Representing ISO-Annotated Dynamic Information in UMR

Kiyong Lee
Korea University, Seoul
ikiyong@gmail.com

Harry Bunt
Tilburg University, NL
harry.bunt@tilburguniversity.edu

James Pustejovsky
Brandeis University, Waltham
jamesp@brandeis.edu

Alex C. Fang
City University of Hong Kong
alex.fang@cityu.edu.hk

Chongwon Park
University of Minnesota Duluth
cpark2@d.umn.edu

Abstract

The ISO working group on semantic annotation aims to adopt the UMR formalism to represent dynamic information involving motions and their embedding grounds. The paper details how ISO's XML-based temporal and spatial annotations, involving motions and spatio-temporally conditioned event-paths, will be converted to AMR or UMR forms. It also attempts to enrich the representation of dynamic information with the integrated spatio-temporal annotation scheme that accommodates first-order dynamic logic, as briefly noted. The main motivation of such an effort is to make spatio-temporal annotations and related dynamic information easily understandable by artificial agents like robots to act. Our approach bridges ISO's richly specified standards with the task-oriented expressiveness of UMR and dynamic logic. This integration paves the way for seamless downstream use of spatio-temporal annotations in dialogue systems, simulation environments, and embodied agents.

Key Words: dynamic information, dynamic space, embedding ground, motion, spatio-temporal annotation, UMR

1 Introduction

We propose and explore the use of UMR in ISO's new project on *motion in dynamic space* (ISO/PWI 24617-18). Given Pustejovsky et al. (2019)'s use of AMR, the adoption of AMR for ISO's annotation standards is not novel. Furthermore, the adoption of AMR or UMR has been motivated by the rapid rise in their use in computational linguistics over the past decade; they simplify computational annotation processes while maintaining scalability, unencumbered by extensive syntactic pre-analysis.

As pointed out in Pustejovsky et al. (2019), the strength of AMR lies in its focus on the *predicative* core of a sentence while presenting an intuitive representation for semantic interpretation. More importantly, treating predicates as the root of each AMR structure facilitates annotation processes, just as the event-based temporal annotation of ISO-TimeML and the motion-based spatial annotation of ISO-Space are anchored to eventuality and motions, respectively.

The proposed project's scope for annotating motions embedded in spatio-temporal domains encompasses motions, space, time, and the embedding ground of a motion, called *dynamic space*. We aim to enrich this annotation scheme by augmenting the categorization of spatial and temporal entities with first-order dynamic logic and an iterative program procedure.

The paper will develop as follows. We discuss representing semantic annotations of language in Section 2. In Section 3, we demonstrate how ISO's dual annotation structures are represented in UMR. Section 4 introduces Spatio-Temporal Markup Language (Pustejovsky and Moszkowicz, 2011) and Generative Lexicon-based AMR (GLAMAR) (Tu et al., 2024) to treat motion-oriented dynamic information with the notion of sub-events. The dynamic logic formulates constraints on the iterative process of motions. The paper ends with concluding remarks.

2 Representing Semantic Annotations of Language

2.1 Abstract Annotation Scheme vs Concrete Physical Representation Format

Following Bunt (2010), the ISO SemAF group has divided the specification of each annotation scheme into two sub-components. The first sub-component *abstract syntax* formally defines the annotation structures of the scheme in abstract (set-theoretic) terms while reflecting its conceptual design based on a metamodel. In contrast, the other sub-component, *concrete syntax*, has adopted XML as the physical format for representing annotation

Proceedings of the 6th International Workshop on Designing Meaning Representations, pages 49–58
August 4, 2025.

structures. As depicted in Figure 1, a variety of concrete syntaxes is possible for representing annotation structures. Still, each of them must conform to the proposed abstract syntax while ideally retaining their logical equivalence. Hence, each concrete specification of representing annotation structures depends totally on the abstract syntax of an annotation language.

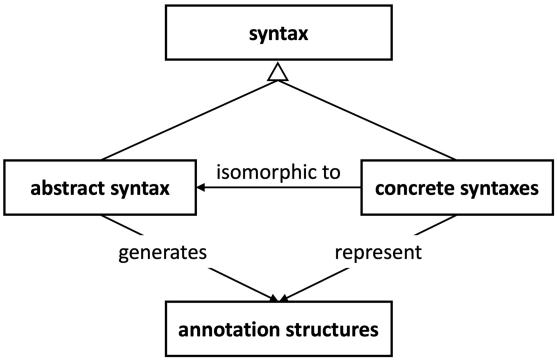

Figure 1: Syntax of an Annotation Language: Abstract vs. Concrete (Lee, 2023)

While introducing the two ISO standards, ISO-TimeML and ISO-Space, Pustejovsky (2017a) and Pustejovsky (2017b) have adopted two different representation formats. XML was adopted to represent annotation structures in ISO-TimeML, but a predicate-logic-like format was adopted in ISO-Space. Nevertheless, the representation of annotation structures in both representation formats conforms to their respective abstract specifications (syntaxes) of temporal and spatial annotations.

