Capturing Motion in ISO-SpaceBank

James Pustejovsky
Brandeis University
jamesp@cs.brandeis.edu

Zachary Yocum
Brandeis University
zyocum@brandeis.edu

Abstract

This paper presents the first description of the motion subcorpus of ISO-SpaceBank (MotionBank) and discusses how motion-events are represented in ISO-Space 1.5, a specification language for the representation of spatial information in language. We present data from this subcorpus with examples from the pilot annotation, focusing specifically on the annotation of motion-events and their various participants. These data inform further discussion of outstanding issues concerning semantic annotation, such as quantification and measurement. We address these questions briefly as they impact the design of ISO-Space.

1 Introduction

The goal of ISO-Space is to provide a specification of an annotation language for encoding spatial and spatiotemporal information as expressed in natural language texts. Section 2 enumerates the elements of syntax in ISO-Space 1.5. Section 3 presents data from the MotionBank pilot annoation effort (a subcorpus of ISO-SpaceBank). In the subsequent discussion we focus specifically on relations pertaining to motion, and discuss only limited aspects of topological, orientational, and measurement relations. Section 4 contains discussion of outstanding issues and how they may be tackled.

ISO-Space is being developed as a comprehensive foundation for the annotation of spatial information in natural language text. While there are clearly many issues remaining, we have attempted to follow a strict methodology of specification development, as adopted by ISO TC37/SC4 and outlined in (Bunt, 2010) and (Ide and Romary, 2004), and as implemented with the development of ISO-TimeML (Pustejovsky et al., 2005) and others in the family of SemAF standards.

As reported in (Pustejovsky et al., 2013), ISO-Space is designed to capture both spatial and spatiotemporal information. While still in development, it is clear that the conceptual inventory for spatial language annotation must at least include the following notions:

(1) a. Locations (regions, spatial objects): Geographic and geopolitical places.
b. Entities participating in spatial relations.
c. Paths: routes, lines, turns, arcs.
d. Topological relations: *in*, *connected*.
e. Direction and Orientation: *North*, *down*.
f. Time and space measurements: *20 miles away*, *for two hours*.
g. Object properties: intrinsic orientation, dimensionality.
h. Frames of Reference: absolute, intrinsic, relative.
i. Motion: tracking objects over time.

In the following discussion, we report on the annotation of motion-events and participants, as part of the developing ISO-SpaceBank corpus, and discuss the issues arising with incorporating movement within a spatial representation language.

2 ISO-Space 1.5

In this section, we present a brief description of the ISO-Space 1.5 specification. Note that examples are annotated only with those syntactic elements and attributes which are relevant to the discussion.

2.1 Location Tags

Place Tag The attributes for the PLACE tag are largely inherited from SpatialML (Mani et al., 2010), with some minor additions. This tag is used to annotate geographic entities like lakes and mountains, as well as administrative entities like towns and counties.

(2) a. I camped next to the municipal [**building**$_{pl1}$].
PLACE(id=pl1, form=NOM, dcl=FALSE, countable=TRUE)
b. I traveled north to northern [**Lago Maracaibo**$_{pl2}$].
PLACE(id=pl2, form=NAM, dcl=FALSE, countable=TRUE)

The `form` attribute distinguishes nominal forms (2a) from regions with proper names (2b).

The ISO-Space `mod` attribute is included here because it is substantially different from its counterpart in SpatialML (MITRE, 2007).[1] The ISO-Space `mod` attribute is intended to capture cases like *tall building*, *long trail*, or *the higher observation deck*, where *tall*, *long* and *higher* do not constrain the location of the entity but they do contribute spatial information.

ISO-Space locations tags includes a Document Creation Location or `dcl` attribute. The DCL is a special location that serves as the "narrative location". If a document includes a `dcl`, it is generally specified at the beginning of the text, similarly to the manner in which a Document Creation Time is specified in TimeML (Pustejovsky et al., 2005).

The `countable` attribute is used to distinguish regions referred to with countable sortals (*cities, lakes*) and mass sortals (*highlands, countryside*).

