context (e.g., PATH type MOTION_SIGNAL tags). For instance, the MOTION tag for *biked* in, *John biked*, without additional specification, would be classified as MOVE, however, in, *John biked across town*, *biked* would be classified as CROSS due to the additional path of motion information conferred by the prepositional phrase *across town*.

If a MOTION tag's `motion_class` attribute is annotated as MOVE, this indicates that the event structure is unclear or underspecified. All that is required for the MOVE class is that there is some event-path that is introduced. The MOVE class, as such, could be considered a base-case, and the event structures of all other motion classes are more specific. For instance, annotating a MOTION with `motion_class` MOVE_EXTERNAL stipulates that at every phase of the event the `mover` and `ground` are disconnected or externally connected.

For the rest of the event structure descriptions in this section, we will adopt the following notation convention. Event-paths are represented using a comma separated tuples denoting spatial relations between the `mover` and some point along the event-path. The point along the path that is the "focus point", i.e., the salient location with respect to which the motion is framed, is dependent on the motion class. E.g., while the salient point for the LEAVE class is the `source`, the focus for the REACH class is the `goal`. These tuple elements will consist of RCC8^{+}[10] relations. The order of the elements in the tuples represent a temporal ordering for the event structure; the first element

[9]Refer to Section 3.3 for further details regarding MOTION_SIGNAL tags

[10]Refer to Table 11 in Section 5.1 for further details regarding the RCC8^{+}relation types.

Class	Path Focus	Event Structure
MOVE	undefined	undefined
MOVE_EXTERNAL	`ground`	(DC, DC, DC) or (EC, EC, EC)
MOVE_INTERNAL	`ground`	(IN, IN, IN)
LEAVE	`source`	$(NTPP, (TPP, PO, EC), DC)$ or $(TPP, (PO, EC), DC)$ or (EQ, EC, DC)
REACH	`goal`	$(DC, (EC, PO, TPP), NTPP)$ or $(DC, (EC, PO), TPP)$ or (DC, EC, EQ)
DETACH	`source`	(PO, EC, DC) or (EC, DC)
HIT	`goal`	(DC, EC, PO) or (DC, EC)
CROSS	`midPoint`	$(DC, (EC, (PO, [IN\|EQ], PO), EC), DC)$ or $(EC, (PO, [IN\|EQ], PO), EC)$ or $(PO, [IN\|EQ], PO)$ or $(TPP, NTPP, TPP)$
FOLLOW	`pathID`	$(IN, NTPP, IN)$
DEVIATE	`pathID`	(IN, EC, DC)

Table 6: Motion Class Event Structures

describes some pre-state at the beginning or start of the event-path and the last element describes a post-state at the end of the event-path. If the tuple contains intermediate elements between the first and last, those elements describe a state (or series of states) at some intermediate point(s) on the event-path. In some cases, the tuple elements may be represented by a disjunction of a number of RCC8$^+$relation types within square brackets, with | denoting logical disjunction. Additionally, there are some event structures whose elements consist of a complex sub-event, which is also represented as a comma separated tuple of RCC8$^+$relations.

The way to read the event structure representations in Table 6, such as for the LEAVE motion class, $([IN|EQ], (PO, EC), DC)$, would be as follows. At the beginning of the event-path the mover is either inside the ground or occupies the same space as the ground. Then there is a sub-event where the mover is first partially-overlaps the ground and subsequently is externally connected to the ground. Finally, at the end of the path, the mover is disconnected from the ground.

Note: It is important to point out that these event structures are described in terms of the RCC8$^+$relations whose arguments are spatial objects themselves. The FOLLOW and DEVIATE classes of motion require that the path-focus is a PATH. For all other classes of motion, the path-focus argument must be coerced to a two-dimensional region in order to interpret the event structure frames.

Note: Additionally, there is a constraint on the CROSS motion class in that this class requires the event-path to traverse across the `midPoint` such that the `source` and `goal` locations along the event-path would satisfy an OLINK orienting the `goal` across from the `source` relative to the `midPoint`. That is, an event-path that simply overlaps some region, but returns to the `source`, or otherwise fails to traverse across a bisectional axis through the `midPoint` that separates the `source` and `goal`, would not be classified as a CROSS motion.

Finally, the `motion_sense` attribute distinguishes between different kinds of interpretations

of motion-events. The `literal` sense covers motion verbs that describe dynamic motion-events involving a mover whose location changes over time and space. The FICTIVE sense covers atemporal motion-events, i.e., events where the `mover` object introduces a static-path. The INTRINSIC_CHANGE sense attribute covers motion-events that involve temporal or dynamic change in the intrinsic spatial structure or spatial configuration of an object over space. Table 7 lists some examples of each of the senses of motion.

Motion Sense Value	Examples
LITERAL	*John biked*, *the ball rolled*, *the balloon rose*
FICTIVE	*the river ran*, *the road climbed*, *the mountains rose*
INTRINSIC_CHANGE	*the glacier receded*, *the river rose*, *the balloon expanded*

Table 7: `motion_sense` Attribute Values

3.3 Motion Signal

A motion signal is taken to be a particle, preposition, verb, or adverb that encodes path or manner information about a motion-event. The MOTION_SIGNAL tag specifies the kind of information that the term contributes: either path-of-motion or manner-of-motion information. The MOTION_SIGNAL type PATH indicates path-of-motion information, such as the underlined terms in the following motion expressions:

(52) a. John went <u>to</u> the store.

b. John departed <u>from</u> the airport.

c. John swam <u>around</u> the pool.

d. John crawled <u>into</u> bed.

MOTION_SIGNAL tags of type MANNER contribute manner-of-motion information:

(53) a. John left the garage <u>by car</u>.

b. John arrived <u>on foot</u>.

c. John rode his <u>bike</u>.

d. John turned the corner <u>sprinting</u>.

3.3.1 Motion Signal Extents

When a PATH type MOTION occurs in conjunction with a MOTION_SIGNAL contributing path-of-motion information, the extent of the MOTION_SIGNAL tag should be limited to the preposition, particle, or other individual term alone. Example (54) illustrates a sentence where the entire extent of the PP *to the store* is not tagged—only *to*.

(54) John walked [**to**$_{ms1}$] the store.

When a MANNER or COMPOUND type MOTION appears in conjunction with a prepositional MOTION_SIGNAL that supplies manner-of-motion information, as in Example (55), the MOTION_SIGNAL tag extent should subsume the entire prepositional or adverbial phrase.

(55) John left the garage [**by car**$_{ms2}$].

Additionally, in certain situations, verbs such as *start*, *begin*, *initiate*, *continue*, *end*, *finish*, *complete*, and *conclude*, which typically express grammatical aspect of subordinate verbs—i.e., how actions relates to the flow of time—may be captured with the MOTION_SIGNAL tag. In many instances these verbs express the temporal aspect of a NONMOTION_EVENT (e.g., *The party* ***begins*** *at 7pm.*), in which case they are not spatially relevant and would not be tagged. In other instances these verbs may introduce fictive motion (e.g., *The river* ***begins*** *at the lake.*). However, in situations where a nominal MOTION occurs as a complement of these aspectual verbs, it is appropriate to tag the aspectual verb as a PATH motion signal to capture the path-of-motion information it contributes. Example (56) illustrates these sorts of MOTION_SIGNAL tags, which indicate points along event-paths associated with nominal motion.

(56) a. The [**trek**$_{m1}$] to the peak [**starts**$_{ms1}$] at the base camp.

b. The museum [**tour**$_{m2}$] [**concluded**$_{ms2}$] at the gift shop.

c. John's [**run**$_{m3}$] [**continued**$_{ms3}$] through the tunnel.

3.3.2 Motion Signal Attributes

Motion Signal have only two relevant attributes:

Attribute	Value
`id`	`a1`, `a2`, `a3`, ...
`motion_signal_type`	MANNER, PATH

Table 8: Attributes for MOTION_SIGNAL

Annotators must choose whether the `motion_signal_type` is PATH or MANNER:

1. MANNER: Used when the MOTION_SIGNAL supplies information about the MANNER of the MOTION it is associated with

2. PATH: Used when the MOTION_SIGNAL supplies information about the PATH of the MOTION it is associated with

Example (57) shows the attribute values for the above examples.

(57) a. John walked [**to**$_{ms1}$] the store.
MOTION_SIGNAL (`id`=`a1`, `motion_signal_type`=PATH)

b. John left the garage [**by car**$_{ms2}$].
MOTION_SIGNAL (`id`=`a2`, `motion_signal_type`=MANNER)

c. John arrived [**in**$_{ms3}$] Boston.
MOTION_SIGNAL (`id`=`a3`, `motion_signal_type`=PATH)

It may not always be obvious when a preposition is acting as a MOTION_SIGNAL or a SPATIAL_SIGNAL. Annotators should consider what semantic information the preposition is contributing. In Example (57c), *in* is acting as a motion signal in this context, although the preposition *in* may act as a SPATIAL_SIGNAL in other contexts. In this instance, *in* is supplying path-of-motion information about the motion-event—namely *arrived*—by indicating the `goal` component of the event-path—namely *Boston*.

4 Spatial Signals and Measurements

4.1 Spatial Signals

A SPATIAL_SIGNAL is a word that supplies information to a spatial link. For example, the spatial signals are highlighted in each of the sentences in Example (58).

(58) a. The cup is [**on**$_{s1}$] the table.

b. Boston is [**north of**$_{s2}$] New York.

c. Danielle was headed [**west-northwest**$_{s3}$] at near 17 mph (28 kph).

d. The new skyscraper [**at**$_{s4}$] 111 Huntington Avenue was completed in 2002, [**directly across**$_{s5}$] the street from The Colonnade Hotel.

In general, spatial signals are prepositions or prepopistional phrases that reveal the particular relationship between two location tag elements, thereby helping the annotator decide what kind of links should be used and what the values for attributes in those links should be. Recall that MOTION_SIGNAL tags and SPATIAL_SIGNAL tags have different functions: `spatial_signal` tags always supply information about topological or qualitative spatial relations between other elements, and MOTION_SIGNAL tags capture information specifically about the path or manner of a motion-event.

4.1.1 Spatial Signal Extents

The extents for spatial signals are usually one word prepositions and are generally easy to spot. Example (59) illustrates some markable extents captured with the SPATIAL_SIGNAL tag type.

4.1.2 Spatial Signal Attributes

Spatial signals have three attributes associated with them, as shown in Table 9.

Attribute	Value
`id`	`s1`, `s2`, `s3`, ...
`cluster`	identifies the sense of the preposition
`semantic_type`	DIRECTIONAL, TOPOLOGICAL, DIR_TOP

Table 9: Attributes for SPATIAL_SIGNAL

Of these attributes, only `semantic_type` must be filled in by the annotator for this annotation task. The `semantic_type` refers to what kinds of links are introduced by the spatial signal. This attribute has three possible values as follows:

1. DIRECTIONAL: Introduces an OLINK (refer to Section 5.2)
2. TOPOLOGICAL: Introduces a QSLINK (refer to Section 5.1)
3. DIR_TOP: Introduces both a QSLINK and an OLINK

Example (59) shows the attribute values for the above examples.

(59) a. The cup is [**on**$_{s1}$] the table.
SPATIAL_SIGNAL(`id=s1`, `semantic_type`=DIR_TOP)

b. Boston is [**north of**$_{s2}$] New York.
SPATIAL_SIGNAL(`id=s2`, `semantic_type`=DIRECTIONAL)

c. Danielle was headed [**west-northwest**$_{s3}$] at near 17 mph (28 kph).
SPATIAL_SIGNAL(`id=s3`, `semantic_type`=DIRECTIONAL)

d. The new skyscraper at 111 Huntington Avenue was completed in 2002, [**directly across**$_{s4}$] the street from The Colonnade Hotel.
SPATIAL_SIGNAL(`id=s4`, `semantic_type`=DIR_TOP)

The remaining attribute, `cluster`, will not be used for this Pilot Annotation. This attribute refers to the sense of the spatial signal as it appears in a sense inventory. It is expected that the signal's sense will indicate what link tag types are introduced by the signal. So, in the future, the annotator won't have to fill in the `semantic_type` attribute if he or she knows the sense number for the signal. For now, though, annotators should ignore this attribute.

4.2 Measures

A MEASURE is a special kind of spatial signal that captures distances and dimensions and introduces a measure link (i.e., an MLINK[11]). MEASURE tags consist of a numerical component and a unit component as shown in Example (60a) through Example (60d), or consist of a relative measurement term such as in Example (60e).

4.2.1 Measure Extents

The extent for the MEASURE tag includes the numerical component and the unit component. The sentences in Example (60) each contain a MEASURE tag.

(60) a. John walked for [**5 miles**$_{me1}$].

b. The field is [**100 yards**$_{me2}$] long.

c. Danielle's center was about [**710 miles**$_{me3}$] ([**1,145 kilometers**$_{me4}$]) east of the northern Leeward Islands.

d. At a mere [**25 stories**$_{me5}$], it is overshadowed by the other two.

e. Arriving in the town of Juanjui, [**near**$_{me6}$] the park, I learned ...

[11] Refer to Section 5.4 for more details regarding MLINK tags.

4.2.2 Measure Attributes

The attributes for the MEASURE tag are fairly straightforward as shown in Table 10. The `value` attribute should have a numerical value for the numerical component of the MEASURE. The unit of measurement should be stored in the `unit` attribute, as shown in Example (61).

Note: There are exceptional cases where distances are described in relative terms. In Example (60e), for instance, *near* has been tagged as a MEASURE, though its `unit` attribute remains unspecified. Other relative spatial terms, such as *close* or *far*, may also act in this fashion, though they are also capable of acting as spatial modifiers that would fill a `mod` attribute for a location tag (e.g., the underlined adjectives in, *the near side of the lake* or, *the far mountains*).