Example (refexTS briefly shows how they represent annotation structures.

(1) a. Data with categorized identifiers:
John$_{se1}$ left$_{e1/m1}$ Boston$_{pl1}$ yesterday$_{t1}$.

b. ISO-TimeML (Pustejovsky, 2017a):
```
<EVENT id="e1" target="w2"
pred="LEAVE" tense="PAST"/>
<TIMEX3 id="t1" target="w4"
type="DATE" value="2025-02-16"/>
<TLINK eventID="e1"
relatedToTime="t1"
relType="IS_INCLUDED"/>
```

c. ISO-Space (Pustejovsky, 2017b):
```
SPATIAL_ENTITY(id=se1,
type=PERSON,form=NAM)
MOTION(id=m1, target=w2,
motion_class=LEAVE, tense=PAST)
PLACE(id=pl1, target=w3,
cvt=CITY, form=NAM)
MOVELINK(id=mvli, trigger=m1,
mover=se1, source=pl1)
```

Both ISO-TimeML and ISO-Space focus on predicates, which can be either events or motions. `TLINK` relates the event of leaving to the time *yesterday*. Triggered by the motion *left*$_{m1}$, `MOVELINK` relates the spatial entity *John*$_{se1}$ to the source *Boston*$_{pl1}$.

2.2 UMR as a New Representation Format

UMR adopts the AMR formalism but extends its sentence-level representation to the document level (UMR, 2022). Consider first the sentence-level representation as in Example 2.

(2) a. Data:
```
(s / sentence
(The man left Boston yesterday
before it rained.))
```

b. AMR Format:
```
(l / leave-01
  :ARG0 (m / man)
  :source (b / Boston)
  :temporal (y / yesterday)
  :temporal (b1 / before
    :op1 (r / rain))
```

The AMR formalism represents abstract semantic concepts and relations that include event participant roles, such as `ARG0` or `actor`. In the AMR format, as in (2b) above, the slash (/) indicates semantic *concepts* while the colon (:) indicates a value of a semantic *relation*. In addition to argument roles, these relations form triplets bound to a governing concept (e.g., `l / leave-01 :ARG0 (m / man)`).

UMR then adds a document-level representation to the sentence-level representation. For example, the sentence-level representation can be extended to a document-level representation such as Example 3 be added:

(3) UMR Document-level Representation
```
(s / sentence)
(d / document-level
   :temporal (sr :before sl))
```

Linked to the sentence-level representation (2), the document-level representation (3) relates the rain event sr to the event of John's departure sl, interpreted as stating that John's departure occurred before the rain.

3 Representing ISO's Dual Annotation Structures in UMR

3.1 Dual Structures of Annotation

ISO's SemAF annotation schemes formally define annotation structures, each divided into two substructures: *entity structures* and *link structures*. Entity structures are anchored to markables in segmented communicative or textual data while marking them up for specific purposes, such as annotating temporal or spatial information in language. In contrast, link structures each relate an entity structure to a set of other entity structures.

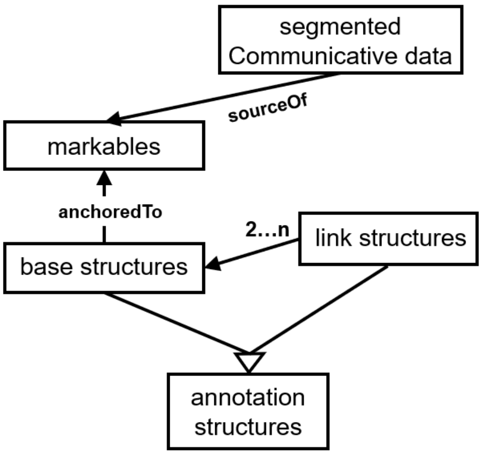

Figure 2: Two-level Annotation Structures

3.2 Temporal Link

In ISO-TimeML, the temporal link relates two entity structures annotating events temporally.

(4) a. John left$_{e1}$ before$_{s1}$ it rained$_{e2}$.

b. Temporal Annotation:

```
<EVENT id="e1", target="w2"
pred="LEAVE"/>
<SIGNAL id=""s1', target="w3"/>
<EVENT id="e2", target="w5"
pred="rain"/>
<TLINK eventID="e1",
relatedToEvent="e2",
relType="BEFORE",
signalID="s1"/>
```

TLINK can be represented in UMR at its document level, as shown earlier in Example 3.

