

Annotation of Tense and Aspect Semantics for Sentential AMR

Lucia Donatelli

Georgetown University
led66@georgetown.edu

Michael Regan

University of New Mexico
University of Colorado Boulder
michael.regan@colorado.edu

William Croft

University of New Mexico
wcroft@unm.edu

Nathan Schneider

Georgetown University
nathan.schneider@georgetown.edu

Abstract

Although English grammar encodes a number of semantic contrasts with tense and aspect marking, these semantics are currently ignored by Abstract Meaning Representation (AMR) annotations. This paper extends sentence-level AMR to include a coarse-grained treatment of tense and aspect semantics. The proposed framework augments the representation of finite predications to include a four-way temporal distinction (event time before, up to, at, or after speech time) and several aspectual distinctions (including static vs. dynamic, habitual vs. episodic, and telic vs. atelic). This will enable AMR to be used for NLP tasks and applications that require sophisticated reasoning about time and event structure.

1 Introduction

The Abstract Meaning Representation (AMR) is a readable and compact framework for broad-coverage semantic annotation of English sentences (Banarescu et al., 2013).¹ AMR aims to abstract away from syntactic idiosyncrasies such that sentences with the same basic meaning are represented by the same AMR graph. This paper extends existing AMR to include a coarse-grained representation of tense and aspect. Example (1) shows a sentence with its annotation from the current AMR corpus alongside our proposed additions for tense (in blue) and aspect (in purple):

- (1) “Your brother’s in the hospital and he’s not going to last the night.”
- | CURRENT | NEW |
|--------------------------------|-----------------------------------|
| (a / and | |
| :op1 (b / be-located-at-91 | :stable - :time (n2 / now) |
| :ARG1 (p / person | |
| :ARG0-of (h / have-rel-role-91 | |
| :ARG1 (y / you | |
| :ARG2 (b2 / brother))) | |
| :ARG2 (h2 / hospital)) | |
| :op2 (l / last-01 :polarity - | :ongoing - |
| | :time (a / after :op1 (n3 / now)) |
| :ARG1 p | |
| :ARG2 (d / date-entity | |
| :dayperiod (n / night))) | |

In AMR, each *concept* (entity or predicate) is tied to a variable which uniquely identifies a graph node. In PENMAN notation (Matthiessen and Bateman, 1991), a slash links a variable to its concept, and names of relations/roles (edges in the graph) are preceded by a colon. PropBank framesets of semantic roles (Palmer et al., 2005) that account for argument structure play a central role in AMR design; sentential AMR additionally includes entity typing and wikification, as well as entity and event coreference within sentences.²

This work is licensed under a Creative Commons Attribution 4.0 International License. License details: <http://creativecommons.org/licenses/by/4.0/>

¹<http://amr.isi.edu/>; data released at <https://amr.isi.edu/download/amr-bank-struct-v1.6.txt> (*Little Prince*) and <https://catalog.ldc.upenn.edu/LDC2017T10>

²*Multi-sentence* AMR is an additional layer developed to indicate cross-sentence coreference links within documents (O’Gorman et al., 2018, <https://github.com/timjogorman/Multisentences-AMR-guidelines/blob/master/ms-amr.md>).

The current annotation in (1) specifies entities and propositional structure³ but notably omits the present time meaning of the copula and the future meaning of “*going to*.” It also does not specify whether these eventualities are stative (temporary or permanent) or dynamic, and if the latter whether they are in progress, progressed to completion, or terminated prematurely. The current AMR in (1) would thus be identical for the sentence “*Your brother was in the hospital and did not last the night.*” The distinction between these two interpretations could have vital importance for practical scenarios in which automated decision-making systems operate, as well as for information extraction applications that involve identifying and situating linguistically described events both in time and in relation to one another.