Path Tag The PATH tag is used to capture locations where the focus is on the potential for traversal or functions as a boundary. This includes common nouns as in (3a) and (3b), as well as proper names as in (3c). The attributes of the PATH tag are a subset of the attributes of the PLACE tag, but with the additional `beginID`, `endID`, and `midIDs` attributes. The PATH tag is intended to capture only noneventive paths, which are treated as inherently nondirectional. As such, the `beginID` and `endID` attributes simply indicate bounding points rather than directionality. Table 1 summarizes the attributes for the PATH tag.

Attribute	Value
id	p1, p2, p3, ...
beginID	ID of a location tag
endID	ID of a location tag
midIDs	list of IDs of midpoint locations
form	NAM or NOM
elevation	a MEASURE ID
mod	a spatially relevant modifier
countable	TRUE or FALSE
quant	a generalized quantifier

Table 1: PATH Tag Attributes.

(3) a. ... I arrived at the end of the [**road**$_{p1}$].
b. ... a massive mountain [**range**$_{p2}$] that hugs the west [**coast**$_{p3}$] of Mexico.
c. I followed the [**Pacific Coast Highway**$_{p4}$] along the coastal mountains ...

Non-Consuming Location Tags It is often useful to identify locations that are not mentioned explicitly in the text. In such cases, ISO-Space allows for non-consuming location tags. For example, a non-consuming PLACE tag would be necessary in the case of *John climbed to 9,000 feet* where the elevation *9,000 feet* indirectly references a location that is not associated with any extent in the text.

2.2 Non-Location Tags

While location tags essentially designate a region of space that can be related to other regions on space, ISO-Space allows for non-location elements of a text to be coerced into behaving like a region of space so that they may participate in the same kinds of relationships. There are three of these kinds of non-location tags that may behave like locations in ISO-Space: SPATIAL_E, EVENT and MOTION.[2]

Spatial Entity The SPATIAL_E (spatial entity) tag is intended to capture any entity that is both located in space and participates in an ISO-Space link tag, as illustrated in (4). Attributes include: `id`, `form`, `mod`, `countable`, and `quant`.

(4) [**David**$_{se1}$] passed three [**cars**$_{se2}$] on the road.

[1]Given this discrepancy with SpatialML, it is likely that the ISO-Space annotator will have to perform some "clean-up" of the PLACE elements that are inherited from a SpatialML annotation. This issue will be taken up in the annotation guidelines, though, as it is not relevant to this specification.

[2]Note that, depending on the annotation task, annotating these tags may not be the responsibility of the ISO-Space annotator. Instead, capturing this kind of information may be left to other annotation schemes and it will be left to the ISO-Space annotator to recognize when such an element should participate in an ISO-Space link tag.

Event The EVENT tag captures events that do not involve a change of location but are directly related to another ISO-Space element by way of a link. Events are inherited directly from the ISO-TimeML annotation scheme (Pustejovsky et al., 2005) and require no further specification in ISO-Space.

Spatial Signal The SPATIAL_SIGNAL tag captures relation words or phrases that supply information to an ISO-Space link tag. Signals are typically prepositions or other function words that specify the particular relationship between two ISO-Space elements. Attributes include: `id`, `cluster`, and `semantic_type`.

Adjunct The ADJUNCT tag captures additional *event-path* or *manner-of-motion* information that is not contributed directly by a motion verb, but rather by a satellite word or phrase. PATH motion adjuncts are often prepositions (e.g, *to* and *from*). Adjuncts of type MANNER supply manner of motion information (e.g., *by car*). Notice in (5d) that multiple adjuncts may contribute to a single motion.

(5) a. John walked [**to**$_{a1}$] the store.
b. John left [**for**$_{a2}$] Boston.
c. John traveled [**by car**$_{a3}$].
d. John arrived [**by bike**$_{a4}$] [**at**$_{a5}$] the trailhead.

Measure The MEASURE tag is used to capture distances and dimensions for use in an MLINK or to fill the `elevation` attribute for a location tag. See (Pustejovsky et al., 2013) for more details.

2.3 Spatial Relation Links

There are four relationship tags in ISO-Space defined as follows:

(6) a. QSLINK – for qualitative spatial relations;
b. OLINK – for orientation relations;
c. MOVELINK – for movement relations;
d. MLINK – for dimensions of a region or the distance between locations.