Attribute	Value
`id`	`me1`, `me2`, `me3`, ...
`value`	number component
`unit`	measurement phrase component

Table 10: Attributes for MEASURE

(61) a. John walked for [**5 miles**$_{me1}$].
MEASURE (`id`=`me1`, `value`="5", `unit`="miles")

b. The field is [**100 yards**$_{me2}$] long.
MEASURE (`id`=`me2`, `value`="100", `unit`="yards")

c. Danielle's center was about [**710 miles**$_{me3}$] ([**1,145 kilometers**$_{me4}$]) east of the northern Leeward Islands.
MEASURE (`id`=`me3`, `value`="710", `unit`="miles")
MEASURE (`id`=`me4`, `value`="1145", `unit`="kilometers")

d. At a mere [**25 stories**$_{me5}$], it is overshadowed by the other two.
MEASURE (`id`=`me5`, `value`="25", `unit`="stories")

e. The city has sunk [**6 meters**$_{me6}$] over the past decade.
MEASURE (`id`=`me6`, `value`="6", `unit`="meters")

f. The hot dog stand [**near**$_{me7}$] Macy's.
MEASURE (`id`=`me7`, `value`="NEAR", `unit`=∅)

Note in Example (61e) that the MEASURE tag `value` attribute is not negative. In such cases the directionality is contributed by the motion verb *sunk*, not the MEASURE tag. The value of the MEASURE tag measures a dimension of the event-path of the sinking event, which cannot be negative. Note, however that the `value` attribute for MEASURE tags may take a negative value when identifying elevations when specified as an offset on some scale (e.g., *500 ft below sea level* would necessitate a `value` of "-500"). Annotators should not fill separating commas (or other extraneous notation) in attributes that hold numerical values. E.g., in Example (61c), the `value` for `me4` is "1145" not "1,145". For non-integer numerical values, use decimal notation, not fractions (e.g., "0.5" not "1/2").

5 Spatial Relationships

Thus far, all of the tags that have been discussed, with the exception of METALINK tags, have involved tagging some spatially relevant span of text. The remainder of the tag types capture information about spatial relationships between those tagged elements. There are four link tag types (not counting METALINK, which is not spatial in nature). The link tags are:

1. QSLINK — qualitative spatial links;
2. OLINK — orientation information;
3. MOVELINK — movement links;
4. MLINK — defining the dimensions of a location.

Each of these links captures unique information about the relationships shared between spatial objects. Note that link tags have no extents themselves. Links typically hold the IDs of two spatial objects, the IDs of any other tags that supply further information to the link, and some additional attributes for describing the nature of the relationship between the objects mentioned in the link. In a way, the tags discussed in previous sections in this document can be thought of as "ingredients" for creating these links.

The remainder of this section describes each of the four link tag types in detail. In addition, the examples in this section are more complete so they should provide additional information for the extent tags as well.

5.1 Qualitative Spatial Links

A qualitative spatial link captures the topological relationship between two spatial objects. For this reason, they are triggered by SPATIAL_SIGNAL tags with a `semantic_type` of either TOPOLOGICAL or DIR_TOP. Topological information primarily refers to containment and connection relations between two regions. The possible relationships come from a field of research called Qualitative Spatial Reasoning (QSR), which primarily deals with how abstract objects relate. Since most of the spatial objects that are mentioned in natural language text are not abstract, however, QSR is generally insufficient for fully capturing the intended relationship between the objects. For that reason, both QSLINK and OLINK tags may be required to fully capture spatial relationships.

For example, consider the sentence: *The cup is on the table.* The SPATIAL_SIGNAL *on* in this sentence tells us that the cup is in direct contact with the table. This is **topological** information. However, a simple "direct contact" relationship does not say whether the cup is sitting on top of the table (the likely intended relationship) or if it is somehow clinging to the side of or hanging from beneath the table (not likely, but possible). To capture this aspect of the relationship, an OLINK is required. This is discussed in Section 5.2. For now, though, let us focus on qualitative spatial relation (QSR) based relationships.

SpaceEvaluses the Region Connection Calculus (RCC) as the basis for its qualitative spatial relationships. RCC is concerned with how regions (spatial objects) are *connected* to each other. RCC8$^+$, a variant of RCC8, which consists of 8 basic relations, is used as a basis for the possible relationships between regions. The RCC8 along with `in` and `out` will be referred to as RCC8$^+$. Table 11 defines the different relationships that RCC8$^+$captures and Figure 1 shows an abstract example of the RCC8 relations.

The objects participating in a spatial relation with one another are referred to by different labels, but for SpaceEval, the argument labels are identified as `trajector` or `landmark`. The `trajector` is

Value	Description
DC	disconnected
EC	externally connected
PO	partial overlap
EQ	equal
TPP	tangential proper part
TPPi	tangential proper part inverse
NTPP	non-tangential proper part
NTPPi	non-tangential proper part inverse
IN	disjunction of TTP and NTTP
OUT	disjunction of EC and DC

Table 11: RCC8$^+$Relations

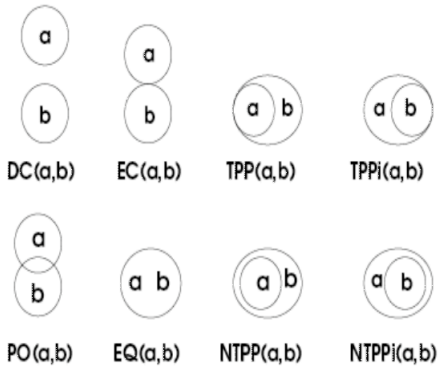

Figure 1: RCC8 Relations

the object *being related* to the `landmark` while the `landmark` is what the `trajector` is *being related to*. It is not a universal rule, but, often, the `trajector` is a movable object while the `landmark` tends to be more static. In the cup and table example above, the cup is the `trajector` while the table is the `landmark`. The next section includes several examples that should help clarify this distinction.

5.1.1 Qualitative Spatial Link Attributes

Table 12 shows the attributes for the QSLINK tag. As usual, the `id` attribute is assigned automatically, but the annotator must fill in the `trajector`, `landmark`, `trigger`, and `relType` values.

Both `trajector` and `landmark` can hold the ID of a PLACE, PATH, SPATIAL_ENTITY, NONMOTION_EVENT or MOTION tag. When an entity that is not a PLACE participates in a QSLINK, it is actually being coerced into behaving like a location. That is, rather than saying that a spatial entity is in some relationship to another object, a QSLINK really signifies that the region that the spatial entity occupies is related to the region referenced by another object. Remember that the `trajector` is the object *being related* and the `landmark` is the object that the `trajector` is *being related to*.

The `trigger` value takes the ID of SPATIAL_SIGNAL with a `semantic_type` of TOPOLOGICAL or DIR_TOP. Keep in mind that signals of this type always introduce a QSLINK, but the `trigger` attribute is optional because it is possible to have a QSLINK that is not associated with any particular

Attribute	Value
id	qsl1, qsl2, qsl3, ...
relType	DC, EC, PO, EQ, TPP, TPPi, NTPP, NTPPi, IN, OUT
trajector	ID of location/entity/event tag that is being related
landmark	ID of the location/entity/event tag that is being related to
trigger	ID of the spatial signal that triggered the link

Table 12: Attributes for QSLINK

SPATIAL_SIGNAL in the text.

The relType attribute is used to specify the type of qualitative spatial relationship that exists between the trajector and the ground. The relType attribute takes as its value any of the RCC8 relations as well as the additional value IN, which is the disjunction of TPP and NTPP. A relType of IN should be used when it is not clear whether TPP or NTPP is the correct relType, but one of those two values must apply. The value EQ is special in that it is used to say that two spatial objects occupy the same space, or in other words, that they have the same location. Table 13 displays the possible relType values with some natural language examples with trajector objects marked as $[trajector_t]$ and ground objects as $[landmark_l]$.

Note: The EQ relType is not used to indicate that two spatial objects are actually identical. I.e., in cases where the same spatial object is referenced multiple times in the text via anaphora, has its name mentioned multiple times, these would be cases of coreference. [12].

Value	Example
DC	the $[grill_t]$ outside of the $[house_l]$
EC	the $[cup_t]$ on the $[table_l]$
PO	$[Russia_{t/l}]$ and $[Asia_{t/l}]$
EQ	the $[The\ White\ House_{t/l}]$ and $[1600\ Pennsylvania\ Avenue_{t/l}]$
TPP	the $[coast_t]$ of $[Delaware_l]$
TPPi	
NTPP	the $[island_t]$ in the $[lake_l]$
NTPPi	
IN	the $[bookcase_t]$ in the $[room_l]$

Table 13: QSLINK relType values

The sentences in Example (62) demonstrate QSLINK tags. For additional examples, refer to Section 8.

(62) a. The [**book**$_{se1}$] is [**on**$_{s1}$] the [**table**$_{se2}$].
SPATIAL_SIGNAL (id=s1, cluster="on-1", semantic_type=DIR_TOP)
QSLINK (id=qsl1, trajector=se1, landmark=se2, trigger=s1, relType=EC)

[12]To indicate that two spatial objects mentioned in the text are actually identical the METALINK tag is used. Refer to Section 2.5 for further details.

b. The light [**switch**$_{se3}$] is [**on**$_{s2}$] the [**wall**$_{se4}$].
SPATIAL_SIGNAL (id=s1, cluster="on-2", semantic_type=DIR_TOP)
QSLINK (id=qsl2, trajector=se3, landmark=se4, trigger=s2, relType=PO)

c. A thick green [**rainforest**$_{pl2}$] grew up [**around**$_{s4}$] the [**road**$_{p1}$].
SPATIAL_SIGNAL (id=s4, semantic_type=TOPOLOGICAL)
QSLINK (id=qsl4, trajector=p1, landmark=pl2, trigger=s4, relType=IN)

Notice that while the same spatial signal is used in both of these examples, the relType value for each differs. This is because the signal *on* is being used in a slightly different sense in each of the examples. It is also noteworthy here that the semantic_type for these examples dictates that an OLINK be supplied in addition to these QSLINK tags. OLINK tags are described in Section 5.2.

5.2 Orientation Link

The OLINK tag covers those relationships that occur between two locations that are non-topological in nature. Orientation links essentially fill in the spatial relations that QSLINK tags cannot capture. This includes three different types of information based on the three frames of reference as follows:

1. **Absolute**: This frame of reference can be considered the "bird's eye" view.
2. **Intrinsic**: This frame of reference is used when some part of a spatial object has an intrinsic orientation such as a TV, which has an intrinsic front.
3. **Relative**: This frame of reference is used when the relationship being described depends on a particular entity's point of view.

Once the frame of reference for the OLINK has been identified, the annotator must also supply a reference point. For ABSOLUTE OLINK tags, the referencePt attribute value must match the relType value which must be a cardinal direction. For INTRINSIC OLINK tags, the referencePt takes the ID filled for the landmark attribute. For RELATIVE OLINK tags, the referencePt is either the ID of a spatial entity from whose viewpoint the relation is described, or the special value "VIEWER", which is used to indicate that the author did not explicitly declare who is viewing the relationship, but it is still a relative frame of reference (e.g., *the table on David's left* vs. *the table on the left*).

OLINK tags also capture projective information. Consider the sentences in Example (63).

(63) a. The [**helicopter**$_{se1}$] is [**above**$_{s1}$] the [**town**$_{pl1}$].

b. The [**hill**$_{pl2}$] is [**above**$_{s2}$] the [**town**$_{pl3}$].

c. The [**city of Boston**$_{pl4}$] is [**north of**$_{s3}$] [**Stoughton, MA**$_{pl5}$].

d. The [**city of Boston**$_{pl6}$] is [**north of**$_{s4}$] [**New York City**$_{pl7}$].

Both Examples (63a) and (63b) use the same SPATIAL_SIGNAL word, *above*. However, in Example (63a), the likely interpretation is that the *helicopter* is located directly above the *town*. This is not the most salient interpretation for Example (63b); hills usually do not fly or hover above towns in the same way that the helicopters do. To distinguish between these two interpretations, we say that the OLINK in Example (63b) has a projective interpretation in which we imagine that the region associated with the town projects outwards beyond its normal limits. It is this projected region,

which associated with the town, that the hill is actually located above. So both of these sentences should have nearly identical OLINK tags created for them, except that the `projective` attribute value for the link for Example (63b) would be flagged as TRUE and FALSE for Example (63a).

The issue of projectivity may also arise for OLINK tags involving any of the four cardinal directions. In Example (63c) the relation between *Boston* and *Stoughton* would not be projective because Boston is directly north of *Stoughton*. The relation in Example (63d), however, would be projective because *Boston* is indirectly north of *New York City*. The actual relationship in Example (63d) could be described with a `relType` of "NORTHEAST", although this information isn't directly accessible from the language—that kind of world-knowledge would have to be looked up from a gazetteer entry which is not part of the pilot annotation task. Example (64) includes annotations that illustrate this distinction in further detail.

5.2.1 Orientation Link Attributes

Table 14 shows the attributes for the OLINK tag.

Attribute	Value
`id`	`ol1`, `ol2`, `ol3`, ...
`relType`	ABOVE, BEHIND, NEXT TO, NORTH OF, ...
`trajector`	ID of the location/entity/event tag that is being related
`landmark`	ID of the location/entity/event tag that is being related to
`trigger`	ID of the spatial signal that triggered the link
`frame_type`	ABSOLUTE, INTRINSIC, RELATIVE
`referencePt`	cardinal direction, ground entity, viewer entity
`projective`	TRUE, FALSE

Table 14: Attributes for OLINK

As with QSLINK, the `trajector` and `landmark` attributes can hold the ID of any location/entity/event tag. The `trigger`, which is optional, must be a SPATIAL_SIGNAL with a `semantic_type` of either DIRECTIONAL or DIR_TOP. The `projective` attribute can have a value of either TRUE for projective interpretations or FALSE for non-projective cases. The `relType` attribute currently has an open set of values, some of which are named in Table 14. Annotators should try to stick to this set of values, but may annotate additional values if none of those are appropriate.

Perhaps more so than any other tag, the attributes of the OLINK tag are dependent on each other. The value for the `frame_type` attribute determines what the `referencePt` value should be, thus the `frame_type` should be filled first. Table 15 shows the consequences for each `frame_type` value.

Frame Type Value	Effect
ABSOLUTE	`referencePt`=value of `relType`
INTRINSIC	`referencePt`=value of `ground`
RELATIVE	`referencePt`=VIEWER or tag ID specifying the viewpoint

Table 15: Impact of `frame_type` Values on `referencePt`

Example (64) shows several different kinds of OLINK tags. Once again, only the tag in question is shown in these annotations though many of them also have accompanying QSLINK tags.