3.3 Quantification and Scope

Pustejovsky et al. (2019) demonstrated how quantifier scoping in ISO-Space could be treated in UMR. Example 5 shows how ISO (2014) annotates quantifier scoping.

(5) a. A computer$_{se1}$ is on$_{ss1}$ every desk$_{se2}$.

b.
```
<spatialEntity
 id="se1" pred="computer"
 quant="1">
<spatialEntity id="se2"
 pred="desk" quant="every"
 scopes="se1"/>
<event id="e1"
 pred="isLocated"/>
<sRelation id="sr1"
 pred="on"/>
<qsLink figure="se1"
 ground="se2"
 relType="on", trigger="sr1"/>
<scopeLink figure="se2"
 ground="se1"
 relType="wider"/>
```

The attribute @scopes in <spatial Entity id="se2"/> is not an inherent property of entities but is contextually marked up.

In Example 6, UMR represents quantifier scoping at the document, better called *discourse*, level.

(6) Quantifier Scoping in UMR:

```
(s / sentence
  :text "A Computer is on every desk"
  (i / be-located-at-91
    :theme (c / computer
              :quantity 1)
    :location (d / desk
                  :quantity every)))
(d/ discourse level
    :scope (sc :wide sd))
```

The last line in (6), following the UMR guidelines, is to be interpreted as follows: sc indexes the argument of sentence s denoted by c, i.e., *a computer*, while sd indexes the argument of sentence *s* denoted by , i.e., *every desk*. This then can be paraphrased as "every desk (sd) has a wide scope over a computer (sc)".

In the UMR format, Gysel et al. (2021) treats scope by introducing an inverse relation *pred-of* that indicates a predicate like *answer-01* as in Example 7 is a predicate under the scope node.

(7) "Someone didn't answer all the questions."

```
(a /answer-01
   :ARG0 (p /person)
   :ARG1 (q /question
            :quant All :polarity -)
   :pred-of (s / scope
               :ARG0 p :ARG1 q))
```

The scope node indicates that *someone* takes wider scope over *(not) all the questions*.

As in Example 6, we may also represent the scopal relation in Example 7 at the document, better called *discourse* level of UMR.

(8) "Someone didn't answer all the questions."

```
(a /answer-01
   :ARG0 (p /person)
   :ARG1 (q /question
            :quant All :polarity -))
(d /discourse level
   :scope (q :wide p))
```

Unlike Representation 7, Representation 8 explicitly states *someone* p has wide scope over *(not) all the questions* q. Such a discourse-level representation can thus accommodate other types of scopal relations, *dual* and *equal*, which Bunt et al. (2018) claim to be necessary for quantification in general.

With the scopal relations thus specified, Representations 7 and 8 both yield an identical first-order logical form, yielding an identical interpretation:

(9) $\exists p[person(p) \wedge \neg\forall q[question(q) \rightarrow \exists a[answer\text{-}01(a) \wedge \text{ARG0}(a,p) \wedge \text{ARG1}(a,q)]]]$

3.4 Treating Non-consuming Tags

SpatialML (MITRE, 2010), from which ISO-Space originated, introduces so-called *non-consuming tags* for assumed places.

(10) a. Raw Data:
We drove 50 miles east of Boston. The next day, we drove 100 miles north.[1]

b. Three Non-consuming PLACE Tags:
We drove $\text{PLACE}_{pl1:target}$ 50 miles east of $\text{Boston}_{pl2:source}$. The next day, we drove $\text{PLACE}_{pl1:source}$ $\text{PLACE}_{pl3:target}$ 100 miles north.

c. RLINK in SpatialML:

```
<RLINK id=5 source=pl2:Boston
target=pl1 distance=2:50 miles
direction=E signals=2 3/>
<RLINK id=9 source=pl1 target=pl3
distance=6:100 miles
direction=N signals=6 7
```

We can identify a non-consuming tag as an implicit argument to a relation (e.g., an event) that is not syntactically realized.

Every motion triggers a trajectory that a moving object traverses. ISO-Space (ISO, 2020) has thus introduced a non-consuming tag, called *event path*, for trajectories to replace RLINK in SpatialML (MITRE, 2010). Consider Example 11 to see how it is annotated by ISO (2020).