Excluding semantics expressed by morphosyntactic elements such as inflection and articles has been a deliberate design decision in AMR thus far: to ensure the human annotation process would scale, it was decided that AMR should focus on the semantics of content words and relations between them.⁴ Nevertheless, the AMR design leaves room to fold in additional aspects of meaning, and as AMR is refined, we believe that the semantics of tense and aspect can be incorporated at scale by asking annotators to specify key pieces of information about eventualities as explained here.

This paper identifies and formalizes the linguistic phenomena that convey temporal and aspectual information at the sentence level, taking into account both lexical aspect and morphosyntactic cues in finite verbs/predications. Tense and aspect have been widely studied in theoretical and computational linguistics (§2). Here, we synthesize this work and develop an approach that we believe embodies AMR’s guiding principles of simplicity and transparency (§3). Finally, we discuss constructions that merit special attention as evidenced through initial pilot annotation work (§4), and we conclude with ongoing troubleshooting and future work to update the annotation scheme (§5).

2 The Case for Tense and Aspect Semantics in AMR

Developing a framework for semantic representation of tense and aspect necessitates striking a balance between the elements that grammatically encode relevant information and the more flexible pragmatic effects that such information may have. There has been a vast amount of research on temporality and aspectuality, both theoretical (Reichenbach, 1947; Vendler, 1957; Comrie, 1976; Comrie, 1985; Langacker, 1982; Dowty, 1986; Hinrichs, 1986; Moens and Steedman, 1988; Klein, 1994; Chang, 1997; Chang et al., 1998; Partee, 1999; Allen et al., 2008; Croft, 2012) and applied to corpora (see §2.2). However, we wish to capture basic distinctions of practical value for NLP applications, and so we must distill this research into a small number of annotation categories.

Here we present elements of existing AMR that relate to the new annotations; we then consider how previous and existing frameworks apart from AMR have attempted to define and specify the temporal nature of eventualities in text. These serve as points of comparison for how to integrate both theoretical and cross-linguistic findings into a versatile, albeit English-centric, annotation design for AMR.

2.1 Tense and Aspect Relevant Semantics in Current AMR

Current AMR does not explicitly formalize the semantics of tense and aspect. However, temporal or aspectual properties of eventualities are represented in the AMR if lexically explicit in the sentence.

AMR uses PropBank framesets of semantic roles as a central part of its overall annotation scheme. Though these framesets may represent eventualities in the syntax, the correspondence is not one-to-one, making it necessary to reference the original sentence for precise event identification.⁵ When these frames are specified with lexically overt temporal expressions (typically in the form of prepositional phrases and adverbials, marked using the AMR role :time), event identification is facilitated.⁶ Such expressions may imply a specific event time in relation to the discourse, as with adverbials including *yesterday* or

³This includes both the PropBank frameset *last-01* and the AMR-specific frameset *be-located-at-91*.

⁴Grammatical elements such as mood, modality, and negation were included in AMR 1.0, as they were identified as more “contentful” than other pieces of morphosyntax at the time.

⁵For example, *teacher* is annotated as (p / person :ARG0-of (t2 / teach-01)), where the frame *teach-01* is not an identifiable eventuality or suitable for current tense and aspect annotation without morphological decomposition.

⁶:time is used for all temporal location modifiers, however, as in expressions like *former Soviet Union (:time former)* and is not limited to eventualities.

tomorrow; other times, these temporal relations are ambiguous, such as *in June* (in the past or future). The aspectual interpretation of these eventualities remains unspecified, but it may fall out from the participant roles expressed: “*I ran a marathon*” conveys a different aspect than “*I ran for two hours.*”

Apart from `:time`, AMR annotates additional “non-core” roles; these roles may convey temporal or aspectual information, but they crucially depend on the presence of explicit lexical material to be included in the annotation. For example, the role `:duration` specifies the temporal length of an event; `:purpose` may specify the future-oriented endpoint of an event; `:cause` may specify the (typically) past-oriented reason for an event; `:condition` or the frame `have-condition-91` may specify the non-temporal conditional antecedent of an event; `:subevent` indicates a part-whole event relation; and `:frequency` indicates how often something may occur. In the absence of explicit lexical material, such non-core roles are accordingly absent in AMR frames though they may exist implicitly in the discourse or via reasoning.