(64) a. [**Boston**$_{pl1}$] is [**north of**$_{s1}$] [**New York City**$_{pl2}$].
OLINK (id=ol1, trajector=pl1, ground=pl2, trigger=s1, relType="NORTH", frame_type=ABSOLUTE, referencePt=NORTH, projective=TRUE)

b. The [**dog**$_{se1}$] is [**in front of**$_{s2}$] the [**couch**$_{se2}$].
OLINK (id=ol2, trajector=se1, landmark=se2, trigger=s2, relType="FRONT", frame_type=INTRINSIC, referencePt=se2, projective=FALSE)

c. The [**dog**$_{se3}$] is [**next to**$_{s3}$] the [**tree**$_{se4}$].
OLINK (id=ol3, trajector=se3, landmark=se4, trigger=s3, relType="NEXT TO", frame_type=RELATIVE, referencePt=VIEWER, projective=FALSE)

d. The [**hill**$_{pl3}$] is [**above**$_{s4}$] the [**town**$_{pl4}$].
OLINK (id=ol4, trajector=pl3, landmark=pl4, trigger=s4, relType="ABOVE", frame_type=INTRINSIC, referencePt=pl4, projective=TRUE)

e. The [**helicopter**$_{se5}$] is [**above**$_{s5}$] the [**town**$_{pl6}$].
OLINK (id=ol5, trajector=se5, landmark=pl6, trigger=s5, relType="ABOVE", frame_type=INTRINSIC, referencePt=pl6, projective=FALSE)

f. The [**book**$_{se1}$] is [**on**$_{s1}$] the [**table**$_{se2}$].
OLINK (id=ol1, trajector=se1, landmark=se2, trigger=s1, relType="ABOVE", frame_type=INTRINSIC, referencePt=se2, projective=FALSE)

g. The [**gum**$_{se1}$] is [**on**$_{s1}$] the [**table**$_{se2}$].
OLINK (id=ol1, trajector=se1, landmark=se2, trigger=s1, relType="BELOW", frame_type=INTRINSIC, referencePt=se2, projective=FALSE)

h. The new tropical [**depression**$_{se8}$] was about 430 miles (690 kilometers) [**west**$_{s4}$] of the southernmost [**Cape VerdeIslands**$_{pl8}$]
OLINK (id=ol8, trajector=se8, landmark=pl8, trigger=s4, relType="WEST", frame_type=ABSOLUTE, referencePt="WEST", projective=TRUE)

5.3 Movement Links

The MOVELINK tag connects motion-events with mover-participants. The other attributes of the MOVELINK tag are then used to specify any evident information about components of the event-path as well as any motion signals. MOVELINK tags are always introduced by a triggering MOTION tag. Therefore, whenever an annotator tags an extent with the MOTION tag, they are committing to also creating a corresponding MOVELINK. The annotation for the MOVELINK depends on the `motion_type` of the MOTION (i.e., MANNER, PATH, or COMPOUND). A bare-manner motion verb (e.g., *David cycles seriously*) still triggers a MOVELINK, though most of the attributes will be underspecified since there is no evident event-path information. At the other extereme, it's possible for PATH or COMPOUND type motions to make use of the full range of MOVELINK attributes.

5.3.1 Movement Link Attributes

Table 16 shows the attributes for the MOVELINK tag.

The `trigger` value of a MOVELINK is filled by the MOTION tag ID of the motion which is being linked to the `mover` participant.

Attribute	Value
id	mvl1, mvl2, mvl3, ...
trigger	ID of a MOTION that triggered the link
source	ID of a location/entity/event tag at the beginning of the event-path
goal	ID of a location/entity/event tag at the end of the event-path
midPoint	ID(s) of event-path midpoint location/entity/event tags
mover	ID of the locatin/entity/event tag whose location changes
ground	ID of a location/entity/event tag that the mover participant's motion is relative to
goal_reached	TRUE, FALSE, UNCERTAIN
pathID	ID of a PATH tag that is identical to the event-path of the trigger MOTION
motion_signalID	ID(s) of (an) MOTION_SIGNAL tag(s) that contributes path or manner information to the trigger MOTION

Table 16: MOVELINK Tag Attributes.

The source, goal, midPoint, and ground attributes are used when the trigger is a PATH or COMPOUND type MOTION. Motions of these types always include some information about the path traversed by the mover (i.e., the event-path). This information is stored in the MOVELINK's source, goal, midPoint, and landmark attributes. The values for these attributes may be filled by any location tag or any tag which can be coerced to act as a location including PLACE, PATH, SPATIAL_ENTITY, NONMOTION_EVENT, and MOTION tags, though they will most often be filled with IDs of PLACE or PATH tags. When creating MOVELINK tags, don't look across sentence boundaries to find source, goal, ground or other event-path information (allow post-processing to fill in that kind of information).

The mover attribute specifies the tag element which participates in the MOTION event by changing location. The mover attribute usually takes an ID of a SPATIAL_ENTITY, though it may also be filled by a location tag or event tag coerced to a location.

The goal_reached attribute, which can have a value of TRUE, FALSE, or UNCERTAIN, is used for those cases when it is not clear from the text whether the identified goal location was reached. If there is no goal location associated with the event then this attribute will be left unspecified. For instance, in *John arrived in Boston*, the goal_reached attribute would be set to TRUE. To take another example, in *John left for Boston*, *Boston* appears to be the goal of the MOTION, but the reader does not know if *John* ever really got there. In such a case, the goal_reached attribute should be set to UNCERTAIN. Marking goal_reached as UNCERTAIN stipulates that the annotator is unsure of *John*'s location within the narrative after the *left* MOTION event has occurred. Note that it does not fall within the purview of the SpaceEvaltask to capture negation or modality. In the case of *John* ***did not*** *arrive in Boston*, the negation does not affect the goal_reached attribute value, which should be set to TRUE. Even though, within the spatial narrative, it is true that *John*'s location would not be *Boston* after the motion-event had completed, if the verb phrase headed by *arrive* were not negated, the goal location of *Boston* would have been reached. Contrastively, in *John approached Boston*, goal_reached would be filled as FALSE because upon completing of the approach motion-event, *John* would still not have reached *Boston*.

The motion_signalID attribute takes the IDs of any MOTION_SIGNAL tags contributing to the manner-of-motion or the event-path of the triggering MOTION. motion_signalID is an optional

attribute because not all motion verbs are accompanied by motion signals. For example, in *John traveled by car*, the phrase *by car* is a motion AJDUNCT, but for *John traveled for three days*, there is no motion AJDUNCT.

Depending on the `motion_class` of the trigger MOTION, certain attributes of the MOVELINK tag, which define the event-path, will be required. E.g., in Example (65e), the `motion_class` for the MOTION *jump* is MOVE_EXTERNAL. This `motion_class` requires that the `landmark` attribute is filled by the identifier of the PATH tag for *fence* to capture the fact that the location of the *fence* is what *John jumped* relative to. The only `motion_class` that may remain totally underspecified is the MOVE class, although it is not obligated to be underspecified. I.e., a MOVE class motion may have a `source`, `goal`, `midPoint` or other attributes specified, but only if the motion-event structure does not fit any of the more specific classes. Table 17 lists which MOVELINK attributes are requisite for each of the different classes of MOTION.

MOVELINK tags triggered by FOLLOW class MOTION tags require the `pathID` attribute to be filled, such as in Example (65c). In such a case, the `pathID` attribute for MOVELINK is used to link the PATH that is traversed by the `mover` to the MOTION to specify an explicit event-path. Essentially this specifies that the PATH *Massachusetts Turnpike* is identical to the event-path for the *drove* MOTION. Note, however, that there may also be information about the event-path supplied by way of the `source`, `midPoint`, and `goal` or `landmark` attributes (in this instance, the GOAL happens to be *Worcester*).

Motion Class	**Required Attributes**
MOVE	None
MOVE_EXTERNAL	`landmark`
MOVE_INTERNAL	`landmark`
LEAVE	`source`
REACH	`goal`
DETACH	`source`
HIT	`goal`
FOLLOW	`pathID`
DEVIATE	`source`
CROSS	`source`,`midPoint`,`goal`

Table 17: MOVELINK Attributes Required by Classes of Motion

Example (65) illustrates how to annotate MOVELINK tags. Since the MOTION tag that triggers a movement link informs the MOVELINK tag's attributes, the MOTION tags are also included in the example annotations.

(65) a. [**John**$_{se1}$] [**walked**$_{m1}$] [**from**$_{ms1}$] [**Boston**$_{pl1}$] [**to**$_{ms2}$] [**Cambridge**$_{pl2}$].
MOTION (`id`=m1, `motion_type`=COMPOUND, `motion_class`=MOVE, motion_sense=LITERAL)
MOVELINK (`id`=mvl1, `trigger`=m1, `mover`=se1, `source`=pl1, `goal`=pl2, `goal_reached`=TRUE, `motion_signalID`=a1,a2)

b. [**John**$_{se2}$] [**traveled**$_{m2}$] [**by car**$_{ms3}$].
MOTION (`id`=m2, `motion_type`=MANNER, `motion_class`=MOVE, `motion_sense`=LITERAL)
MOVELINK (`id`=mvl2, `trigger`=m2, `mover`=se2, `motion_signalID`=a3)

c. [**John**$_{se3}$] [**drove**$_{m3}$] [**to**$_{ms4}$] [**Worcester**$_{pl3}$] [**on**$_{s1}$] the [**Massachusetts Turnpike**$_{p1}$].
MOTION (id=m3, motion_type=COMPOUND, motion_class=FOLLOW, motion_sense=LITERAL)
MOVELINK (id=mvl3, trigger=m3, mover=se3, goal=pl3, goal_reached=TRUE, motion_signalID=a4, pathID=p1)

d. [**John**$_{se4}$] [**left**$_{m4}$] [**for**$_{ms5}$] [**Boston**$_{pl3}$].
MOTION (id=m4, motion_type=PATH, motion_class=LEAVE, motion_sense=LITERAL)
MOVELINK (id=mvl4, trigger=m4, mover=se4, goal=pl3, goal_reached=UNCERTAIN, motion_signalID=a5)

e. [**John**$_{se5}$] [**jumped**$_{m5}$] [**over**$_{ms6}$] the [**fence**$_{p2}$].
MOTION (id=m5, motion_type=COMPOUND, motion_class=MOVE_EXTERNAL, motion_sense=LITERAL)
MOVELINK (id=mvl5, trigger=m5, mover=se5, ground=p2, ajdunctID=a6)

f. [**John**$_{se6}$] [**walked**$_{m6}$] [**off**$_{ms7}$] the [**path**$_{pl4}$].
MOTION (id=m6, motion_type=COMPOUND, motion_type=DEVIATE, motion_sense=LITERAL)
MOVELINK (id=mvl6, trigger=m6, source=p4, mover=se6, motion_signalID=a7)

g. The [**brook**$_{p3}$] [**runs**$_{m7}$] [**along**$_{ms8}$] the [**road**$_{p4}$].
MOTION (id=m7, motion_type=PATH, motion_class=FOLLOW, motion_sense=FICTIVE)
MOVELINK (id=mvl7, trigger=m7, goal=p4, motion_signalID=a8)

h. The [**glacier**$_{p5}$] [**crept**$_{m8}$] [**down**$_{ms9}$] the [**valley**$_{p6}$].
MOTION (id=m8, motion_type=MANNER, motion_class=FOLLOW, motion_sense=INTRINSIC_CHANGE, motion_signalID=a9)
MOVELINK (id=mvl8, trigger=m8, mover=p5)

i. The [**clouds**$_{se7}$] [**spread**$_{m9}$] [**over**$_{s2}$] the Peruvian [**coast**$_{p7}$].
MOTION (id=m9, motion_type=MANNER, motion_class=MOVE, motion_sense=INTRINSIC_CHANGE)
MOVELINK (id=mvl8, trigger=m8, mover=se7)

Pay special attention to Examples (65h) and (65i). INTRINSIC_CHANGE MOTION tags will always have a motion_type of MANNER. This is due to the fact that the location of an entity is an extrinsic property. The location of the mover entities participating in INTRINSIC_CHANGE motion-events are not undergoing any change of location, but rather there is some dynamic change in an intrinsic characteristic—such as the size, shape, or conformation—of the mover participant.

5.3.2 Special Section: Ergative Motion Verbs

Be extra careful when annotating MOVELINK tags triggered by ergative motion verbs. Ergative verbs act differently depending on whether they are used transitively or intransitively. Note the sentences in Example (66) where the mover entities have been marked in **boldface**. Note that when *flew* is used transitively, as in Example (66c), it is the object 'being flown' that fulfills the role of the mover, but when used intransitively, as in Example (66a), it is the 'flyer' that is the mover. That is, in Example (66c), it is not the case that ***John*** *flew over the harbor*. Depending on

the context, it may be entailed that both the subject and direct object of a ditransitive, ergative, motion verb happen to undergo the same change of location. This is the case in Example (66f). In those cases it is appropriate to list multiple `mover` entities in the comment field as discussed in Section 10 or create a separate MOVELINK for each `mover` participant.

(66) a. **John** flew over the harbor.

b. The remote-controlled **plane** flew over the harbor.

c. John flew the remote-controlled **plane** over the harbor.

d. **John** drove to the airport.

e. The **taxi** drove to the airport.

f. The **taxi** drove **John** to the airport.

5.3.3 Special Section: Underspecified Mover Participants

When creating MOVELINK tags it is necessary to link from a MOTION tag to some other `mover` element.[13] In some cases there may not be an explicit `mover` participant to link to. Two situations where this problem arises are nominalized motion-events and motion-events that fall under certain grammatical or narrative moods (e.g., imperative, simulated etc.). In such cases where the `mover` is underspecified, and it is not evident from the context who or what is moving, it is acceptable to link from the MOTION tag to itself. When performing a link from a MOTION tag to itself, leave the `mover` attribute unspecified (i.e., blank). You should then write a short comment describing the `mover` participant. Example (67) demonstrates how to create MOVELINK tags for underspecfied `mover` participants and what informative comments might look like.

If the mover is underspecified, but it is implicit from the context that the `mover` participant is someone or something that is mentioned elsewhere in the text, the annotator may create a non-consuming tag to fill the role of the `mover`. For these situations it will be necessary to use a METALINK tag to establish the relation between the non-consuming tag and the tag that it refers to.

(67) a. There was some incredible night [**biking**$_{m1}$] [**in**$_{ms1}$] the [**Atacama Desert**$_{pl1}$].
MOTION (`id`=`m1`, `extent`="biking" `motion_type`=COMPOUND, `motion_class`=MOVE_INTERNAL, `motion_sense`=LITERAL)
MOVELINK (`id`=`mvl1`, `fromText`="biking", `toText`="biking", `trigger`=`m1`, `mover`=∅[14], `ground`=`pl1`, `motion_signalID`=`a1`, `comment`="The author is describing biking, but it is not clear who participated as the mover.")

Note: If, from context, it is clear that the author is participating in the *biking*, and the author mentions him or herself explicitly elsewhere in the text, then it would be appropriate to create a non-consuming tag to represent the author as participating as the `mover`. Then the non-consuming tag would need to be linked to the explicit mention via a METALINK.