(11) a. Categorized word-segmented Data:
$\text{John}_{x1:w1}$ $\text{drove}_{e1:w2}$ 50_{w3} miles_{w4} east_{w5} of_{w6} $\text{Boston}_{pl1:w7}$.
$\emptyset_{pl2:goal}$ $\emptyset_{ep1}$

b. entity structures:

```
<ENTITY id="x1" target="w1"
type="PERSON" name="John"/>
<EVENT id="e1" target="w2"
pred="DRIVE"/>
<PLACE id="pl1" target="w7"
type="CITY" name="Boston"/>
<PLACE id="pl2"/>
<EVENT_PATH id=ep1 mover="x1"
source="pl1" goal="pl2"
direction="E" distance="50 mi"
trigger="m1"/>
```

c. Link structure:

```
<MOVELINK figure="x1"
ground="ep1"
reltype="TRAVERSES"/>
```

Annotation 11 contains two non-consuming tags: $\emptyset_{pl2:goal}$ and $\emptyset_{ep1}$. The first tag refers to the goal, the second one to the event path created by the motion of John's driving.

Example 12 shows how these non-consuming tags are represented in UMR.

(12) Representing an event-path in UMR:

```
Data (John drove 50 miles east of Boston.)
Predciate-structure level
 (d / drive-01
   :ARG0 (p / person
```

[1]Taken from MITRE (2010), Section 15.

```
            :name (n / name
                   :op1 "John"))
   :distance (q / distance-quantity
                :quant 50
                :unit (m / mile))
   :direction (e / east )
   :source (c / city
              :name (n2 / name
                     :op1 "Boston"))
   :goal (p1 / place)
   :path (p2 / path
             :dynamic
             :trigger d
             :mover p
             :source c
             :goal p1
             :distance q
             :direction e)
    :aspect Performance
    :modstr FllAff))
Discourse-structure level
   :moveLink (t1 /traverse
                :arg1 p
                :arg2 p2
                :trigger d))
```

As shown at the discourse-structure level of Example 12 above, UMR successfully represents the traversal relation between the mover `p John` and the event-path `p2` triggered by the motion `d` of John's driving.

3.5 Complex entity structures

In ISO (2025), some entity structures are annotated as referring to other entity structures to specify their temporal values. Here is an example:

(13) a. Data:
`We left`$_{e1}$ `[`$_{t11}$ `two weeks]`$_{t12}$ `before Christmas`$_{t2}$`.`

b. Annotation scheme=ISO (2012):

```
<EVENT id="e1" pred="LEAVE"/>
<TIMEX3 id="t1"
target="two weeks"
type="DURATION" value="P2W"
beginPoint="t11" endPoint="t2"/>
<TIMEX3 id="t12"
type="DATE" value="2004-12-25"/>
<TIMEX3 id="t11"
type="DATE" value="2004-12-11"
temporalFunction="TRUE"
anchorTimeID="t1"/>
<EVENT id="e2" pred="Christmas"/>
<TLINK eventID="e1"
relatedToTime="t11"
relType="IS_INCLUDED"/>
<TLINK eventID="e2"
relatedToTime="t12"
relType="IDENTITY"/>
```

The entity structure `<TIMEX3 id=t1>` in (13b) has two attributes, `@beginPoint` and `@endPoint`, which refer to other entity structures for their values. The value of `@beginPoint` is calculated as `2024-12-11`, anchored to the Christmas day `t1`, as annotated in `<TIMEX3 id=t11>` with two attributes `@temporalFunction` and `anchorTimeID`.

AMR can also represent how the value of `@beginPoint` of a time interval, on which the motion of "our levaing" took place, is expressed:

(14)

```
Data (We left two weeks before Christmas.)
Predicate-structure level
 (l/ left-01
   :ARG0 (p / person
            :ref-person 1st
            :ref-number Plural)
   :time (d / date-entity
           :mod (t3 / temporal-interval
              :quant 2 :unit (w / week))
               :start (d1 / date-entity
                        :month 12
                        :day 11)
               :end (d2 / date-entity))
   :temporal (b/ before
              :op1 (n/ name
              :op2 (c/ Christmas
                :date (d2 / date-entity
                        :month 12
                        :day 25))))
   :aspect Performance
   :modstr FllAff)
Discourse-structure level
   :corefence (s / same-date
                :arg1 d
                :arg2 d1)
   :temporal (c / contains
                :arg1 d
                :arg2 l))
```

On the entity structure level, the start of the 2-week duration is dated December 11, for the end of the duration is the same date of Christmas, December 25, as represented on the link structure level.

The departure is also represented as occurring on December 11 at the link structure level.