2.2 Related Tense and Aspect Semantic Annotation Frameworks

Temporality and aspectuality have played an important role in several annotation schemes that have been applied to English corpora for NLP.⁷

TimeML. TimeML (Pustejovsky et al., 2003; Pustejovsky, 2017; Pustejovsky et al., 2017), a markup language for temporal and event expressions, is designed to address four key questions in event and temporal expression markup: (i) how to timestamp an event (identify it and anchor it in time); (ii) how to order events with respect to one another (considering both lexical and discourse properties of ordering); (iii) how to reason with contextually underspecified events (e.g. “*last week*” and “*two weeks before*”); and (iv) how to reason about the persistence of events (i.e. how long the event or its outcome lasts). TimeML precisely identifies events and their temporal relations within a text. This ability to reason with contextually underspecified events at the sentence level is a main concern of the current framework.

Situation entity labeling. Friedrich and Palmer (2014) and Friedrich et al. (2016) present a method for automatically identifying the type of a situation entity (SE) using a system of seven SE types (Smith, 2008). The authors assume SEs to be expressed at the clause level: *state*, *event*, *report*, *generic sentence*, *generalizing sentence*, *question*, and *imperative*. The labeling of SE types is a non-trivial task even for humans, and the annotation trial encountered many borderline cases—an issue the current framework also deals with. The clause-level focus of annotation is the most similar to the current proposed framework, and the current framework makes reference to a few of the SE types.

Richer Event Description (RED). The RED corpus (O’Gorman et al., 2016) captures coreference, bridging, temporal, causal, and subevent relations in the same annotation; the annotations provide an integrated sense of how the events in a particular document relate to each other with the intent of developing systems that learn rich interactions between events. The annotation scheme itself is based on a fine-grained analysis of aspectual structure combined with an analysis of physical event types following proposals in theoretical linguistics (Vendler, 1957; Croft, 2012). For example, *states* are further specified as (i) transitory, (ii) inherent permanent, (iii) acquired permanent, or (iv) point. By decomposing complex events in a clause, the overall dynamic causal network of entities interacting over time as described in a text may be modeled. Though the RED framework is more fine-grained and discourse-oriented than the current proposed framework, the need to specify event-event relationships for a complete event structure of the sentence is an important takeaway. Challenging areas (event relations conveyed via presupposition, modality, idioms, and/or inference) overlap with challenges for the current guidelines, as well.

Causal and Temporal Relation Scheme (CaTeRS). CaTeRS (Mostafazadeh et al., 2016) is a semantic annotation framework of event-event relations in commonsense stories designed to capture a comprehensive set of causal and temporal relations. The authors define an event to be any lexical entry under the following ontology types in the TRIPS ontology (Allen et al., 2008): Event-of-state; Event-of-change; Event-type; Physical-condition; Occurring; and Natural-phenomenon. Temporal relations are specified between events following the Interval Algebra of (Allen, 1984). Causal relations are further specified

⁷Different languages face different hurdles: e.g., tense is not overt in Chinese grammar, which has prompted work on disambiguation using contextual cues (Zhang and Xue, 2014). The approaches listed here assume different stances with regards to their cross-linguistic validity, for which the reader is directed to the original papers.