Qualitative Spatial Link QSLINKs are used in ISO-Space to capture topological relationships between tag elements captured in the annotation. The `relType` attribute values come from an extension to the RCC8 set of relations that was first used by SpatialML. The possible RCC8+ values include the RCC8 values (Randell et al., 1992), in addition to IN, a disjunction of TPP and NTPP (cf. Table 2).

Relation	Description
DC	Disconnected
EC	External Connection
PO	Partial Overlap
EQ	Equal
TPP	Tangential Proper Part
TPP$_i$	Inverse of TTP
NTTP	Non-Tangential Proper Part
NTTP$_i$	Inverse of NTTP
IN	Disjunction of TTP and NTTP

Table 2: RCC8+ Relations.

It is worth noting that while the QSLINK tag is used exclusively for capturing topological relationships, which are only possible between two regions, the `figure` and `ground` attributes can accept IDs for both PLACEs and PATHs, which are more traditional regions, as well as SPATIAL_ES, EVENTs, and MOTIONs. In the latter cases, it is actually the region of space that is associated with the location of the entity or event that participates in the QSLINK. That is, the entity or event is coerced to a region for the purposes of interpreting this link.

In practice, a QSLINK is triggered by a SPATIAL_SIGNAL with a `semantic_type` of TOPOLOGICAL or DIR_TOP (cf. (7) below).

(7) [**The book**$_{se1}$] is [**on**$_{s1}$] [**the table**$_{se2}$].
SPATIAL_SIGNAL(`id=s1`, `cluster=`"on-1", `semantic_type=` DIR_TOP)
QSLINK(`id=qsl1`, `figure=sne1`, `ground=sne2`, `trigger=s1`, `relType=`EC)

Orientation Link Orientation links describe non-topological relationships. A SPATIAL_SIGNAL with a DIRECTIONAL `semantic_type` triggers such a link. In contrast to qualitative spatial relations, OLINK relations are built around a specific frame of reference type and a reference point. The attributes for OLINK are listed in Table 3.

The `referencePt` value depends on the `frame_type` of the link. The ABSOLUTE frame type stipulates that the `referencePt` is a cardinal direction. For INTRINSIC OLINKs, the `referencePt` is the same identifier that is given in the `ground` attribute. For RELATIVE OLINKs, the identifier for the viewer should be provided as to the `referencePt`. If the viewer is not explicit in the text, the special value "VIEWER" should be used. Examples of this link are illustrated in (8).

(8) a. [**Boston**$_{pl1}$] is [**north of**$_{s1}$] [**New York City**$_{pl2}$].

Attribute	Value
id	ol1, ol2, ol3,...
relType	ABOVE, BELOW, FRONT, NORTH,...
figure	ID of the location/entity/event that is being related to the ground
ground	ID of the location/entity/event that is being related to by the figure
trigger	ID of a SPATIAL_SIGNAL that triggered the link
frame_type	ABSOLUTE, INTRINSIC or RELATIVE
referencePt	ground location/entity/event ID, cardinal direction, or viewer entity ID
projective	TRUE or FALSE

Table 3: OLINK Attributes.

OLINK(id=ol1, figure=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, figure=sne1, ground=sne2, trigger=s2, relType="FRONT", frame_type=INTRINSIC, referencePt=sne2, projective=FALSE)

Measure Link Measurement relationships are captured with the MLINK tag, as first proposed for ISO-TimeML (Pustejovsky et al., 2010). Currently, this tag describes either the relationship between two spatial objects or the dimensions of a single object (cf. Table 4).

Attribute	Value
id	ml1, ml2, ml3,...
figure	ID of the location/entity/event event that is being related to the ground
ground	ID of the location/entity/event that is being related to by the figure
relType	DISTANCE, LENGTH, WIDTH, HEIGHT, or GENERAL_DIMENSION
val	a MEASURE ID or NEAR, FAR, TALLER, SHORTER,
endPoint1	ID of a location/entity/event at one end of a stative path
endPoint2	ID of a location/entity/event at the other end of a stative path

Table 4: MLINK Attributes.

When an MLINK is used to capture an internal dimension of an object as in (9b) or (9c), the ID of that object should appear in the figure attribute. The annotator may either repeat the identifier in the ground attribute or leave the ground unspecified.