4 Motion-oriented Dynamic Information

4.1 Overview

Pustejovsky and Moszkowicz (2011) combined TimeML (Pustejovsky et al., 2005) and SpatialML (Mani et al., 2010) into the Spatio-temporal Markup Language (STML) to annotate dynamic information involving motions and motion paths in language. Now, STML can be updated to ISO (2012) and ISO (2020), which have formally defined the notion of event paths triggered by motions. An event path, triggered by a motion, is traversed by a moving object and is thus defined as a nonempty finite directed sequence of spatio-temporally delimited positions of a moving object. Dynamic Interval Temporal Logic (DITL) was adopted as the semantics of STML for reasoning with programs.

We work with an excerpt from a travelogue through Central America, taken from Pustejovsky and Moszkowicz (2011):

(15) Sample Raw Data:
John left San Cristobal de Las Casas four days ago. He arrived in Ocosingo that day. The next day, John biked to Agua Azul and played in the waterfalls for 4 hours. He spent the next day at the ruins of Palenque and drove to the border with Guatemala the following day.

We first show, in Subsection 4.2, how STML annotations in XML are represented in UMR.

4.2 Representing STML Annotations in UMR

For illustration, we take the first sentence from Data 15 and segment it into words and mark up their category identifiers.

(16) Sample Data: Categorized Segmentation
$_{S1}$[John$_{se1:w1}$ left$_{m1:w2}$ [San Cristobal de Las Casas]$_{pl1:w3}$ four days$_{t1:w4-5}$ ago$_{s1:w6}$].

We now apply STML to annotate Sample Data 16 in XML.

(17)
```
<annotation id="a1" aScheme="STML">
<spatialEntity id="se1" target="w1"
  type="person" name="John"/>
<motion id="m1" target="w2"
  type="transition" pred="leave"/>
<place id="pl1" target="w3"
  cvt="town" form="name"/>
<timeX3 id="t1" target="w4-5"
  type="duration"
  value="4" unit="day"
  start="t11" end="t12"/>
<timeX3 id=""t11"target=""
  type="date" value="2025-03-08"/>
<timeX3 id="t12" target=""
  type="date" value="2025-03-12"
  trigger="s1"/>
<signal id="s1" target="w6:ago"/>
<eventPath id="ep1" target=""
  start="<pl1,t11>"
  end="<unknown,t12>" trigger="m1"/>
<tLink id="tL1" eventID="m1"
  relatedToTime="t11"
  relType="DURING"/>
<moveLink id="mvL1" figure="se1"
ground="ep1" relType="traverses"/>
</annotation>
```

Annotation 17 above represents the information about John's departure from San Cristobal, which occurred on the day marked as `t11`. This date represents part of the mover's start position `<pl1, t11>` of a 4-day duration or interval stretched to the present utterance time, today or DCT (document creation time).

Representation 18 now shows how Annotation 17 in XML can convert to UMR:

(18) Data (John left San Cristobal de Las Casas four days ago.)

```
Predicate-structure Level
(l / leave-01
  :ARG0 (s1p / person :name John)
  :time (d / date-entity
          :mod (t1 / temporal-interval
                  :duration (v / value
                              :quant 4
                              :unit day)
                 :start (d1 / date-entity
                               :year 2025
                               :month 3
                               :day 8)
                  :end (t2 / today)))
  :source (s / start-position
             :op1 (l2 / location
                     :name San Cristobal
                           de Las Casas)
```

```
            :op2 d1)
  :aspect Incremental Accomplishment
  :modstr FllAff)
Discourse-structure Level
 (:temporal (b / before
              :arg1 d :arg2 t2)
  :temporal (c / contains
              :arg1 d :arg2 l))
```

John's departure implies a durative performance of eventually reaching a goal. This action also develops incrementally. Hence, UMR marked the aspect of `leave-o1` as *Incremental Accomplishment* in UMR Representation 18, while the ISO annotation schemes fail to do so.

Temporal Interval vs Duration In Example 18, the concept `:time` refers to the occurrence time of the motion *leave*, whereas the concept `:duration` is its modifier. In Example 19, on the other hand, the duration *four hours* modifies John's activity of playing directly, meaning that it lasted four hours, while *the next day* was the time of its occurrence.

(19) Data: (The next day, John biked to Agua Azul and played in the waterfalls for 4 hours.)

```
Predicate-structure Level
(b \ bike
   :ARG0 John
   :time (d / day)
   :duration (t / temporal-quantity
                :quantity 4 unit:day))
```

4.3 Adopting GLAMR

Tu et al. (2024) propose a Generative Lexicon-based AMR (GLAMR) to capture the dynamics associated with change predicates. Adopting GL's subevent structure for verb meaning (Pustejovsky, 1995), a predicate meaning consists of a series of subevent structures related to various transitions triggered by motions or transactions, such as transfer of possessions as in GL-VerbNet (Brown et al., 2019). This structure provides relevant spatio-temporal information on sub-event structures related to various transitions. It also captures the aspectual notions of incremental accomplishment by adding the event structure directly under the topic predicate node, as in Example 20.