	States		Dynamic Events
	:stable +	:stable -	
Episodic	:time ■ now if salient <i>He lives/lived/used to live in Paris.</i>	:time ■ now <i>He was/is living in Paris.</i>	:time ■ now , :ongoing +/-/? , and :complete +/- if telic and realized <i>He went/is going/will go to Paris.</i> <i>He has been to Paris (ever, recently).</i> <i>He has been touring Paris for the past week.</i>
Habitual		:habitual + , and :time ■ now if salient <i>He is in Paris often.</i>	:habitual + , and :time ■ now if salient <i>He goes to Paris often.</i>

Table 1: Overview of tense/aspect annotation scheme by stativity and habituality. Aspectual features are in bold. **:time ■ now** is short for one of: **:time now**, **:time before now**, **:time after now**, **:time up-to now**. For habituals and stable states, **:time** is only annotated if there is a clear relation to the present time, e.g. past time expressed by *used to*.

as *cause (before/overlaps)*; *enable (before/overlaps)*; *prevent (before/overlaps)*; and *cause-to-end (before/overlaps/during)*. For example, “*The [famine]_{e1} ended the [war]_{e2}*” relates *e1* and *e2* as *cause-to-end (overlap)*; these relations are annotated at both the sentence and discourse levels. CaTeRs is noteworthy for the current paper because, while the ability to describe events as individual units may be relatively straightforward, the intersective temporal and causal relations between events may be less obvious.

Tense Sense Disambiguation. The Tense Sense Disambiguation (TSD) (Reichart and Rappoport, 2010) task addresses the multiple possible semantic senses of English grammatical tense constructions, which are listed in a lexicon. For example, the ‘present simple’ construction has several possible meanings, including that of a future event at a definite time (viz. *arrives* in “*My brother arrives in the evening*”) and that of a repeated event (*rises* in “*The sun rises in the East*”). We share the philosophy of TSD that grammatical tense constructions are ambiguous, though our approach is less atomic in nature to make fewer distinctions and create semantic tense/aspect categories more amenable to annotators.

3 Annotation Scheme

3.1 Principles of the Annotation Scheme

As AMR is a semantic annotation scheme, we seek an inventory of categories capturing the semantic dimensions of tense and aspect. Labels of *morphosyntactic* tense and aspect categories like PAST PERFECT and SIMPLE PRESENT would be too superficial as these categories are polysemous. For example, “I leave for Paris tomorrow” presents a *future* event in the guise of the SIMPLE PRESENT. A semantic scheme should abstract away from the morphosyntax and mark it as having future time.

Based on the literature on tense and aspect, we propose a set of key semantic contrasts for annotating English sentences. We believe our approach is in the spirit of AMR, striking an effective balance between the complexity of natural language on the one hand, and simplicity and readability in the meaning representation necessary for its scalability to large corpora on the other. As such, we do not claim to cover all possible distinctions in the realm of tense and aspect, though we keep finer-grained nuances of temporal and aspectual meaning in mind for potential future work. Our annotation principles are as follows: (i) We enhance AMR as **it currently exists** to better capture tense and aspect. Where time and event structure are already encoded (e.g., via temporal adverbials), we do not propose major changes. (ii) AMR abstracts away from morphosyntactic particularities of English. Our extensions therefore capture **semantic** distinctions, going deeper than grammatical tense/aspect categories (e.g. morphological past/present, progressive, perfect, etc.), which are polysemous. (iii) Our extensions are mostly **categorical** in nature (e.g. past/present/future time, static vs. dynamic, habitual vs. episodic) rather than a full calculus of timeline points and intervals. (iv) We focus on the semantic contrasts that are important to **English** grammar. Nevertheless, our design decisions are informed by reference to tense and aspect expression in other languages and the possibility of AMR being extended beyond English (Xue et al., 2014; Li et al., 2016; Migueles-Abraira et al., 2018). (v) Our current approach is limited to the semantics of **finite** verbs/predications. (vi) Our current approach is annotated at/below the level of the **individual sentence**. Additional temporal relations in discourse would need to be annotated at the document level. (vii) As AMR prioritizes broad-coverage corpus annotation, we have developed our scheme with reference to

corpus data. This has required us to consider tense/aspect as it intersects with other semantics and constructions, including narrative tense morphology, generics/habituals, conditionals, and modality.

In line with the neo-Davidsonian approach that AMR takes, we treat tense and aspect as non-core roles or variables that modify eventualities (events/states). The annotation for each category follows current AMR guidelines closely so that it may be integrated seamlessly into existing annotations. The basic distinctions are summarized in table 1; we describe them below.