(9) a. The new [**tropical depression**$_{se1}$] was about [**430 miles**$_{me1}$] ([**690 kilometers**$_{me2}$]) west of the [**southernmost Cape Verde Island**$_{pl1}$], they said.
MLINK(id=ml1, relType=DISTANCE, figure=sne1, ground=pl1, val=me1)

b. [**The football field**$_{se2}$] is [**100 yards**$_{me2}$] long.
MLINK (id=ml2, relType=LENGTH, figure=sne2, ground=sne2, val=me2)

c. I [**rode**$_{m1}$] [**30 miles**$_{me4}$] yesterday.
MLINK (id=ml6, relType=general_dimen, figure=m1, ground=m1, val=me4)

2.4 Movement

The treatment of movement in ISO-Space draws heavily from the foundations of lexical semantics in (Talmy, 1985) and the motion-event classifications in (Muller, 1998) and (Pustejovsky and Moszkowicz, 2008). There are two ISO-Space tags which capture movement: MOTION and MOVELINK.

Motion Tag The ISO-Space MOTION tag is a species of TimeML event that involves a change of location or spatial configuration. Table 5 lists the attributes of the MOTION tag.

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, STAY
motion_sense	LITERAL, FICTIVE, INTRINSIC_CHANGE

Table 5: MOTION Tag Attributes.

The motion_type attribute refers to the two major strategies for expressing movement in language: *path* and *manner-of-motion* constructions (Talmy, 1985). This is illustrated in (10), where m indicates a manner contributing component, and p indicates a path contributing component. In the first sentence, the motion verb specifies a path whereas in the second the motion verb specifies the manner of motion. The motions in these sentences are actually of the motion_type COMPOUND since they supply both path and manner information.

(10) a. John arrived$_p$ [by foot]$_m$.
b. John hopped$_m$ [out of the room]$_p$.

Motion classes are taken from (Pustejovsky and Moszkowicz, 2008), which in turn are based on those in (Muller, 1998). These classes are associated with a spatial event structure that specifies the spatial relations between the arguments of the motion verb at different phases of the event. Table 6 lists the set of motion classes and their associated motion-event structures.

The `motion_sense` attribute distinguishes between different kinds of interpretations of motion-events. The LITERAL sense covers motion-events where the mover participant's location changes over time. The FICTIVE sense covers cases where the event involves an atemporal, experiential change in an extrinsic property (e.g., elevation or location). The INTRINSIC_CHANGE sense covers motion verbs that describe change in some intrinsic, spatial characteristic (e.g., height, width, length, shape, etc.). The motivation here is to disambiguate language like *the balloon rose above the building* from *the river rose above the levy*, where a LITERAL interpretation—the river's elevation increased—is inappropriate: the location of the eleavtion of the river is supervenient on the change in the volume of the river, therefore signaling an intrinsic change.[3] The `motion_sense` attribute also captures FICTIVE motion interpretations such as, *the mountain rises above the valley*, where there is no temporal interpretation—the mountain's elevation increasing over time—but rather a purely spatial, atemporal interpretaion predicating spatial characteristics of the mountain over some region.

Movelink Tag MOVELINK tags, which are introduced by MOTION tags, capture information about the path or course a particular motion takes. Table 7 lists the attributes of the MOVELINK link.

The event structures for MOVE_EXTERNAL and MOVE_INTERNAL motion-events require a `ground` location relative to which the motion of the `mover` participant occurs. This location is identified with the `ground` attribute introduced in Table 7 and its use is demonstrated in Example (11a).

Another attribute introduced in Table 7 is `adjunctID`. This attribute takes the identifier of an ATTRIBUTE tag that contributes path or manner information about the event-path of the MOVELINK's triggering motion-event. The use of the `adjunctID` attribute is demonstrated in Example (11b)

[3]While this could be an instance of a metonymic sense extension, such as *the kettle boiled* (per a reviewer's suggestion), we believe this is more specific to the entailments associated with an intrinsic change in an object's spatial extent.

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 whose whose location changes
`ground`	ID of a location/entity/event tag that the `mover`'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 triggering MOTION
`adjunctID`	IDs of any ADJUNCT tags that contribute path or manner information to the triggering MOTION

Table 7: MOVELINK Tag Attributes.