(20) `t / target (John left San Cristobal de Las Casas four days ago.)`

```
Predicate-structure level
  (l/ leave-01
     :ARG0 (j / john)
     :event-structure (s /subevents
       :E0 (d / do
             :action l)
       :E1 (h / has_position
              :theme j
        :initial_loc (s1 / San Cristobal)
           :initial_time d1)
       :E2 (a / and
           :op1 (m / motion
              :moving-object j
              :trajectory p)
           :op2 (h1 / has_position)
              :polarity -
              :theme j
              :location s2
              :time d2))
     :time (d / date-entity
              :mod (t1 / temporal-interval
             :duration (q1 / temp-quantity
                           :quantity 4
                        :unit (d3 / day))
              :start (d1 / date-entity)
              :end (t3 / today)))
     :event-path (p / positions
        :trigger m
        :moving-object j
        :start (p1 / position
          :location s1
          :time d1
            :op1 (q2 / spatial-quantity
                  :unit meter
                  :quantity 0))
        :next (p2 / position
          :location s2
          :time d1
           :op1 (q3 / spatial-quantity))
        :end (p3 / position))
     :modstr FllAff)

Discourse-structure level
     :temporal (b / before
                 :arg1 d1 :arg2 t3)
     :spatial   (g / greaterThan
                 :arg1 q3 :arg2 q2))
```

The event-structure and the event-path share values, but from different perspectives. The sub-event `E2` triggers the `event-path` as a trajectory of a moving object `j`. John's position changed as he moved: he was no longer in San Cristobal's initial location `s1` but moved to the next location `s2`, while all these sub-events occurred on the same day.

At the discourse or link structure level, two relations are represented: temporal and spatial. The temporal relation states that the day `d1` of John's departure from San Cristobal preceded the DCT `t3, today`, while the duration says there was a four-day interval between the departure day and the DCT. The spatial relation then states that the event-path length has lengthened from `q2` to `q3` while the mover moved from the start location `s1` to the next location `s2` or `s1+1`.

4.4 Applying Dynamic Interval Temporal Logic

DITL[2] formalizes the dynamic aspectual notion of incremental accomplishment in UMR as a program in DITL. Pustejovsky and Moszkowicz (2011) (page 16) formulates the notion of a directed motion leaving a trail as a program, represented with minor modifications in DITL, as in:

(21) Motion Leaving a Trail:
$move_{tr}(x) =_{df} pos(x) := y, b := y,$
$p := (b); (y := z, y \neq z, p := (p, z))^+$

This program states that the trail path p stretches as the beginning point b of the mover x by the Kleene iteration $+$ (more than one occurrence), as the mover x moves on. Then, the motion-triggered dynamic path p will be a sequence of x's positions, incremented iteratively as time progresses. Here, the notion of position $pos(x)$, defined as a complex function from time to $loc(x)$, which is the location of a moving object x, replaces the notion of $loc(x)$.

4.5 Dynamic Space as Minimal Embedding Ground

The spaces in which dynamic paths stretch out are also constrained by their embedding ground. Climbing over a hill creates a path tangential to the surface shape of the hill. In contrast, flying over a hill may create a path almost tangential but detached from it.

(22) Minimal Embedding Grounds
a. John climbed *over* the hill.
b. The helicopter flew *over* the hill.
c. Joh swam *around* the lake.
d. John walked *around* the lake.

[2]Mani and Pustejovsky (2012) has a fuller version of introducing DITL.

Swimming around a lake means it takes in the water, whereas running around the lake means a circular activity outside the lake. Despite the same use of spatial relators like *over* and *around*, each action or activity is characterized by a different embedding ground. Hence, the fine-grained characterization of motions or their paths should be specified with the type of embedding ground in both ISO semantic annotations and UMR.

5 Concluding Remarks

There are two commonalities between ISO SemAF standards and UMR. First, both ISO-TimeML and ISO-Space emphasize the role of events and motions. Such a focus fits well into the structure of AMR and UMR, both of which stress the predicative core of propositional content.

Second, the dual annotation structure of ISO semantic annotation frameworks such as ISO-TimeML and ISO-Space conforms perfectly to the dual level of UMR, sentence (predicate structure)-level and document (discourse)-level.

There are, however, some differences. First, ISO SemAF uses a semantic role link, tagged `SRLINK`, to assign participant semantic roles to events. By following neo-Davidsonian semantics, AMR/UMR treats them as relations between event instances and their arguments or adjuncts. ISO's semantic link needs to be applied repeatedly to assign a series of participant roles. AMR/UMR, in contrast, directly copies a series of those roles associated with each predicate from available linguistic resources such as PropBank.