[13] Refer to Section 10 for details on link creation in MAE.

[14] The ∅ symbol is used here to indicate an underspecified attribute value that would be left blank.

b. [**Take**$_{m1}$] the [**stairs**$_{p1}$] [**to**$_{ms1}$] the [**roof**$_{pl1}$].
MOTION (`id`=`m1`, `extent`="Take", `motion_type`=PATH, `motion_class`=FOLLOW, `motion_sense`=LITERAL)
MOVELINK (`id`=`mvl1`, `fromText`="take", `toText`="take", `trigger`=`m1`, `mover`=∅, `goal`=`pl1`, `pathID`=`p1`, `motion_signalID`=`a1`, `comment`="This is an imperative construction where the mover participant would be anyone that follows the direction.")

5.4 Measure Links

The MLINK tag serves two purposes. First, it can be used to capture the distance between two locations as in *The bone is two feet from the dog.* Such relationships are commonly accompanied by a MEASURE extent, but this is not a requirement. For example, the phrase *the hot dog stand near Macy's* also introduces an MLINK since *near* is interpreted on a scale.

In addition to relating two spatial objects, measure links can also be used to describe the dimensions of a single object. Locations, spatial entities, and even events possess spatial dimensions that may be captured by an MLINK tag. A typical case where the MLINK tag is used is when the length dimension of a location is described as in *The football field is 100 yards long.* Note, however, that the MLINK tag can also capture dimensions of motion-events as in *I rode 30 miles* (Examples (68b) and (68c)). In such a case the MLINK is actually specifying a dimension of the event-path introduced by the MOTION.

5.4.1 Measure Link Attributes

The attributes for the MLINK tag are presented in Table 18.

Attribute	Value
`id`	`ml1`, `ml2`, `ml3`, ...
`trajector`	ID of a location/entity/event tag
`landmark`	ID of the related location/entity/event tag
`relType`	DISTANCE, LENGTH, WIDTH, HEIGHT, or GENERAL_DIMENSION
`val`	NEAR, FAR, TALLER, SHORTER, or ID of a MEASURE tag
`endPoint1`	ID of a location/entity/event tag at one end of a stative-path
`endPoint2`	ID of a location/entity/event tag at the other end of a stative-path

Table 18: Attributes for MLINK

When the MLINK tag is used to describe the relationship between two spatial objects, their IDs are filled in the `trajector` and `landmark` attributes. In the other MLINK usage, in which only one spatial object is described, that object's ID should be filled in the `trajector` attribute and either repeated as the `landmark` or the `landmark` attribute should be left unspecified.

The `relType` attribute describes what dimension is being measured with the MLINK. The possible values are DISTANCE, LENGTH, WIDTH, HEIGHT, or GENERAL_DIMENSION. Table 19 describes how to choose the appropriate `relType` value depending on the dimension being measured.

The `val` attribute describes the actual measurement. Its value is filled with the ID of a MEASURE tag. For now, the `relType` has a closed set of possible values, but this may change as the pilot annotation proceeds. If the annotator believes an MLINK is appropriate but is not satisfied with the possible values for the link attributes, he or she should comment on this in the MLINK's annotation.

When a stative path is used to describe the dimensions of an object, any endpoints that bound the object should appear in the `endPoint1` and `endPoint2` attributes. As usual, the values for

relType Value	Description
DISTANCE	distance between two spatial objects
LENGTH	intrinsic length of a single spatial object
WIDTH	intrinsic width of a single spatial object
HEIGHT	intrinsic height of a single spatial object
GENERAL_DIMENSION	the dimension being measured is not clear

Table 19: MLINK relType Values

these attributes can be the ID of any location tag (i.e., places, paths, events, etc.). Example (68) shows annotations for several MLINK tags.

(68) a. The new [**tropical depression**$_{se1}$] was about [**430 miles**$_{me1}$] ([**690 kilometers**$_{me2}$]) west of the [**southernmost Cape Verde Island**$_{pl1}$], forecasters said.
MLINK (id=ml1, relType=DISTANCE, trajector=se1, ground=pl1, val=me1)
MLINK (id=ml2, relType=DISTANCE, trajector=se1, ground=pl1, val=me2)

b. [**The football field**$_{se2}$] is [**100 yards**$_{me3}$] long.
MLINK (id=ml3, relType=LENGTH, trajector=se2, ground=se2, val=me3)

c. I [**rode**$_{m1}$] [**30 miles**$_{me4}$] yesterday.
MLINK (id=ml4, relType=GENERAL_DIMENSION, trajector=m1, ground=m1, val=me4)

d. [**The office**$_{pl3}$] is [**25 feet**$_{me5}$] wide from [**the bookcase**$_{se3}$] to [**the white board**$_{se4}$].
MLINK (id=ml6, relType=WIDTH, trajector=pl3, ground=pl3, val=me5, endPoint1=se3, endPoint2=se4)

e. [**The hot dog stand**$_{se5}$] [**near**$_{me6}$] [**Macy's**$_{se6}$].
MLINK (id=ml7, relType=DISTANCE, trajector=se5, ground=se6, val=me6)

6 Special Section: Quantification

This section is intended to cover a number of attributes which handle quantification over spatial elements. These attributes, namely gquant, countable, and scopes, are common to many different tag types so they have been consolidated here. The tag types which share these attributes include location tags, namely PLACE and PATH tags, as well as any tags that may be coerced to act like a location including SPATIAL_ENTITY, NONMOTION_EVENT and MOTION tags. Table 20 reiterates the attributes for handling quantification.

Attribute	Value
COUNTABLE	TRUE or FALSE
GQUANT	a generalized quantifier
SCOPES	an ID of a location/entity/event tag that is the *scopee* in a *scopes*(*scoper*, *scopee*) relation

Table 20: Quantification Attributes

A value of TRUE is usually filled for the countable attribute for nouns that are individually countable such as *trees*, *lakes*, *roads* or *trips*. A value of FALSE might be used for mass terms such

as *water*, or *countryside*. Recall that languages such as English possess mechanisms for coercing mass-terms to act like count-terms (e.g., in English, *There are three waters in the fridge*) and vice versa (e.g., *Collectively, the trees surround the building*).

The `gquant` attribute takes a generalized quantifier as its value. Generalized quantifiers are most commonly introduced by determiner phrases, though they may be implicit as well, as is often the case for generic references (e.g., *snow* is taken to be universally quantified in *Snow is white*). Determiners may indicate universal or existential quantification depending on context (e.g., in English, the indefinite article *a(n)* may introduce universal or existential quantification depending on the context; the same is the case for the grammatically plural articles *some* and *any*). There are also many determiners in English that indicate a vaguely specified quantity, e.g., *some, several, many, few, most*. Numerals (e.g., *half, one, two, one hundred*, ...) or other numerical terms (e.g., *both, a dozen*) may indicate a specific quantity. Other expressions may introduce a constraint on the quantity, e.g., *more than a few, nearly every, most*.

The `scopes` attribute specifies a scoping relation between the *scoper*, which is the tag element whose `scopes` attribute is being filled, and the *scopee*, which is the tag element whose ID is filled as the value. That is, if the `scopes` attribute for a `spatial_entity` tag `se1` were filled with the value `pl1`, we can represent this relation in terms of a $scopes(scoper, scopee)$ function such that $scopes(\texttt{se1}, \texttt{pl1})$ meaning that the quantifier for `se1` has scope over the quantifier for `pl1`.

The following sentences in Example (69) are aimed to help illustrate how to capture scoping relations with the `scopes` attribute:

(69) a. A [**computer**$_{se1}$] [**on**$_{s1}$] every [**desk**$_{se2}$].
SPATIAL_ENTITY (`id`=`se1`, `text`="computer", `form`=NOM, `countable`=TRUE, `gquant`="1", `scopes`=∅[15])
SPATIAL_ENTITY (`id`=`se2`, `text`="desk", `form`=NOM, `countable`=TRUE, `gquant`="every", `scopes`=`se1`)
SIGNAL (`id`=`s1`, `semantic_type`=DIR_TOP)
QSLINK (`id`=`qsl1`, `relType`=EC, `figure`=`se1`, `ground`=`se2`, `trigger`=`s1`)
OLINK (`id`=`ol1`, `relType`="above", `figure`=`se1`, `ground`=`se2`, `trigger`=`s1`, `frame_type`=INTRINSIC, `referencePt`=`se2`, `projective`=FALSE)
$\forall se_2\ \exists_1 se_1\ [se_2 \in \text{DESKS} \wedge se_1 \in \text{COMPUTERS} \rightarrow \text{EC}(se_2, se_1) \wedge \text{ABOVE}(se_2, se_1)]$

b. [**I**$_{se1}$] [**rode**$_{m1}$] [**along**$_{ms1}$] some steep [**roads**$_{p1}$].
SPATIAL_ENTITY (`id`=`se1`, `text`="I", `form`=`nom`, `countable`=`true`, `gquant`="1", `scopes`=`p1`)
PATH (`id`=`p1`, `text`="roads", `form`=NOM, `countable`=TRUE, `gquant`="some", `scopes`=`m1`)
MOTION (`id`=`m1`, `text`="rode", `motion_type`=COMPOUND, `motion_class`=FOLLOW, `motion_sense`=LITERAL, `countable`=TRUE, `gquant`="exists", `scopes`=∅)
MOTION_SIGNAL (`id`=`a1`, `motion_signal_type`=`path`)
MOVELINK (`id`=`mvl1`, `trigger`=`m1`, `mover`=`se1`, `pathID`=`p1`, `motion_signalID`=`a1`)
$\exists_1 se_1\ \exists_{some} p_1\ \exists m_1\ [se_1 \in \text{INDIVIDUALS} \wedge p_1 \in \{\text{STEEP} \cap \text{ROADS}\} \wedge m_1 \in \text{RIDINGS} \wedge follows(m_1, p_1) \wedge mover(m_1, se_1)]$

c. ... every [**car**$_{se2}$] that [**passed**$_{m1}$] [**me**$_{se1}$] had at least 3 or more [**people**$_{se3}$] [**inside**$_{s1}$].
SPATIAL_ENTITY (`id`=`se2`, `text`="car", `form`=NOM, `countable`=TRUE, `gquant`="every", `scopes`=`m1`)
SPATIAL_ENTITY (`id`=`se1`, `text`="me", `form`=NOM, `countable`=TRUE, `gquant`="1", `scopes`=`se2`)

[15] The symbol ∅ is used here to indicate an unspecified attribute value.

SPATIAL_ENTITY (id=se3, text="people", form=NOM, countable=TRUE, gquant=gte3, scopes=∅)
SPATIAL_SIGNAL (id=s1, text="inside", semantic_type=TOPOLOGICAL)
MOTION (id=m1, text="passed", motion_type=PATH, motion_class=MOVE_EXTERNAL, motion_sense=LITERAL, countable=TRUE, gquant="exists", scopes=se3)
MOVELINK (id=mvl1, trigger=m1, mover=se2, ground=se1)
QSLINK (id=qsl1, relType=IN, figure=se3, ground=se2)
$\exists_1 se_1\ \forall se_2\ \exists m_1\ \exists_{\geq 3} se_3\ [se_1 \in \text{INDIVIDUALS} \wedge se_2 \in \text{CARS} \wedge m_1 \in \text{PASSINGS} \wedge se_3 \in \text{PEOPLE} \wedge moves\_external(m_1, se_1) \wedge mover(m_1, se_2) \rightarrow \text{IN}(se_2, se_3)]$

Partial logical translation has been provided following each annotation in Example (69) to make them easier to interpret. Note that any time multiple quantified variables are introduced there exists a possibility for scoping ambiguities. Like any lexical or syntactic ambiguity, not all quantificational scoping ambiguities will necessarily resolve to a single plausible interpretation. Part of the annotator's responsiblity is to assume an interpretation which they find most appropriate and ensure that their annotation is consistent with that interpretation.

The interpretation assumed for the annotation in Example (69a) is one under which for every se_2 there exists some se_1 such that if se_2 is a *desk* and se_1 is a *computer* then se_1 is *on* se_2. The other possible interpretation is the one under which it holds that there exists some *computer* for every *desk* such that that particular *computer* is *on* every *desk*. This second interpretation was discarded in this case since it seems unlikely, at least pragmatically, that one particular computer would be on every desk. Provided an appropriate context, however, that interpretation might be plausible. E.g., given a domain of discourse where the set of desks is relatively small, and those desks are arranged in such a way that their surfaces are contiguous, a particular computer could potentially be on every desk. Under this "particular computer" interpretation, though, the scopes attribute values for the SPATIAL_ENTITY tags for *computer* and *desk* would need to reflect that scoping relation such that $scopes(se_1, se_2)$. The "particular computer" interpreation is not reflected in the annotation provided where the scoping relation is such that $scopes(se_2, se_1)$.

The annotation in Example (69b) corresponds to an individual interpretation where there is a distinct *riding* motion-event for each *steep road*. This is contrastive with a collective interpretation, which would require that the countable attribute for the PATH tag for *roads* (p1) be set to FALSE. Such a collective interpretation would be one for which a single *riding* event occurred in which the *steep roads*, collectively, were traversed. Pragmatically, for this to be plausible, one would have to imagine the roads to be configured in a parallel series. Under the individual interpretation, which corresponds to the annotation example, it holds that there exists an entity se_1 for some paths p_1, and for p_1 there exists a motion-event m_1 such that the se_1 is an individual (referenced by I) and m_1 is a *riding* and p_1 are *steep roads* and those m_1 traverse p_1 and the mover participant in m_1 is the individual se_1.

The interpretation that has been annotated in Example (69c) is another individual interpretation where the countable attributes for each entity have been annotated as TRUE. The gquant attribute for se_2 takes the value "every" due to the determiner *every* while se_1 and se_3 are quantified existentially. For se_1 there is an explicit, definite determiner (*3 or more*) which is captured with the gquant value of "geq3", and for se_3 the gquant attribute takes a value of "1" since *me* is a grammatically singular pronoun. Under this interpretation it holds that there exists an individual entity se_1 for every entity se_2 and for every se_2 there exists a motion-event m_1 and for m_1 there exist *3 or more* entities se_3 such that if se_1 is an individual (referenced by *me*) and se_2 is a *car*

m_1 is a *passing* and se_3 are *people* and m_1 is a *passing by* of se_1 and the mover participant of m_1 is se_2, then se_3 are *inside* se_2.