(11) a. ... [**we**$_{se1}$] [**passed**$_{m1}$] [**glaciers**$_{p1}$] and [**snowfields**$_{pl1}$] ...
SPATIAL_E (`id=sne1, form=`NOM, `countable=`TRUE)
MOTION (`id=m1, motion_type=`PATH, `motion_class=`MOVE_EXTERNAL, `motion_sense=`LITERAL)
MOVELINK (`id=mvl1, trigger=m1, mover=sne1, ground=p1`)
MOVELINK (`id=mvl2, trigger=m1, mover=sne1, ground=pl1`)

b. [**I**$_{se2}$] [**biked**$_{m2}$] [**into**$_{a1}$] a [**town**$_{pl2}$] at 4pm.
SPATIAL_E (`id=sne2, form=`NOM, `countable=`TRUE)
MOTION (`id=m2, motion_type=`COMPOUND, `motion_class=`REACH, `motion_sense=`LITERAL)
MOVELINK (`id=mvl3, trigger=m2, goal=pl2, mover=sne2, goal_reached=yes, adjunctID=a1`)

2.5 Annotation vs. Axioms

It is important to note that ISO-Space's inventory of explicit representations does not capture the whole picture. Some representations are introduced at the level of abstract syntax by specific axiomatic rules. We introduce the assumed premises for motion briefly, and defer details to the final paper.

Mover Participants The first axiom pertaining to motion in ISO-Space is that, for every motion-event,

Value	Requisite Attributes	Event Structure
MOVE	mover	$_{begin}[location\_of(\mathtt{mover})] \not\sim\ _{end}[location\_of(\mathtt{mover})]$
MOVE_EXTERNAL	mover, ground	$_{begin...end}[\{\mathrm{DC} \wedge \mathrm{EC}\}(\mathtt{mover}, \mathtt{ground})]$
MOVE_INTERNAL	mover, ground	$_{begin...end}[\mathrm{IN}(\mathtt{mover}, \mathtt{ground})]$
LEAVE	mover, source	$_{begin}[\mathrm{IN}(\mathtt{mover}, \mathtt{source})],\ _{end}[\{\mathrm{DC} \wedge \mathrm{EC}\}(\mathtt{mover}, \mathtt{source})]$
REACH	mover, goal	$_{begin}[\mathrm{DC}(\mathtt{mover}, \mathtt{goal})],\ _{end}[\mathrm{IN}(\mathtt{mover}, \mathtt{goal})]$
DETACH	mover, source	$_{begin}[\mathrm{EC}(\mathtt{mover}, \mathtt{source})],\ _{end}[\mathrm{DC}(\mathtt{mover}, \mathtt{source})]$
HIT	mover, goal	$_{begin}[\mathrm{DC}(\mathtt{mover}, \mathtt{goal})],\ _{end}[\mathrm{EC}(\mathtt{mover}, \mathtt{goal})]$
FOLLOW	mover, pathID	$_{begin...end}[path\_of(\mathtt{mover}) \sim \mathtt{pathID}]$
DEVIATE	mover, pathID	$_{begin}[path\_of(\mathtt{mover}) \sim \mathtt{pathID}],\ _{end}[path\_of(\mathtt{mover}) \not\sim \mathtt{pathID}]$
CROSS	mover, source, midPoints, goal	$_{begin}[\mathrm{IN}(\mathtt{mover}, \mathtt{source})],\ _{mid}[\mathrm{IN}(\mathtt{mover}, \mathtt{midPoints})],\ _{end}[\mathrm{IN}(\mathtt{mover}, \mathtt{goal})]$
STAY	mover, ground	$_{begin...end}[\{\{\mathrm{RCC8+}\}, \{\mathrm{OLINK}\}\}(\mathtt{mover}, \mathtt{ground})]$

Table 6: Motion Class Event Structures

there exists an entity which fulfills the role of mover for that event. The mover is that participant in the motion-event which undergoes a change in its location. That is to say:

(12) $\forall e \exists x[motion\text{-}event(e) \rightarrow mover(x, e)]$

Event Paths The other essential component of ISO-Space that is generated axiomatically is the event-path created by the mover associated with a motion-event. That is to say:

(13) $\forall e \exists p[motion\text{-}event(e) \rightarrow [event\text{-}path(p) \wedge loc(e, p)]]$

Previous versions of the ISO-Space specification included an event-path tag as part of the concrete syntax, distinct from the non-eventive PATH tag. In fact, the source, goal, midPoint and pathID attributes of the MOVELINK tag presume an event-path (although these attributes are often underspecified). The primary motivation for the removal of event-paths as their own category in the concrete syntax is that our abstract syntax axiomatically introduces an event-path for each motion-event.[4]

This decision simplifies the annotation task in that annotators need only identify features of the event-path if the language contributes information about the path of traversal. A bare-manner motion verb, as in *David* ***cycles*** *seriously*, for instance, introduces a completely underspecified event-path. Thus, the following annotation in 14 would be sufficient.

(14) [**David**$_{se1}$] [**cycles**$_{m1}$] seriously.
SPATIAL_E (id=sne1, text="David", form=NAM)
MOTION (id=m1, text="cycles", motion_type=MANNER, motion_class=MOVE, motion_sense=LITERAL)
MOVELINK (id=mvl1, trigger=m1, source=∅, goal=∅, midPoint=∅, mover=sne1, ground=∅, goal_reached=∅, pathID=∅, adjunctID=∅)

[4]Discussions from participants at ISA-7 and ISA-8 were instrumental in leading to this modification in the specification.

3 ISO-SpaceBank Subcorpus Data

The data in this section are tabulated from the pilot annotation of MotionBank, a subcorpus of ISO-SpaceBank consisting of 50 entries (20,877 word tokens) from a travel blog whose author cycled across the Americas. Table 8 presents a breakdown of the tag counts for each ISO-Space tag type. Table 9 lists the counts for each class of motion over the same subcorpus by frequency.

Tag Type	Frequency
PLACE	1313
SPATIAL_E	856
MOVELINK	834
MOTION	794
SPATIAL_SIGNAL	558
ADJUNCT	407
PATH	294
EVENT	186
total	**5308**

Table 8: Tag Counts

To best illustrate the annotation of motion and the various participants, we present one detailed example in full. Sentence (15) is spatially quite rich and it is also notable for the figurative language that is employed. The first item of note is the non-consuming place tag that has been created. In this case the MEASURE ID of *over 6,000 feet* fills the elevation attribute of the non-consuming place tag. The ID of this PLACE tag is then used later to fill the goal location for the MOVELINK triggered by m3 (*climbs*).

The second thing to note is that the motion_sense attributes for all the MOTION

Motion Class	Frequency
MOVE	183
REACH	177
STAY	130
HIT	62
LEAVE	56
FOLLOW	54
CROSS	54
MOVE_INTERNAL	39
MOVE_EXTERNAL	26
DETACH	11
DEVIATE	2
Total	**794**

Table 9: Motion Class Counts

tags are FICTIVE. This is because the *road* is fulfilling the role of `mover` and the annotator assumed figurative, atemporal interpretations for the *Departing*, *climbs*, and *climb* motion-events.

(15) a. [**Departing**$_{m2}$] [**Copala**$_{pl11}$], the [**road**$_{p1}$] [**climbs**$_{m3}$] [**to**$_{a1}$] [**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, text="Copala", form=NAM, dcl=FALSE, num=SING)
PLACE (id=pl12, text=∅, elevation=me5, dcl=FALSE, num=SING)
PATH (id=p1, midIDs={pl11, pl12}, form=NOM)
PATH (id=p2, text="ridgeline", form=NOM, countable=TRUE)
MEASURE (id=me5, text="over 6,000 feet", value="gt 6000", unit="feet")
MEASURE (id=me6, text="30 miles", value="30", unit="feet")
MLINK (id=ml5, figure=m3, GROUND=m3, relType=GENERAL_DIMENSION, val=m6, endPoint1=pl11, endPoint2=pl12)
MOTION (id=m2, text="Departing", motion_type=PATH, class=LEAVE, motion_sense=FICTIVE)
MOVELINK (id=mvl2, trigger=m2, source=pl11, mover=p1, pathID=p1)
MOTION (id=m3, text="climbs", class=MOVE, motion_sense=FICTIVE)
ADJUNCT (id=a1, text="to", type=PATH)
MOVELINK (id=mvl3, trigger=m3, source=pl11, goal=pl12, mover=p1, goal_reached=TRUE, pathID=p1, ajdunctID=a1)
MOTION (id=m4, text="climb", class=MOVE, motion_sense=FICTIVE)
MOVELINK (id=mvl4, trigger=m4, source=pl12, mover=p1, pathID=p1)
SPATIAL_SIGNAL (id=s8, text="hugging", semantic_type=DIR_TOP)
QSLINK (id=qsl8, relType=DC, figure=p1, ground=p2, trigger=s8)