3.2 Annotation Procedure

Annotators are given a set of guidelines that explain how to identify annotation targets (§3.3); how to use table 1 and table 2 and apply labels for tense (§3.4) and aspect (§3.5); and how to address challenging constructions such as modality, certain auxiliary verbs, AMR special frames, and negation (§4). Thus far, annotators have used the online AMR editor, adding new relations presented here to existing AMRs.⁸ Several initial pilot annotations with both expert and novice annotators (further detailed in §4) have helped to refine both our annotation procedure and our overall labeling scheme.

3.3 Annotation Targets

To keep the annotation task manageable, we only consider as sites for tense/aspect annotation those concepts in the AMR that correspond to finite predications in the syntax.⁹ This restriction is motivated by the understanding that finite clauses explicitly describe eventualities on the temporal dimension, and thus tense-aspect category values are relevant (Langacker, 1987).

We understand eventualities to include all kinds of events: states, activities, achievements, accomplishments, and processes (Bach, 1986). Though labels themselves are annotated on the AMR concept that corresponds to the finite verb, they refer to the event structure denoted by the entire predication as evident in the frameset (which corresponds to the verb’s argument structure). Thus, while the sentence “*I read a book last night while sitting*” possesses only one annotation target (read-01), “*I read a book last night while I was sitting*” possesses two targets (read-01, sit-01). Note that a verb’s arguments may be relevant to the aspectual annotation on the annotation target: e.g., “*I read a book last night*” is telic whereas “*I read poems last night*” is atelic due to the nature of the complement. All annotation targets receive explicit aspectual labels, and some will additionally receive explicit time labels.

3.4 Time

We annotate the location of an eventuality in time relative to speech time.¹⁰ This may be signaled grammatically with auxiliaries and/or verbal morphology, and/or with temporal adverbials. We annotate this information on an AMR annotation target using the :time role (note that another :time annotation may already be present to reflect temporal information expressed adverbially).

Present Time. The current annotation overtly marks expressions of well-defined semantically present tense when the event time cooccurs with the speech time. Target expressions are verbs in the present form. They are annotated by adding :time (n / now) under the AMR annotation target.¹¹

(2) “What are you doing here?”¹²

(b / be-located-at-91
:time (n / now))

Future Time. We mark expressions of semantically future tense when the event time occurs after the speech time (disregarding any notions of modality in such constructions). Typical targets are (i) the use of the auxiliary “will” in front of the bare verb; and (ii) the use of the present tense with a future-oriented

⁸<https://amr.isi.edu/editor.html>

⁹As mentioned above, many of the PropBank predicates in AMR correspond to eventualities expressed with nonfinite clauses, deverbal nouns, and so forth. In future work, we will explore extending our annotations to such concepts.

¹⁰Hence the label “time” instead of “tense.”

¹¹Aspectual annotations are not shown in this section.

¹²Many sentences provided here come from *The Little Prince*, which has been annotated for English AMR. For purposes of clarity and space, we omit parts of the AMR that are not relevant to the current annotation framework.

temporal adverbial, as “*I leave for Paris tomorrow.*” The grammaticalized expression *going to* can also have a future meaning, as in (1). Targets are annotated with `:time (a / after :op1 (n / now))`:

- (3) “I will try to make my portraits.”
(t / try-01
:time (a / after
:op1 (n / now)))

Past Time. We mark expressions of semantically past tense when the event time occurs prior to speech time. Typical target expressions are past tense verbs, though not all uses of morphological past tense reflect past time.¹³ Targets are annotated with `:time (b / before :op1 (n / now))`.