Secondly, the degree of granularity in AMR /UMR differs from ISO SemAF in treating dialogue acts, discourses, and quantification. Such differences can, however, be fixed with minor but time-consuming modifications. AMR/UMR requires additional structural modifications to represent dialogue and discourse structure in a richer and more expressive fashion, one accommodating the needs of dialogue and discourse understanding in NLP. Developing such further extensions to UMR based on the work carried out within the ISO working group is an exciting challenge, and promises to better integrate standards specifications within the family of AMR representations.

Limitations

The scope of this paper is restricted. It mainly compares the representation of two ISO SemAF standards, ISO-TimeML and ISO-Space, with UMR. Our future work should be extended to other ISO standards on dialogues, discourses, quantification, and quantitative information in general. It should include studying details in annotating the tense, aspect, and modality of predicates, the specification of which varies much from language to language.

We have intentionally avoided evaluating UMR. We have accepted the review by Bos (2016) for its semantic adequacy and some articles, such as Van Gysel et al. (2021), for learnability, scalability, or applicability to computing applications. This paper did not compare computational application or scalability between ISO SemAF and AMR/UMR. This is mainly because ISO SemAF has focused on the abstract and theoretical formulation of semantic annotation structures rather than on issues of direct use in industrial applications.

We have not yet experimented with the possibility of amalgamating UMR with DRT or its subsequent extensions for semantic representation. One interesting proposal is to treat events like *walk* not as a functional type $e \rightarrow t$ but a basic type e in DRSs. We then have $[instance(e, walk), instance(j, John), actor(e, j)]$ in DRS as well as in UMR, instead of $[walk(e), John(x), actor(e, x)]$ in DRT. With this proposal accepted, we think the UMR logical format and the DRT representation format are identical.

The focus of this paper on attempting to convert XML-represented annotations to AMR/UMR is motivated by the fact that most of the ISO SemAF standards use XML as their representation format (although the DialogueBank (Bunt et al., 2016), a multilingual resource of dialogues annotated according to ISO 24617-2:2012 also uses two alternative representation formats and supports the conversion among them.) This has made all ISO SemAF standards interoperable with other ISO annotation standards on the other linguistic levels, such as lexicology, morphology, syntax, and data construction, all based on XML and the TEI Guidelines for using XML for text processing.

We understand UMR is at a developing stage and may remain as such. Our ISO working group on semantic annotation believes that some of our standards cover semantic issues such as dialogues, discourse theories, and quantification in much more breadth and depth and hopes to contribute to the editing of UMR guidelines in the future. The ISO semantics group will learn much in the area of computational applications through continued interactions with the UMR group.

Ethics Statement

All authors believe this work contributes to advancing natural language understanding, enabling more accurate and robust analysis of human-produced text. We collectively hope it will help expand equitable access to information and improve Human-Computer Interaction. At the same time, we emphasize the need for ongoing monitoring of societal impacts, particularly regarding the potential amplification of harmful stereotypes or disinformation.

Acknowledgements

The authors would like to express their gratitude to everyone whose efforts made this work possible. In particular, the lead author, Kiyong Lee, extends sincere appreciation to his co-authors for their insightful discussions and countless hours spent refining the manuscript. We are likewise deeply indebted to the three anonymous reviewers and the decision committee, whose careful reading and thoughtful critiques greatly enhanced the clarity and overall quality of the paper. Any remaining shortcomings are, of course, our own.

References

Johan Bos. 2016. Squib: Expressive power of abstract meaning representations. *Computational Linguistics*, 42(3):527–535.

Susan Windisch Brown, Julia Bonn, James Gung, Annie Zaenen, James Pustejovsky, and Martha Palmer. 2019. VerbNet representations: Subevent semantics for transfer verbs. In *Proceedings of the First International Workshop on Designing Meaning Representations*, pages 154–163, Florence, Italy. Association for Computational Linguistics.

Harry Bunt. 2010. A methodology for designing semantic annotation languages exploiting syntactic-semantic iso-morphisms. In *Proceedings of ICGL 2010, the Second International Conference on Global Interoperability for Language Resources*, pages 29–45, City University of Hong Kong.

Harry Bunt, Volha Petukhova, Andrei Malchanau, Kars Wijnhoven, and Alex Fang. 2016. The DialogBank. In *Proceedings of the Tenth International Conference on Language Resources and Evaluation (LREC'16)*, pages 3151–3158, Portorož, Slovenia. European Language Resources Association (ELRA).