Pay special attention to the `scopes` attributes in the above examples. If there are n quantified variables, then there may be as many as $n-1$ scoping relations. For Example (69a) there is one scoping relation which can be represented in terms of a $scopes(scoper, scopee)$ funciton where $scopes(se_2, se_1)$. For Example (69b) there are two scoping relations: $scopes(se_1, p_1)$ and $scopes(p_1, m_1)$. In Example (69c) there are three scoping relations: $scopes(se_1, se_2) \wedge scopes(se_2, m_1) \wedge scopes(m_1, se_3)$. Note how the ordering of the quantifiers in the logical formulae, after applying quantifier raising (QR), reflect these scoping relations in the annotations.

7 Phased Approach for the Pilot Annotation

This section outlines a phased approach for conducting the annotation. In this approach there will be tow annotation phases: one for extent tags and one for link tags. In Phase 1, annotators will tag PLACE, PATH, SPATIAL_ENTITY, NONMOTION_EVENT, MOTION_SIGNAL, NONMOTION_EVENT, MEASURE, and SPATIAL_SIGNAL extents. In this phase, annotators will also create any non-consuming tags in anticipation of links that will require them. Then a round of adjudication for the extent tags will be performed. Subsequently, in Phase 2, annotators will create all MOVELINK, MLINK, OLINK, QSLINK and METALINK tags. A final round of adjudication will be performed for the link tags after Phase 2.

7.1 Outline of Phases

- Extents Annotation
  - Phase 1
    1. Places
    2. Paths
    3. Spatial Entities
    4. Motions
    5. Motion Signals
    6. Events
    7. Measures
    8. Spatial Signals
- Extents Adjudication
- Links Annotation
  - Phase 2
    1. Movement Links
    2. Measure Links
    3. Orientation Links
    4. Qualitative Spatial Links
    5. Meta Links
- Links Adjudication

This phased approach is intended to divide the annotation into sub-tasks. Due to interdependencies in the sub-tasks, Phase 2 must be built on top of Phase 1. This division of the annotation and adjudication task anticipates multiple annotators and adjudicators working on the phases independently, but also, importantly, is intended to ensure that annotators will never be creating links between unadjudicated extent tags.

While the phases are somewhat independent, there are a few dependency caveats. The following list provides recommendations and exceptions to the order as outlined previously.

- All Phase 2 annotations must be built on top of Phase 1. After a Phase 2 annotation is complete, the resulting annotation ought to include all tags from Phases 1 and 2, combined.
- Phase 2 annotations must be built on top of an adjudicated set of extent tags, thus a Phase 2 annotation will be created on top of a file containing a locked set of extent tags. After the links for Phase 2 have been created, the resulting annotation will be complete, pending the final adjudication of the link tags.
- In the case that a MEASURE tag ID would be used to fill an `elevation` attribute for an extent tag created in Phase 1, then it is recommended that annotators create the `measure` tag immediately so that the `elevation` may be filled concurrently.
- Non-consuming tags[16] should be created in anticipation of their participation of any links. For example, in the event that an annotator creates a MOTION tag in Phase 1, and a non-consuming tag would ultimately participate in the MOVELINK triggered by that MOTION—as is often the case with the CROSS `motion_class`, for example—then the non-consuming tag should be created at the same time as the MOTION tag.
- This applies to spatial signals as well (e.g., *across*). If a SPATIAL_SIGNAL tag annotated in Phase 1 would trigger an OLINK or QSLINK in which a non-consuming location tag would participate then that non-consuming tag should be created at the same time as the SPATIAL_SIGNAL if it was not already created in Phase 1.
- Similarly, if a non-consuming tag would ultimately participate in an MLINK then the non-consuming tag should be created at the same time as the MEASURE tag which would trigger the MLINK.

8 Annotation Examples

In this section, we present several fully annotated examples from real, natural language text. The examples are presented using the usual predicate-argument form. There is also a short explanation of the annotator's choices for each example.

(70) a. ...two palm [**trees**$_{se1}$] [**in front of**$_{s1}$] the [**house**$_{pl1}$] ...
SPATIAL_ENTITY (`id`=`se1`, `extent`="trees", `form`=NOM, `countable`=TRUE)
PLACE(`id`=`pl1`, `extent`="house", `form`=NOM)
SPATIAL_SIGNAL (`id`=`s1`, `extent`="in front of", `semantic_type`=DIR_TOP)
QSLINK (`id`=`qsl1`, `relType`=DC, `trajector`=`se1`, `ground`=`pl1`, `trigger`=`s1`)
OLINK (`id`=`ol1`, `relType`=FRONT, `trajector`=`se1`, `ground`=`pl1`, `trigger`=`s1`, `frame_type`=INTRINSIC, `referencePt`=`pl1`, `projective`=FALSE)

[16] For further details regarding non-consuming tags refer to Section 2.4.

This example is relatively straightforward. Note that *trees* is tagged with the SPATIAL_ENTITY tag instead of the PLACE tag. As discussed in Section 2.3, this may seem confusing, and in some cases it might seem plausible to annotate *trees* with the PLACE tag, but in general *trees* are not the kind of thing that one can *be at* or *take*, though it may be possible in specific circumstances. Note also that it may be unnecessary to have a QSLINK for *in front of*, since this spatial signal deals mostly with orientation and not necessarily a topological relationship. It is plausible that a number of RCC8$^+$types could be used depending on the annotator's chosen interpretation. This annotation assumes that the trees are disconnected from the house, thus the DC `relType` was supplied for the QSLINK. When in doubt, it is better to err on the side of underspecification so as to avoid specifying something that is false.

b. A [**woman**$_{se2}$] and a [**child**$_{se3}$] are [**walking**$_{m1}$] [**over**$_{ms1}$] the [**square**$_{pl2}$] [$\varnothing_{pl3}$] [$\varnothing_{pl4}$] [$\varnothing_{pl5}$] [$\varnothing_{pl6}$]

SPATIAL_ENTITY (`id`=se2, `extent`="woman", `form`=NOM, `mod`=∅, `dcl`=FALSE, `countable`=TRUE)

SPATIAL_ENTITY (`id`=se3, `extent`="child", `form`=NOM, `mod`=∅, `dcl`=FALSE, `countable`=TRUE)

PLACE (`id`=pl2, `extent`="square", `form`=NOM, `mod`=∅, `dcl`=FALSE, `countable`=TRUE)

PLACE (`id`=pl3, `extent`=∅, `form`=NOM, `mod`=∅, `dcl`=FALSE, `countable`=TRUE, `comment`="The source for the [walking,m1] in mvl1")

PLACE (`id`=pl4, `extent`=∅, `form`=NOM, `mod`=∅, `dcl`=FALSE, `countable`=TRUE, `comment`="The goal for the [walking,m1] in mvl1")

PLACE (`id`=pl5, `extent`=∅, `form`=NOM, `mod`=∅, `dcl`=FALSE, `countable`=TRUE, `comment`="The source for the [walking,m1] in mvl2")

PLACE (`id`=pl6, `extent`=∅, `form`=NOM, `mod`=∅, `dcl`=FALSE, `countable`=TRUE, `comment`="The goal for the [walking,m1] in mvl2")

MOTION (`id`=m1, `extent`="walking", `motion_type`=COMPOUND, `motion_class`=CROSS)

MOTION_SIGNAL (`id`=a1, `extent`="over", `motion_signal_type`=PATH)

MOVELINK (`id`=mvl1, `trigger`=m1, `source`=pl3, `goal`=pl4, `midPoint`=pl2, `mover`=se2, `ground`=∅, `goal_reached`=TRUE, `pathID`=∅, `motion_signalID`=a1)

MOVELINK (`id`=mvl2, `trigger`=m1, `source`=pl5, `goal`=pl6, `midPoint`=pl2, `mover`=se3, `ground`=∅, `goal_reached`=TRUE, `pathID`=∅, `motion_signalID`=a1)

OLINK (`id`=ol1, `relType`="ACROSS", `trajector`=pl4, `landmark`=pl2, `trigger`=∅, `frame_type`=RELATIVE, `referencePt`=pl3, `projective`=FALSE)

OLINK (`id`=ol2, `relType`="ACROSS", `trajector`=pl6, `landmark`=pl2, `trigger`=∅, `frame_type`=RELATIVE, `referencePt`=pl5, `projective`=FALSE)

The tricky part here can be how to handle the motion of two entities. Creating two entities and a MOVELINK for each one leads to the suggestion that the entities are not walking over the square together, even though this is likely the case, as lone children are usually accompanied by an adult. However, it could simply be the case that an observer is describing everyone crossing the square; perhaps the woman and the child are not even walking next to each other. To prevent introducing false information into the annotation, it is preferable to underspecify the relationship between the *woman* and the *child*.

Note that *over* is lexically ambiguous. In this annotation it has been tagged as a MOTION_SIGNAL associated with the walking motion-event. This assumes an interpretation where *over* (a1) functions like *across*, suggesting that the motion class for the *walking* motion-event be annotated as CROSS. Recall that MOTION_SIGNAL tags do not trigger OLINK or QSLINK tags. If *over* were interpreted as functioning like the spatial preposition *above*, establishing a directional relation between the *walking* motion-event (m1) and the *square* (pl2), then it should be tagged as a SPATIAL_SIGNAL so as to trigger an OLINK. As is, however, it is contributing path-of-motion information.

Additionally, the non-consuming location tags pl3 and pl4, and pl5 and pl6 were created to specify the source and goal for the MOVELINK tags mvl1 and mvl2. The OLINK ol1 establishes that pl4 is across the *square* (pl2) relative to pl3, and similarly pl6 is across the *square* (pl2) relative to pl5. This is necessary under the looser interpretation where the *woman* and *child* are not necessarily participating in the same motion-event, and thus might begin and end in different locations.

c. The [**Prudential Center**$_{pl3}$], situated [**on**$_{s2}$] [**23 acres**$_{me1}$] ([**93,000 m2**$_{me2}$]), is [**in**$_{s3}$] the [**Back Bay neighborhood**$_{pl4}$] [**at**$_{s4}$] [**800 Boylston Street**$_{pl8}$] and [**houses**$_{s5}$] a [**495,229-square-foot**$_{me3}$] ([**46,008.3 m2**$_{me4}$]) shopping [**mall**$_{pl6}$], the [**Shops at Prudential Center**$_{pl7}$], [**in**$_{s6}$] the [**base**$_{pl9}$]. [$\varnothing_{pl10}$]
PLACE (id=pl3, extent="Prudential Center", form=NAM, mod=∅, dcl=FALSE, countable=TRUE)

PLACE (id=pl4, extent="neighborhood", form=NAM, mod=∅, dcl=FALSE, countable=TRUE)

PLACE (id=pl6, extent="mall", form=NOM, mod=∅, dcl=FALSE, countable=TRUE)

PLACE (id=pl7, extent="Shops at Prudential Center", form=NAM, mod=∅, dcl=FALSE, countable=TRUE)

PLACE (id=pl8, extent="800 Boylston Street", form=NAM, mod=∅, dcl=FALSE, countable=TRUE)

PLACE (id=pl9, extent="base", form=NOM, mod=∅, dcl=FALSE, countable=TRUE)

PLACE (id=pl10, extent=∅, form=NOM, mod=∅, dcl=FALSE, countable=TRUE)

SPATIAL_SIGNAL (id=s2, extent="on", semantic_type=TOPOLOGICAL)

SPATIAL_SIGNAL (id=s3, extent="in", semantic_type=TOPOLOGICAL)

SPATIAL_SIGNAL (id=s4, extent="at", semantic_type=TOPOLOGICAL)

SPATIAL_SIGNAL (id=s5, extent="houses", semantic_type=TOPOLOGICAL)

SPATIAL_SIGNAL (id=s6, extent="in", semantic_type=TOPOLOGICAL)

MEASURE (id=me1, value="23", unit="acres", extent="23 acres")

MEASURE (id=me2, value="93000", unit="square meters", extent="93,000 m2")

MEASURE (id=me3, value="495229", unit="square feet", extent="495,229-square-foot")

MEASURE (id=me4, value="46008.3", unit="square meters", extent="46,008.3 m2")

MLINK (id=ml1, trajector=pl10, relType=GENERAL_DIMENSION, val=me1)

MLINK (id=ml2, trajector=pl10, relType=GENERAL_DIMENSION, val=me2)

MLINK (id=ml3, trajector=pl6, relType=GENERAL_DIMENSION, val=me3)

MLINK (id=ml4, trajector=pl6, relType=GENERAL_DIMENSION, val=me4)

QSLINK (id=qsl2, relType=EQ, trajector=pl3, ground=pl10, trigger=s2)

QSLINK (id=qsl3, relType=IN, trajector=pl3, ground=pl4, trigger=s3)

QSLINK (id=qsl4, relType=EQ, trajector=pl3, ground=pl8, trigger=s4)

QSLINK (id=qsl5, relType=IN, figure=pl6, ground=pl3, trigger=s5)

QSLINK (id=qsl6, relType=IN, figure=pl7, ground=pl9, trigger=s6)

QSLINK (id=qsl7, relType=EQ, figure=pl9, ground=pl3, trigger=∅)

OLINK (id=ol1, relType="BELOW", figure=pl9, ground=pl3, trigger=∅, frame_type=INTRINSIC, referencePt=pl3, projective=FALSE)

This sentence does not feature any motion-events, however there are still many static, spatial relations that need to be accounted for. Note that there are some multiword extent tags in this example. These multiword extents capture locations with proper names, and each name is captured with a single tag. There is also a non-consuming PLACE tag, pl10, which is introduced by the MEASURE tag me1. The location identified by pl10 happens to be identical to *Prudential Center* (pl3), so they were linked together with an EQ relType QSLINK. One other item of note is the PLACE tag for *base*. Although there is no SPATIAL_SIGNAL tag to indicate the relation of the *base* to any other tag, the word *base* itself presumes a topological and directional relation to whatever it is a *base of*. Thus, a QSLINK and OLINK (qsl7 and ol1) were created to capture the fact that the *base* is equivalent to (EQ) the *Prudential Center* (pl3) and intrinsically, beneath it. The trigger attributes for these link tags were left unspecified.

d. The tallest of these [**peaks**$_{pl11}$] have been [**covered**$_{s7}$] in [**snow**$_{se4}$] and [**ice**$_{se5}$] ...
SPATIAL_ENTITY (id=se4, extent="snow", form=NOM, dcl=FALSE, countable=FALSE)

SPATIAL_ENTITY (id=se5, extent="ice", form=NOM, dcl=FALSE, countable=FALSE)