4 Discussion

Several interesting issues arose during the initial motion annotation efforts with ISO-Space. The first concerns how to handle 'simulated' motion-events. Such events are the kind typical in direction-giving language where a direction-giver may specify a path that is intended to be followed without explicitly specifying a mover participant: *Walk 100 meters and turn right after the store*. Initially, this was dealt with by providing an additional `motion_sense` value, called SIMULATED, in order to distinguish such uses from the FICTIVE, LITERAL, and INTRINSIC_CHANGE motion senses. After further corpus investigation, however, we have determined that this is a *narrative modality* rather than a specific sense distinction exploited for motion verbs. This deserves further modeling and we are currently investigating this topic.

Another issue that arises, although interestingly, not represented in the present corpus, involves the use of *extent verbs* (Gawron, 2009). This use is seen in the following: *Past the brook, the road* ***narrows***. This shares semantic elements with the FICTIVE sense, but introduces additional constraints not accompanying those uses (as in *the road climbs*, etc.). This is also currently under further investigation.

It is worth pointing out that quantification presents itself again as an issue. ISO-Space 1.4 provides `countable` and `quant` attributes for location tags, however these features alone remain insufficient for a complete motion-events semantics. Consider (16), for instance. The annotation captures the quantification over *valley* with the PATH tag `p1`, and the MOVELINK (`mvl1`) triggered by *passed* (`m1`) specifies `p1` as a `midPoint` location.

(16) a. ... [**I**$_{se1}$] [**passed**$_{m1}$] through every small, uninhabited [**valley**$_{p1}$] [$\varnothing_{pl1}$] [$\varnothing_{pl2}$][5] ...
SPATIAL_E (id=sne1, text="I", form=NOM)
path (id=p1, text="valley", form=NOM, mod="small", quant="every")

[5]The symbol ∅ is used here to identify non-consuming tags in the text.

MOTION (`id=m1`, `text=`"passed", `motion_type=`PATH, `motion_class=`CROSS, `motion_sense=`LITERAL)

MOVELINK (`id=mvl1`, `trigger=m1`, `source=pl1`, `goal=pl2`, `midPoint=p1`, `mover=sne1`, `goal_reached=`TRUE)

For a proper semantic interpretation, it is essential to produce an interpretation for this sentence where `m1` falls under the scope of the quantifier *every*. That is, for *every valley*, there exists a *passing* motion-event. A partial translation is as follows, where $through$ is a stand in for the appropriate QSLINK relation value.

(17) $\forall p_1 \exists m_1[[valley(p_1) \wedge small(p_1) \rightarrow [pass(m_1) \wedge through(m_1, p_1)]]$

In addressing this issue, ISO-Space 1.5 draws from TimeML's treatment of event quantification in (Bunt and Pustejovsky, 2010; Pustejovsky et al., 2010), to handle examples such as *John taught every Tuesday*. ISO-TimeML captures quantificational scoping relations with a $scopes(scoper, scopee)$ relation. We propose to extend the tag attributes in the ISO-Space with a `scopes` attribute to capture such relations.

Finally, another desideratum that has been made evident by the pilot annotation data is the ability to capture motion when it occurs in nominal form. That is not to say that all motion-event nominals ought to be treated as instances of motion. For example, while *a **vacation** to Mexico* seems to entail travel, *a summer **vacation*** may not. Additionnally, the participants of motion-event nominals are often underspecified. The pilot annotation guidelines did not sufficiently address the possibility of underspecified mover participants, and consequently, the EVENT tag was employed for nominalized motion-events. Examples from MotionBank where this confusion occurred are italicized in the sentences in Example (18).