Harry Bunt, James Pustejovsky, and Kiyong Lee. 2018. Towards an iso standard for the annotation of quantification. In *Proceedings of the 11th International Conference on Language Resources and Evaluation (LREC 2018)*, pages 1787–17949, Miyazaki, Japan. ELRA.

Van Gysel, E. L. Jens, Meagan Vigus, Jayeol Chun, Kenneth Lai, Sarah Moeller, Jiarui Yao, Tim O'Gorman, Andres Cowell, William Croft, Chu-Ren Huang, Jan Hajič, James H. Martin, Stephan Oepen, Martha Palmer, James Pustjovsky, Rosa Vallejos, and Nianwen Xue. 2021. Designing a uniform meaning representation for natural language processing. *KI-Künstliche Intelligenz*, 35:343–360.

ISO. 2012. *ISO 24617-1:2012, Language resource management – Semantic annotation framework (SemAF) – Part 1: Time and events*. The International Organization for Standardization, Geneva. Editorial committee: James Pustejovsky (chair), Branimir Boguraev, Harry Bunt, Nancy Ide, and Kiyong Lee (project leader).

ISO. 2014. *ISO 24617-7:2014 Language resource management – Semantic annotation framework (SemAF) – Part 7: Spatial information*, 1st edition. The International Organization for Standardization, Geneva. Project leaders: James Pustejovsky and Kiyong Lee.

ISO. 2020. *ISO 24617-7:2020 Language resource management – Semantic annotation framework (SemAF) – Part 7: Spatial information*, 2nd edition. The International Organization for Standardization, Geneva. Project leaders: James Pustejovsky and Kiyong Lee.

ISO. 2025. *ISO 24617-12:2025 Language resource management – Semantic annotation framework (SemAF) – Part 12: Quantification*. The International Organization for Standardization, Geneva. Project leader: Harry Bunt.

Kiyong Lee. 2023. *Annotation-Based Semantics for Space and Time in Language*. Cambridge University Press, Cambridge, UK.

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.

Inderjeet Mani and James Pustejovsky. 2012. *Interpreting Motion: Grounded Representation for Spatial Language*. Oxford University Press, Oxford, UK.

MITRE. 2010. *SpatialML 3.1: Annotation Scheme for Marking Spatial Expressions in Natural Language*. The MITRE Corporation, cdoran@Dmitre.org.

James Pustejovsky. 1995. *Generative Lexicon*. The MIT Press, Cambridge, MA.

James Pustejovsky. 2017a. ISO-TimeML and the annotation of temporal information. In *Handbook of Linguistic Annotation*, volume 2, pages 941–988.

James Pustejovsky. 2017b. ISO-Space: Annotating static and dynamic spatial information. In *Handbook of Linguistic Annotation*, volume 2, pages 989–1024.

James Pustejovsky, Robert Ingria, Roser Saurí, Joseé Casta no, Jessica Littman, Rob Gaizauska, Andreas Setzer, Graham Katz, and Inderjeet Mani. 2005. The specification language TimeML. In *The Language of Time: A Reader*, pages 545–557, Oxford. Oxford University Press.

James Pustejovsky and Jessica L. Moszkowicz. 2011. SpatialML: annotation scheme, resources, and evaluation. *Language Resources and Evaluation*, 44:263–280.

James Pustejovsky, Nianwen Xue, and Kenneth Lai. 2019. Modeling quantification and scope in abstract meaning representations. In *Proceedings of the First International Workshop on Designing Meaning Representation*, pages 28–33. Association for Computational Linguistics. Florence, Italy, August 1, 2019.

Jingxuan Tu, Timothy Obiso, Bingyang Ye, Kyeongmin Rim, Keer Xu, Liulu Yue, Susan Windisch Brown, Martha Palmer, and James Pustejovsky. 2024. GLAMR: Augmenting AMR with GL-VerbNet event structure. In *Proceedings of the 2024 Joint International Conference on Computational Linguistics, Language Resources and Evaluation (LREC-COLING 2024)*, pages 7746–7759, Torino, Itay. ELRA and ICCL.

Working Group UMR. 2022. *Uniform Meaning Representation (UMR) 0.9 Specification*. UMR Working Group for Guidelines.

Jens E. L. Van Gysel, Meagan Vigus, Lukas Denk, Andrew Cowell, Rosa Vallejos, Tim O'Gorman, and William Croft. 2021. Theoretical and practical issues in the semantic annotation of four indigenous languages. In *Proceedings of the Joint 15th Linguistic Annotation Workshop (LAW) and 3rd Designing Meaning Representations (DMR) Workshop*, pages 12–22, Punta Cana, Dominican Republic. Association for Computational Linguistics.