PLACE (id=pl11, extent="peaks", form=NOM, mod="tallest", dcl=FALSE, countable=TRUE)

SPATIAL_SIGNAL (id=s7, extent="covered", semantic_type=DIR_TOP)

QSLINK (id=qsl6, relType=EC, trajector=se4, ground=pl11, trigger=s7)

OLINK (id=ol4, relType="ABOVE", trajector=se4, ground=pl11, trigger=s7, frame_type=INTRINSIC, referencePt=pl11, projective=FALSE)

QSLINK (id=qsl7, relType=EC, trajector=se5, ground=pl11, trigger=s7)

OLINK (id=ol4, relType="ABOVE", trajector=se5, ground=pl11, trigger=s7, frame_type=INTRINSIC, referencePt=pl11, projective=FALSE)

This sentence demonstrates a QSLINK triggered by the verb *covered*, which is acting as a SPATIAL_SIGNAL. This SPATIAL_SIGNAL tag triggers two QSLINK tags and two OLINK tags, one each for *snow* and *ice*. Note that there may be a temporal interpretation of the sentence under which the area *covered* by the *snow* and *ice* increased over time, in which case *covered* would be annotated as a MOTION with a motion_sense value of INTRINSIC_CHANGE. For this annotation, however, the interpretation is a static, or atemporal one. Without additional context to inform an annotator's decision, this would be the best option.

e. [**Departing**$_{m2}$] [**Copala**$_{pl11}$], the [**road**$_{p1}$] [**climbs**$_{m3}$] [**to**$_{ms1}$] [**over 6,000 feet**$_{me5}$] in [**30 miles**$_{me6}$], and then continues to [**climb**$_{m4}$] while [**hugging**$_{s8}$] an impressive cliff-lined [**ridgeline**$_{p2}$] literally called 'the spine of the devil.' [$\varnothing_{pl12}$]

PLACE (id=pl11, extent="Copala", form=NAM, elevation=∅, mod=∅, dcl=FALSE, countable=TRUE)

PLACE (id=pl12, extent=∅, elevation=me5, mod=∅, dcl=FALSE, countable=TRUE)

PATH (id=p1, extent="road", beginID=∅, endID=∅, midIDs=pl11,pl12, form=NOM, elevation=∅, mod=∅)

PATH (id=p2, extent="ridgeline", beginID=∅, endID=∅, midIDs=∅, form=NOM, elevation=∅, mod=∅, countable=TRUE)

MEASURE (id=me5, extent="over 6,000 feet", value="gt6000", unit="feet")

MEASURE (id=me6, extent="30 miles", value="30", unit="miles")

MLINK (id=ml5, trajector=m3, ground=m3, relType=GENERAL_DIMENSION, val=me6, endPoint1=pl11, endPoint2=pl12)

MOTION (id=m2, extent="Departing", motion_type=PATH, motion_class=LEAVE, motion_sense=FICTIVE)

MOVELINK (id=mvl2, trigger=m2, source=pl11, goal=∅, midPoint=∅, mover=p1, goal_reached=∅, pathID=p1, motion_signalID=∅)

MOTION (id=m3, extent="climbs", motion_class=MOVE, motion_sense=FICTIVE)

MOTION_SIGNAL (id=ms1, extent="to", motion_signal_type=PATH)

MOVELINK (id=mvl3, trigger=m3, source=pl11, goal=pl12, midPoint=∅, mover=p1, goal_reached=TRUE, pathID=p1, ajdunctID=ms1)

MOTION (id=m4, extent="climb", motion_class=MOVE, motion_sense=FICTIVE)

MOVELINK (id=mvl4, trigger=m4, source=pl12, goal=∅, midPoint=∅, mover=p1, goal_reached=∅, pathID=p1, ajdunctID=∅)

SPATIAL_SIGNAL (id=s8, extent="hugging", semantic_type=TOPOLOGICAL)

QSLINK (id=qsl8, relType=EC, trajector=p1, landmark=p2, trigger=s8)

This sentence is rather dense in terms of spatially relevant information, and it is made even trickier due to the figurative language that is employed. The first item of note is the non-consuming place tag, pl12 that has been created. Note, in this case, that the MEASURE ID of *over 6,000 feet* fills the elevation attribute of pl12. This non-consuming PLACE tag is then used later for the goal for the MOVELINK mvl3 triggered by the first *climb* MOTION tag, m3.

The second thing to note is that the motion_sense attribute values for all the MOTION tags are FICTIVE. This is because the *road* is fulfilling the role of mover in the figurative, atemporal interpretations of the *Departing*, *climbs*, and *climb* motion-events.

f. [**In**$_{s1}$] the [**Peruvian Jungle**$_{pl1}$], [**I**$_{se1}$] [**followed**$_{m1}$] dirt [**roads**$_{p1}$] [**south**$_{ms2}$] [**from**$_{ms1}$] [**Yurimaguas**$_{pl2}$], and planned to [**cross**$_{m2}$] [$\varnothing_{pl0}$] [**into**$_{ms3}$] the [**Andes**$_{p2}$] [**through**$_{ms4}$]

the national park [**Rio Abiseo**$_{pl3}$]. [**Arriving**$_{m3}$] [**in**$_{ms5}$] the [**town of Juanjui**$_{pl4}$], [**near**$_{me1}$] the [**park**$_{pl5}$], [**I**$_{se2}$] learned that my map had lied to me.

MOTION_SIGNAL (id=ms1, extent="From", motion_signal_type=PATH)

PLACE (id=pl1, extent="Peruvian Jungle", form=NAM, countable=TRUE)

SPATIAL_SIGNAL (id=s1, extent="in", semantic_type=TOP)

PLACE (id=pl2, extent="Yurimaguas", form=NAM, countable=TRUE)

SPATIAL_ENTITY (id=se1, extent="I", form=NOM, countable=TRUE)

MOTION (id=m1, extent="followed", motion_class=FOLLOW, motion_sense=LITERAL)

PATH (id=p1, extent="roads", form=NOM, countable=TRUE)

MOTION_SIGNAL (id=ms2, extent="south", motion_signal_type=PATH)

MOTION (id=m2, extent="cross", motion_class=CROSS, motion_sense=LITERAL)

PLACE (id=pl0, extent=∅, form=NOM, countable=TRUE comment="The source location for the crossing motion (m2)")

MOTION_SIGNAL (id=ms3, extent="into", motion_signal_type=PATH)

PATH (id=p2, extent="Andes", form=NAM, countable=FALSE)

MOTION_SIGNAL (id=ms4, extent="through", motion_signal_type=PATH)

PLACE (id=pl3, extent="Rio Abiseo", form=NAM, countable=TRUE)

QSLINK (id=qsl1, relType=IN, figure=m1, ground=pl1, trigger=s1)

MOVELINK (id=mvl1, trigger=m1, source=pl1, mover=SE1, pathID=p1, motion_signalID=a1,a2)

MOVELINK (id=mvl2, trigger=M2, source=pl0, goal=p2, midPoint=pl3, mover=SE1, goal_reached=TRUE, motion_signalID=a3,a4)

QSLINK (id=qsl2, relType=EC, figure=pl0, ground=pl3)

QSLINK (id=qs3, relType=EC, figure=pl3, ground=p2)

OLINK (id=ol1, relType="ACROSS", figure=p2, ground=pl3, frame_type=RELATIVE, referencePt=pl0, projective=FALSE)

MOTION (id=m3, extent="Arriving", motion_class=REACH, motion_sense=LITERAL)

MOTION_SIGNAL (id=ms5, extent="in", motion_signal_type=PATH)

PLACE (id=pl4, extent="town of Juanjui", form=NAM, countable=TRUE)

MEASURE (id=me1, extent="near", value="NEAR", unit=∅)

PLACE (id=pl5, extent="park", form=NOM, countable=TRUE)

SPATIAL_ENTITY (id=se2, extent="I", form=NOM, countable=TRUE)

MOVELINK (id=mvl3, trigger=M3, goal=pl3, mover=SE2, goal_reached=TRUE, motion_signalID=ms5)

MLINK (id=ml1, figure=pl4, ground=pl5, relType=DISTANCE, val=m1)

Note in the first sentence in this example that a QSLINK with a relType of IN (qsl1) was created to relate the *followed* MOTION tag (m1) to the *Peruvian Jungle* (pl1). This is a case

of coercion in which it is actually the location of the *followed* motion-event—essentially the event-path introduced by the motion—that is IN the region associated with `pl1`.

In the second sentence, note that *near* introduces a distal relation between *the town of Juanjui* (`pl4`) and *park* (`pl5`), so it is tagged as a MEASURE (`me1`) that triggers an MLINK (`ml1`) between the two location tags. Since *near* is a deictic spatial term, which is describing the relative distance between two locations, there is no unit of measure that can be filled in the `unit` attribute, so it is left unspecified. The `value` attribute is simply filled as "NEAR".[17]

Additionally, in the context of this example, notice that *south* is acting as a PATH type motion signal (`ms2`). In other contexts *south* may act as a directional SPATIAL_SIGNAL to trigger an OLINK between location tags (e.g., *Mexico is south of the U.S.*). In this instance, however, *south* is acting as a MOTION_SIGNAL because it is specifying a characteristic of the heading of the event-path introduced by *followed* (`m1`), rather than directly relating two location tags.

9 Tips, Tricks, and Strategies for the Pilot Annotation

In this section, we list several useful ideas that have come up over the course of the Pilot Annotation. Some of these are also included in the earlier sections of this document. It is expected that the next version of these Guidelines will incorporate all of these ideas.

1. Most tagged extents will subsequently participate in an link of some kind. If an extent doesn't participate in a link, its ID will probably be filled as an attribute for another tag. If you find that you end up with a "dangling" spatial object (i.e., something that was tagged as a PLACE, PATH, SPATIAL_ENTITY, SPATIAL_SIGNAL, MEASURE, NONMOTION_EVENT, MOTION, or MOTION_SIGNAL but that does not show up in a QSLINK, OLINK, MOVELINK, MLINK or fill an attribute for another tag), go back and reevaluate that tag. It may be that it should not have been annotated in the first place, or it may be that you have missed a link or other tag somewhere along the way. There are exceptions here, but not many. PLACE and PATH tags for locations with proper names should always be created even if they do not participate in any links. Also, if a spatial object participates in a METALINK but no other link, that is OK as long as the object it is related to by the METALINK does, itself, participate in a link.

2. Triggers should only be included in links if they both introduce the link and tell you something about the nature of that link. Not every link needs a trigger, but every MOTION and SPATIAL_SIGNAL triggers a link, and every MEASURE either fills an `elevation` attribute or triggers an MLINK.

3. For light verbs that involve motion, such as *catch a boat*, tag the verb *catch* as a MOTION and the entire extent of the NP, *a boat*, as an motion signal with `motion_signal_type` of MANNER. In some cases we will also want to tag the extent for *boat* as a SPATIAL_ENTITY for the purposes of linking it to other tags, but in general avoid creating extent tags that have overlapping extents.

[17] Open-class attribute values—values that are entered by the annotator, which are not IDs of tag elements—should be entered in all capitals and should not include whitespace. If it is necessary to delimit words within the value, use an underscore ('_') instead of a space.

4. If you're trying to figure out if a preposition like *by* is acting as a motion signal, or is simply an auxiliary verb in a passive voice construction, try unpassive-izing the phrase and it should become clear.

5. In examples like *sunk 6 meters* or *fell 2 feet*, the MEASURE tag `value` attribute should not be filled as negative. In such cases the directionality is contributed by the motion verbs *sunk* and *fell* themselves. The `value` attribute for MEASURE tags may be given a negative value, however, when identifying elevations that have a negative numerical offset (e.g., *500 ft below sea level* would have a value of "-500"). Additionally, don't fill separating commas (or other extraneous notation) in attributes filled with numerical values (e.g., fill "6000" not "6,000"). For non-integer values, use decimal notation, not fractions (e.g., "0.5" not "1/2").

10 Using MAE for Annotation

The following pages include directions for using MAE for annotation. In addition, the following guidelines should be followed for this Pilot Annotation.

1. For now, you should manually fill in the values for all link attributes, including those that come from MAE's `fromID` and `toID` attributes.

2. For those times when you need to provide a list of IDs (e.g, for the `midIDs` attribute in PATH), do so by separating the IDs with commas (e.g., "pl1, pl2"). List the IDs in ascending order even if that order does not correspond to the order they appear in the text.

3. If there is more than one comment in the comment field, separate them by semi-colons.

4. At this time, the specification only allows for single movers, so for multiple movers, list their IDs as a comma-separated list ending with a semicolon in the comment field for the MOVELINK (e.g "MOVERS: se0, se1, se3;").

Techniques in ISO-Space Annotation

by Seth Dworman

I. Getting Started

All annotation for ISO-Space will be done with MAE (Multi-purpose annotation environment), an annotation tool written by Amber Stubbs. The tool may be downloaded at Amber's homepage (http://pages.cs.brandeis.edu/~astubbs/) and found in the section labeled "MAE." Following this leads to additional link (http://code.google.com/p/mae-annotation/) which has the file needed for download. Head to the section labeled "Downloads," which will display the zip file "MAE_v0.9.6.zip." Download it and save it to your computer (recommended to desktop). Now simply extract the files from the zip (see image below), ideally to a folder with the same name. If you do not have software to do the extraction you can easily google it, e.g. WinRAR.

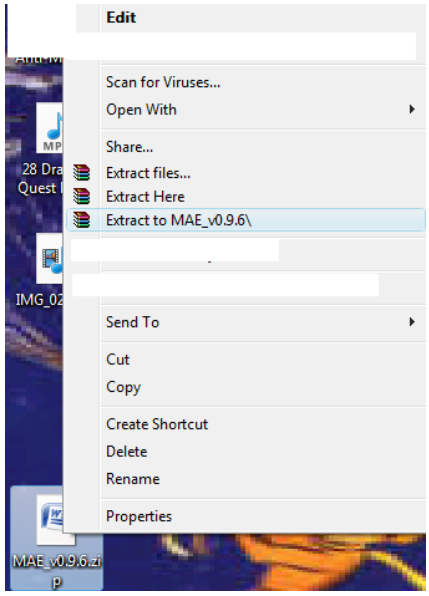

After extracting the files, you will end up with a folder containing said files. Open the folder and in the list of files there should be a "mae_v0.9.6.jar" file. This is the executable for the program—simply double click it to start up MAE (see image below).