- (4) “I pondered over the adventures of the jungle.”
(p / ponder-01
:time (b / before
:op1 (n / now)))

Continuative Time. The English perfect can indicate that an eventuality has been in effect up to the present moment, without specifying whether it will continue.¹⁴ We represent this with *up-to* in the AMR (as contrasted with AMR’s *until*, indicating an endpoint of the eventuality’s duration):

- (5) “The wildfire has spread.” (to its current size, and may not be done spreading)
(s / spread-02
:time (u / up-to
:op1 (n / now)))

Discourse Time. Often, the time reference of an event is sensitive to the structure of the discourse, which is not fully determined by the grammatical tense of the predicate. In such cases, we omit `:time` annotation and leave the sentence temporally underspecified. This occurs in two primary circumstances:

1. *Sequence of tense / embedded clause environments.* For example, in “*He said the elephant was noisy*”, we annotate the matrix clause frame *say-01* with past time, leaving the embedded clause bare.

2. *Ongoing narrative / discourse.* In the course of discussing a situation that takes place in the past or future, the use of tense may not always signal a contrast with the present, but may simply reflect the frame of reference for the situation. For the sentence sequence, “*There was a book on the table. It was in Russian.*” (Klein, 1994), we do not consider the book’s being Russian to be limited to the past. Thus, we omit `:time` annotation on the second sentence, letting it fall out from the discourse.¹⁵

3.5 Aspect

Grammatical aspect may be understood as the *how* of an eventuality, in comparison to the *when* denoted by grammatical tense. Our basic approach to aspect is summarized in table 2. It relies on the following features:¹⁶

:stable + / - We use this role only to refer to states, capturing the canonical distinction between events and states in our annotation. Examples of `:stable +` states are those that are permanent characteristics of individuals or entities that last a lifetime; acquired states that signal a (mostly) irreversible change; and identity relationships. `:stable +` is also used to mark stative generalizations over kinds (Friedrich et al., 2016). `:stable -` states are likely to change, including transitory states that are bounded in time; point states that are bounded and exist at a single point in time; and descriptions that vary by context or time.¹⁷

¹³Clear examples of these non-past time uses are conditionals (“*If our team lost tomorrow, it would be a disaster*”) and polite requests (“*I wanted to ask you...*”)

¹⁴The English perfect has additional uses, which we address in §4.

¹⁵It is often difficult to decide when discourse time is established, how long it is maintained for, and when it needs to be reestablished. We do not have specific guidelines on this point as of now, and look forward to further work on multi-sentence AMR to fully elaborate this standard.

¹⁶Throughout we cite terminology established by Vendler (1957) and utilized by others in linguistic research (Abusch, 1985; Levin, 1993; Rothstein, 2008; Croft, 2012). Readers are directed to these source for clarification on any linguistic terminology.

¹⁷Exceptional copula constructions that are not marked as stative are discussed in §4.

:stable +/-	APPLIES TO ANY AND ALL STATIVE EVENTUALITIES. :stable + → inherent/permanent states <i>Stella is French.</i> :stable - → temporary/transitory/non-inherent states <i>Stella is (currently) in France.</i>
:ongoing +/-/?	APPLIES TO EPISODIC EVENTS (BOTH TELIC AND ATELIC). :ongoing + → unbounded event / event in progress viewed from interior <i>I am looking for a sheep.</i> :ongoing - → bounded event / event viewed as a whole from exterior <i>I walked all over the city yesterday.</i> :ongoing ? → event that has been in progress and may yet continue <i>I have been walking all over the city today.</i>
:complete +/-	APPLIES ONLY TO REALIZED, TELIC EVENTS IN COMBINATION WITH :ongoing +/- . :complete - → realized, telic event that is ongoing/in process, or has ended and not reached completion <i>Lucas mowed the lawn for an hour (but did not finish).</i> :complete + → realized, telic event that has ended and reached completion <i>Lucas read the entire book.</i>
:habitual +	APPLIES TO REGULARLY RECURRENT EVENTUALITIES; CANNOT BE USED WITH :stable +, :ongoing, OR :complete. <i>Boa constrictors swallow their prey whole.</i>

Table 2: Basic annotation for aspectual types. As shown in table 1, the :stable +/- annotation applies to states; :ongoing +/-/? and :complete +/- apply to dynamic episodic events; and :habitual + applies to non-episodic eventualities.