(18) a. The last few days of the *trip* were difficult, including an 8,000 feet *climb* into the Andes.
b. According to Ricardo, bicycle *use* has increased 5 times in the city, and now there are probably between 300,000 and 400,000 *trips* made daily in Bogota by bicycle.
c. Passing through more towns and more *climbs* and *descents* on one lane dirt roads, I eventually climbed into the Cordillera Blanca ...
d. I also received a *tour* of the town from three high school students ...
e. I have now arrived in Yurimaguas, a small city in the jungle, thus ending my two weeks of boat *travel* on the world's largest river system.
f. Many people I have stayed with on this *trip* live in small houses, are poor, own no car, and have little healthcare.

5 Conclusion

In this paper we have presented an initial description of the motion subcorpus of the ISO-Space specification for spatiotemporal and spatial markup of natural language text. Through this discussion we hope to vet some of the remaining issues we have encountered with annotating movement phenemona in natural language. Our expectation is to release the completed MotionBank sub-corpus in June 2013 and subsequently the full SpaceBank corpus in January 2014.

Acknowledgements

This research was supported by grants from the NSF (NSF-IIS 1017765) and the NGA (NURI HM1582-08-1-0018). We would like to thank Jessica Moszkowicz, Marc Verhagen, Harry Bunt, and Kiyong Lee for their contributions to this discussion. We would also like to acknowledge the four anonymous reviewers for their helpful comments. All errors and mistakes are, of course, the responsibilities of the authors.

References

Harry Bunt and J. Pustejovsky. 2010. Annotating temporal and event quantification. In *Proceedings of 5th ISA Workshop*.

Harry Bunt. 2010. A methodology for designing semantic annotation languages exploiting syntactic-semantic iso-morphisms. In *Proceedings of ICGL 2010, Second International Conference on Global Interoperability for Language Resources*.

J.M. Gawron. 2009. The lexical semantics of extent verbs. *San Diego State University*.

Nancy Ide and Laurent Romary. 2004. International standard for a linguistic annotation framework. *Natural Language Engineering*, 10(3-4):211–225.

Inderjeet Mani, Christy Doran, Dave Harris, Janet Hitzeman, Rob Quimby, Justin Richer, Ben Wellner, Scott Mardis, and Seamus Clancy. 2010. Spatialml: annotation scheme, resources, and evaluation. *Language Resources and Evaluation*, 44:263–280. 10.1007/s10579-010-9121-0.

MITRE. 2007. Spatialml: Annotation scheme for marking spatial expressions in natural language. http://sourceforge.net/projects/spatialml/.

Philippe Muller. 1998. A qualitative theory of motion based on spatio-temporal primitives. In Anthony G. Cohn, Lenhart Schubert, and Stuart C. Shapiro, editors, *KR'98: Principles of Knowledge Representation and Reasoning*, pages 131–141. Morgan Kaufmann, San Francisco, California.

James Pustejovsky and Jessica L. Moszkowicz. 2008. Integrating motion predicate classes with spatial and temporal annotations. In *Proceedings of COLING 2008*, Manchester, UK.

James Pustejovsky, Robert Knippen, Jessica Littman, and Roser Saurí. 2005. Temporal and event information in natural language text. *Language Resources and Evaluation*, 39:123–164, May.

James Pustejovsky, Kiyong Lee, Harry Bunt, and Laurent Romary. 2010. Iso-timeml: An international standard for semantic annotation. In *LREC*.

James Pustejovsky, Jessica Moszkowicz, and Marc Verhagen. 2013. A linguistically grounded annotation language for spatial information. Special issue of TAL. Forthcoming.

David Randell, Zhan Cui, and Anthony Cohn. 1992. A spatial logic based on regions and connections. In Morgan Kaufmann, editor, *Proceedings of the 3rd Internation Conference on Knowledge Representation and Reasoning*, pages 165–176, San Mateo.

Leonard Talmy. 1985. Lexicalization patterns: semantic structure in lexical forms. In T. Shopen, editor, *Language typology and semantic description Volume 3: Grammatical categories and the lexicon*, pages 36–149. Cambridge University Press.