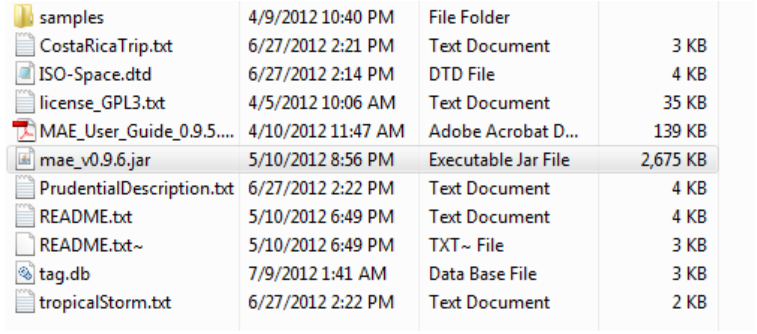

You should finally end up with some kind of window like below:

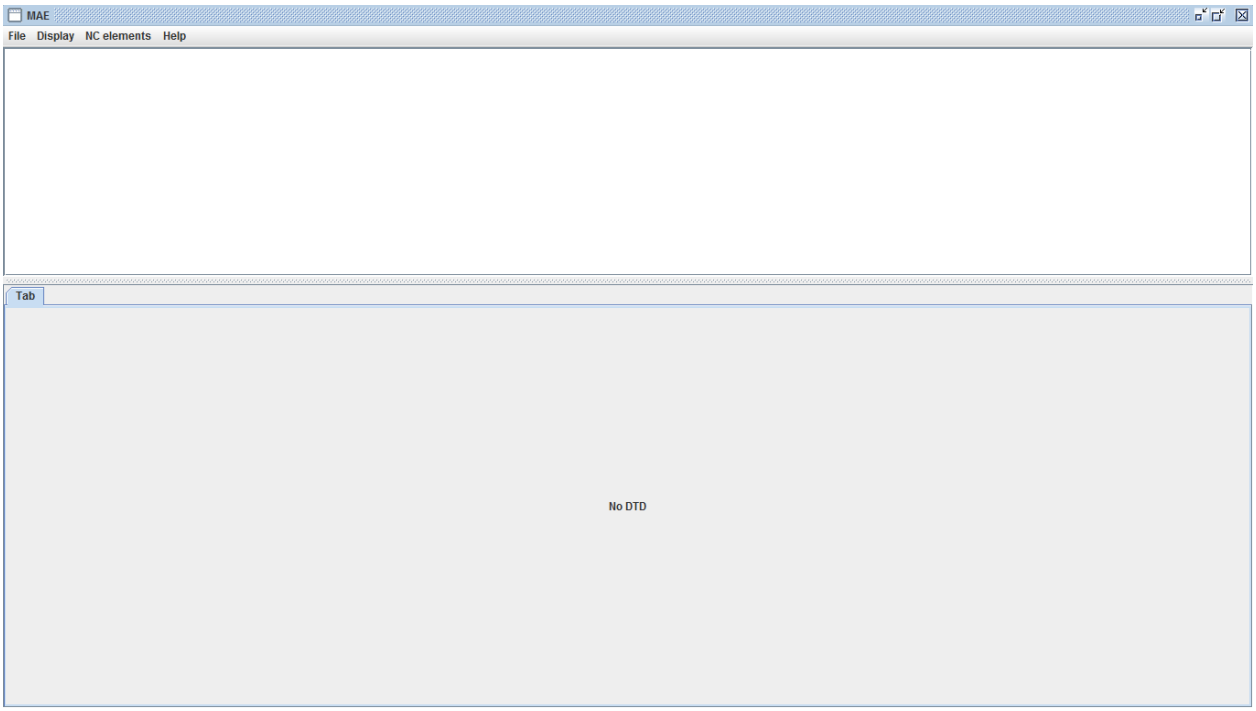

From here text/documents that are to annotated will be opened, but before any such annotation can be done, it is necessary to load the ISO-Space DTD, which contains all the information about the ISO-Space tags and links. You will have received the ISO-Space DTD at some point from Jessica Moszkowicz, at which point now you can simply head to File>Load DTD>ISO-Space.dtd (see image below). (Note the ISO-Space.dtd may not necessarily be in the same folder, but wherever you happened to save it.)

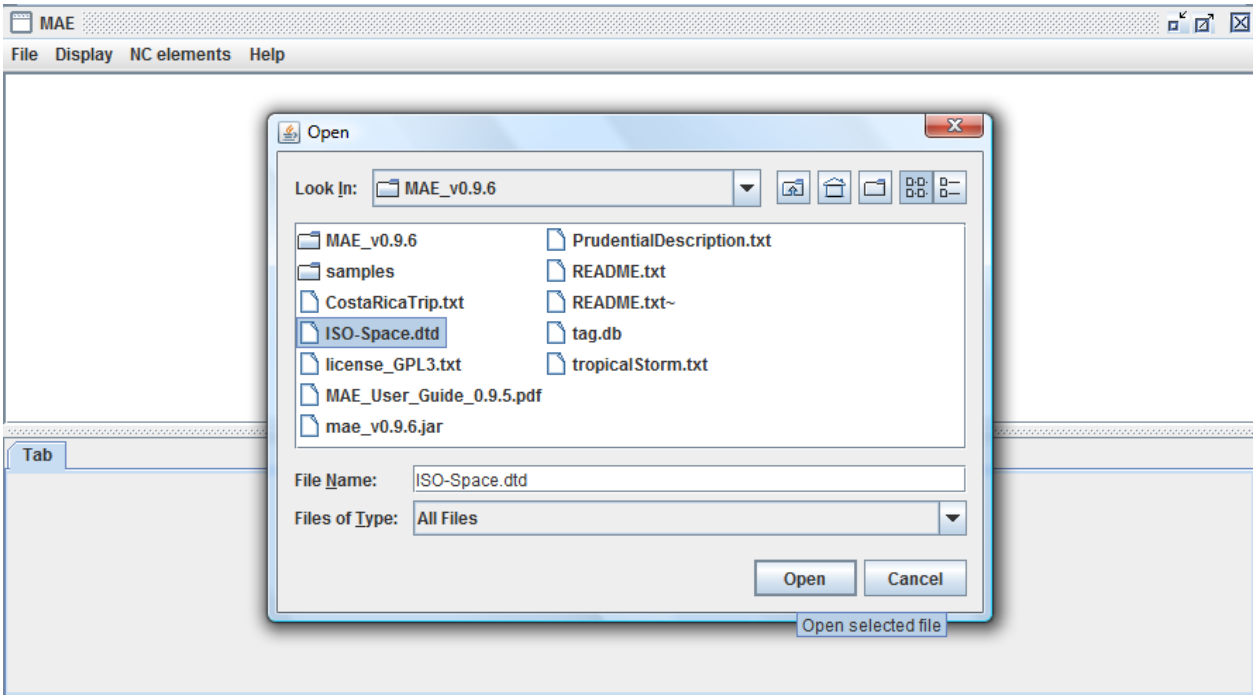

Finally, after loading the ISO-Space DTD, you will almost be ready to annotate! Now, simply head to File>Load File and choose any text/document that you wish to annotate. You should have something similar to below, which means you've correctly set-up MAE for ISO-Space.

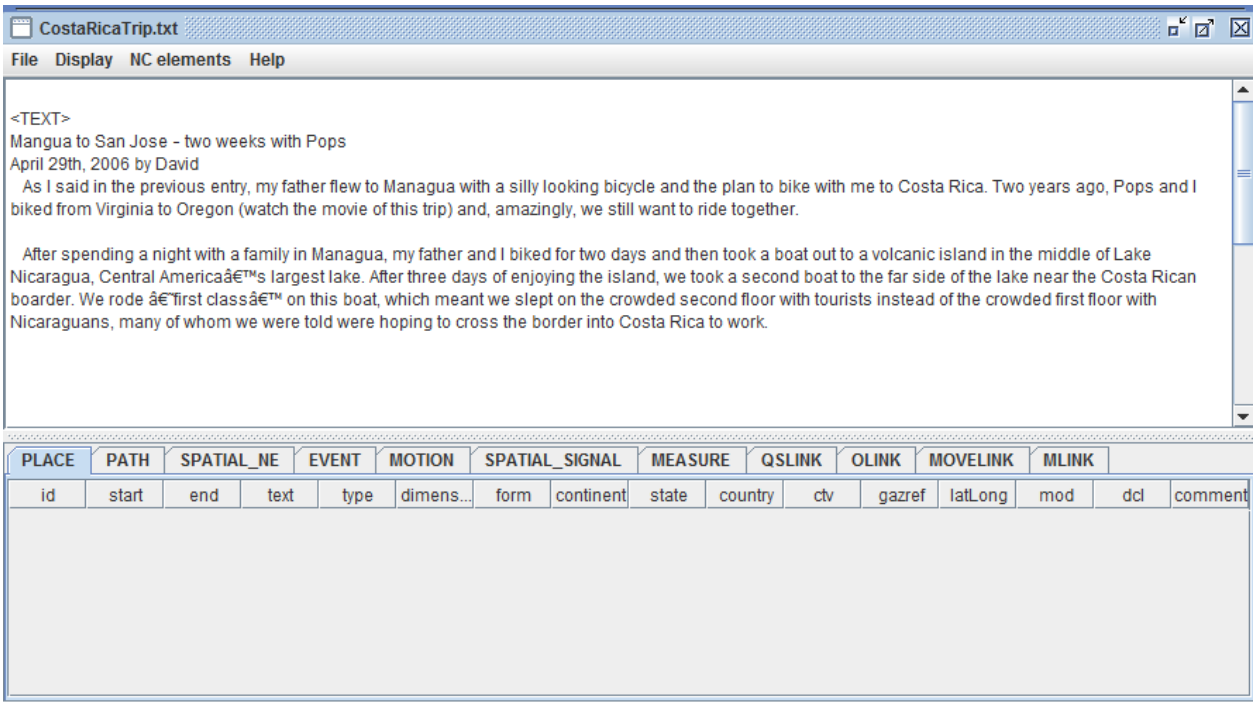

II. Using MAE

Whatever strategy may best suit you, the easiest part of the annotation is selecting what word(s) are places, paths, motions, spatial signals, etc. To tag a phrase, simply highlight the extent you wish to label and then right click on it. A dropdown menu should show up, listing all the possibilities. Simply select the one that is most logical with the extent (remember paths and spatial named entities can be coerced into being places).

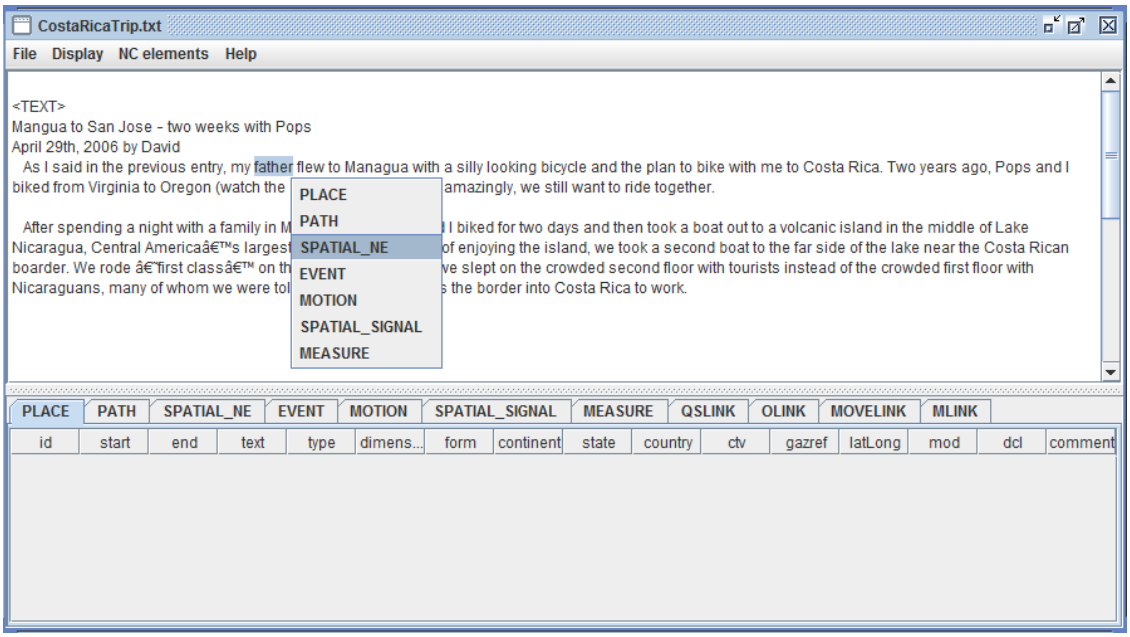

You will know if the extent was correctly labeled if both it appears in one of the tables below (under the respective tag) and the color of the extent in the text has also changed to reflect the tag.

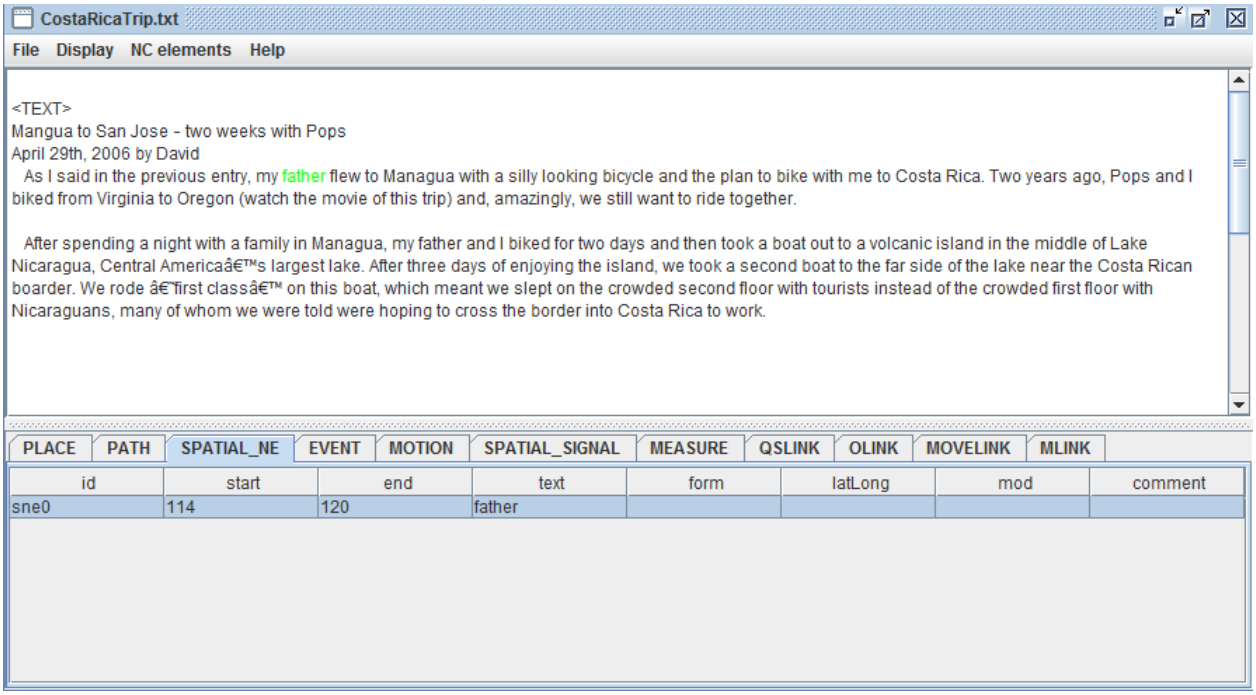

The “id” is used to index each of the different tags, starting with the type (“sne” for spatial named entity) and then the index (starting at 0). All tags and links are automatically indexed for you. However, if you decide to delete an entry in the table (done by right clicking anywhere on the entry and selecting “Remove selected...rows”) the indexing will not reflect the change. Additionally, you will need to use the “id” to reference the extent in links that it may participate in. Finally, note that not all of the attributes for the tag are filled in—if you are required to, you may have to manually fill in such attributes (more relevant for places, paths, spatial_signals, and motions) before moving onto the next extent. To do so, simply click on the empty entry below the attribute column, and either select a possible value or manually enter one in. The same can be done to entries that are already filled/completed (in case you felt there was a mistake).

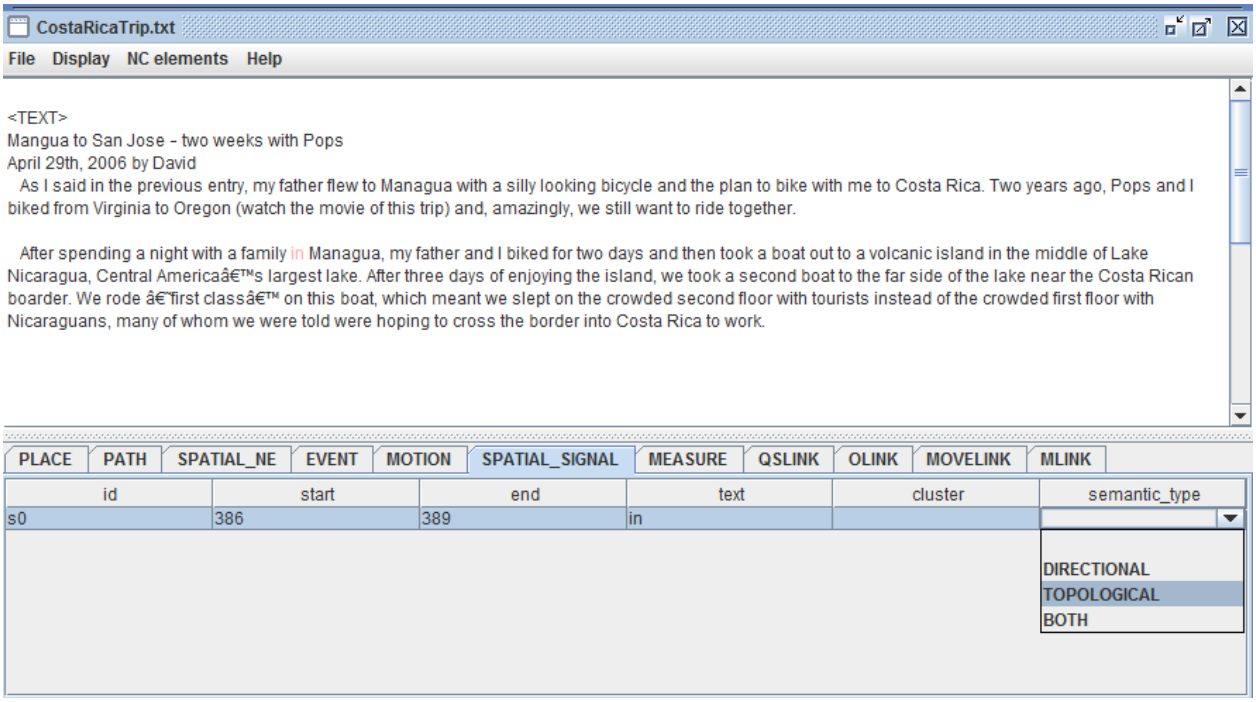

The most important task of the annotation, however, is how the various spatial relationships between the places, paths, motions, etc. are captured. This is done through various links, namely: QSLINK, OLINK, MOVELINK, MLINK, and METALINK. A link essentially will relate several (already annotated) extents to each other in some way. To begin a link, use CTRL+Left Click on the first and second extents you wish to associate in a link.

For QSLINK, OLINK, and MLINK, the first extent you click will be the figure (currently labeled as fromID) and the second extent will be the ground (currently labeled as toID). If an MLINK does not have a ground (e.g. simply a measurement of something), simply click the same extent twice. For MOVELINK, the first extent is the trigger/motion event (again labeled as fromID) and the second extent is the mover (here toID).

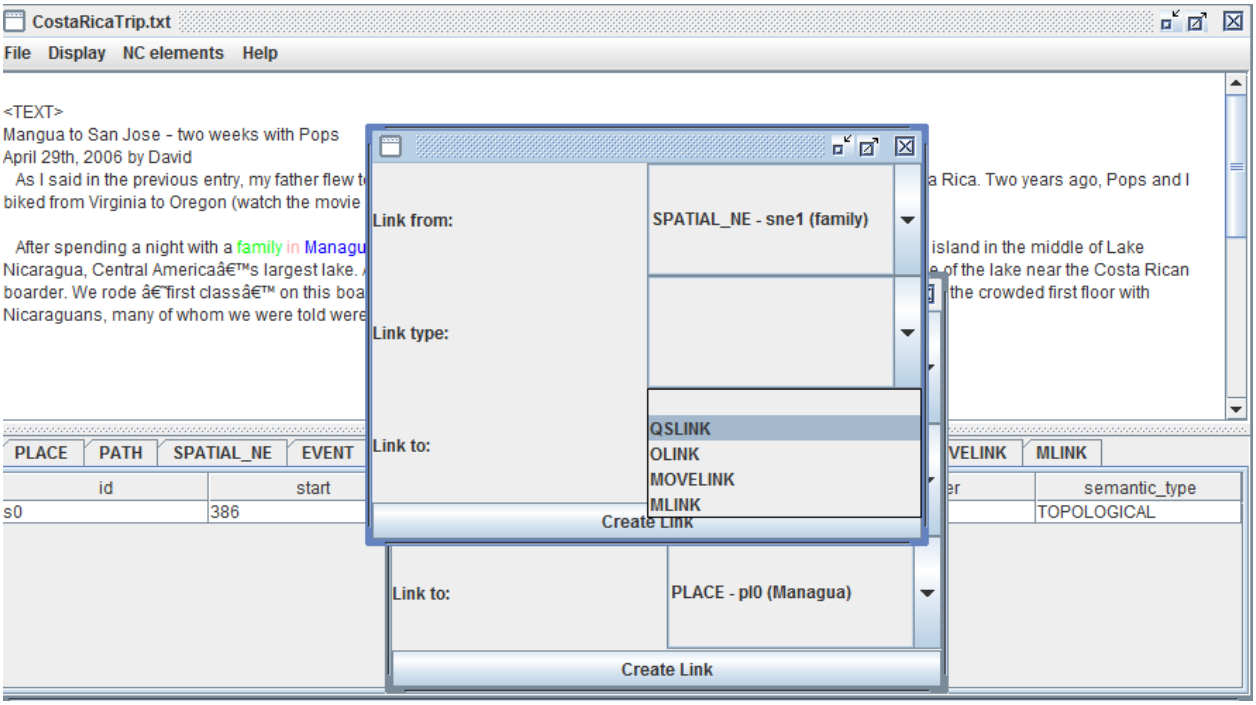

As you may have noticed, where the extents end up in the table for each link is not usually what the function that extent is; however, as long as the tagging is consistent this will not matter, as labels are simply labels. To finish the link, simply head to its entry and fill in the missing information. For most links you will likely need to provide more information about other attributes not provided by the first two extents (trigger, relType, etc.). Some of these attributes require the input of other tags (e.g. a trigger usually needs a spatial signal). Rather than go back through each of the tables of tags and track down each id, there is a way to find the id of the extent by simply highlighted the already tagged extent and right clicking on it. The pop-up menu should display the id of the tag preceded by "remove" (you do not need to remove the tag, just to find its id) as seen below. You will still need to manually enter in the id however.

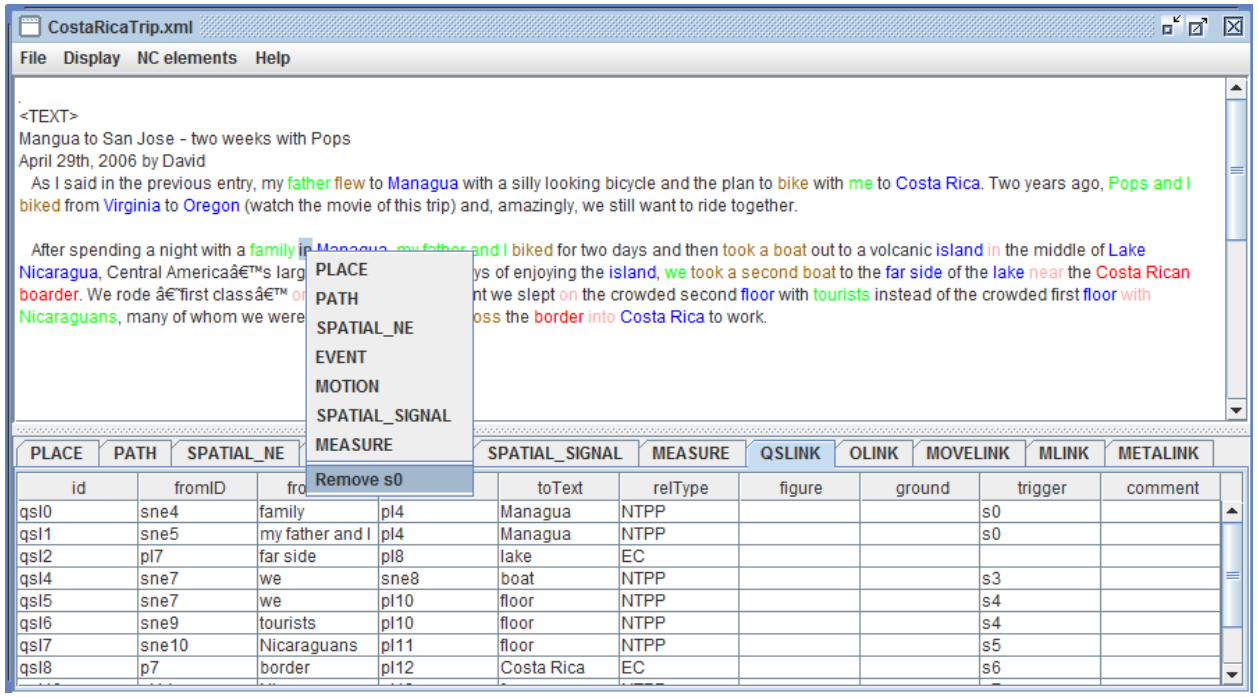

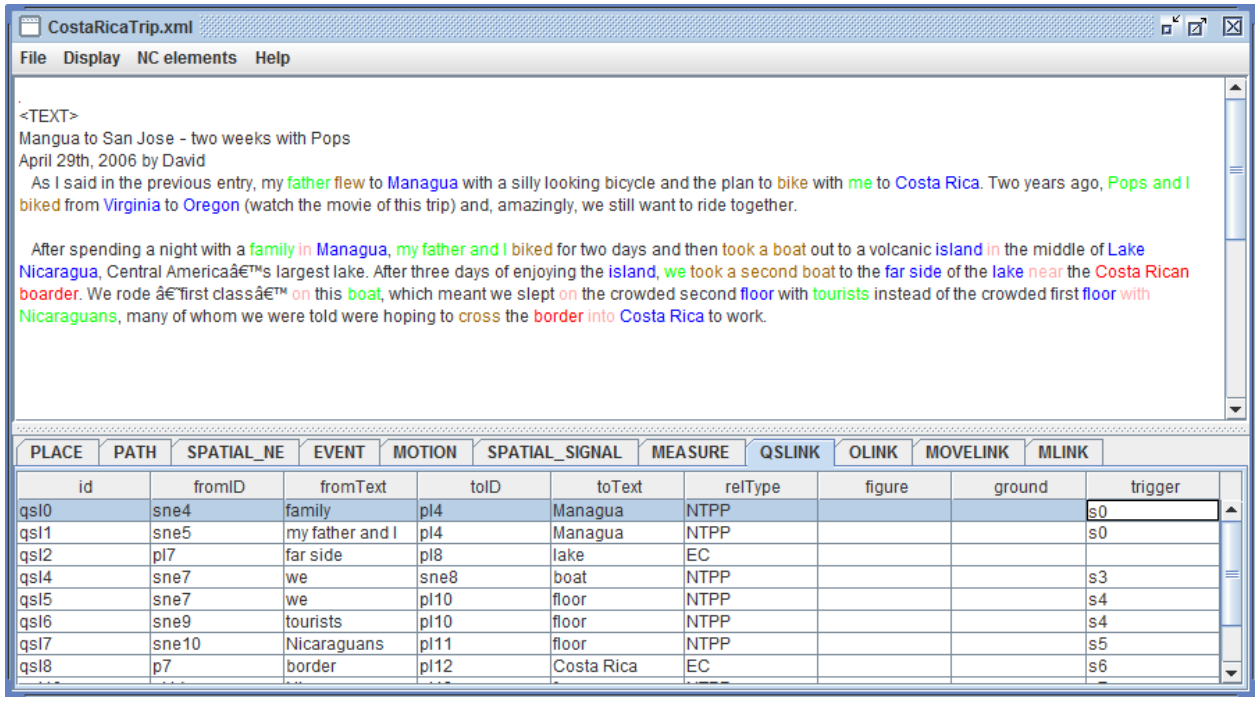

...And that's it! You are now ready to tag extents and set up links, which covers all the annotation work needed for ISO-Space.