As states often possess lasting duration, their temporal nature may be unclear. Below, we annotate :time (above) only if there is a clear semantic time for the state; otherwise, we leave time underspecified.

(6) “It was a picture of a boa constrictor.”¹⁸

(p / picture :stable +
:domain (i / it))

(7) “He was in Turkish costume.”

:ARG2 (c / costume-01 :stable -
:ARG1 (h / he)
:time (b / before
:op1 (n / now)))

So-called “inactive actions,” also known as “progressive states”, may be analyzed as either :stable +/- . This depends on the nature of the individuals or entities involved in the sentence and the nature of the verb. An example of such a contrast is as follows: “*The Sandia mountains lie to the east of Albuquerque*” :stable + versus “*The box is lying/lies on the bed*” :stable - (Croft, 2012).

:ongoing +/-/? This role indicates the interior or exterior perspective on the event signaled by the grammatical tense. :ongoing +, corresponding to the canonical use of the progressive BE + V-ing with dynamic episodic events,¹⁹ signals that the perspective is from the interior of the event, when it is in progress (Comrie, 1976; Portner, 1998).²⁰ :ongoing + can combine with past, present, or future time to mark the progression of both atelic activities with intransitive predicates (“*He was/is/will be eating*”) and telic activities with transitive predicates (“*He was/is/will be eating a sandwich*”).²¹

(8) “He was looking for a sheep.”

¹⁸ :time is absent from this annotation, as it reflects discourse time; see above.

¹⁹ Other uses of the progressive that are *not* :ongoing + include “*The box is lying on the bed*” (:stable -; see above) and the present progressive signaling future time, e.g. “*I am leaving for Paris tomorrow*” (aspectually annotated as :ongoing -).

²⁰ This is similar to the imperfective aspect, though imperfective aspect may be understood to include progressives, iteratives, and habituais (Dowty, 1986; Comrie, 1985). We have a special category for the latter; see :habitual +. Iteratives are simplified here as atelic events and may be annotated as :ongoing +. (“*The light flickers occasionally*”) and :ongoing - (“*The light flickered all day*”).

²¹ Undirected activities used to describe present disposition (“*He’s being a jerk again*”) as well as construals of transitory states (“*Mary’s sitting on the sofa*”) are understood here to be stative (Croft, 2012), and are marked accordingly as :stable -.

```
(l2 / look-01 :ongoing +
      :time (b / before
            :op1 (n / now)))
```

:ongoing ? is used with an episodic event that may continue or not, as in (5); this is signaled in certain constructions using the English perfect (§4). All other episodic events are :ongoing -.

:complete + / - We use this role only with episodic, realized, telic events, i.e. directed events that include a measurable change in the status of one of the verb’s arguments and that have or are taking place.²² As such, the annotation :ongoing +/-/? is a prerequisite for adding :complete +/--. complete + events are commonly associated with English simple past tense constructions for events that have a result state different from the initial state (i.e. prototypical accomplishments and achievements; semelfactives and activities lack marking of :complete +/-). Again, as AMR considers tense and aspect at the sentence level, typical activity verbs (such as *jump*, (9)) may be construed as telic depending on their use in the sentence as a whole.

(9) “I jumped to my feet, completely thunderstruck.”

```
(j / jump-03 :ongoing - :complete +
      :destination (f / foot)
      :time (b / before
            :op1 (n / now)))
```

Likewise, transitive constructions that typically receive telic interpretations due to the nature of their arguments (“*I read a book last night*”) may be construed as atelic (“*I read poems last night*”). Though rare in corpora, an example of :ongoing - :complete - is a sentence of finite temporal duration with a telic verb that has not reached completion: “*I ate lunch for an hour.*” Alternatively, a telic event that is in progress is :ongoing + :complete -: “*I am eating a sandwich right now.*”

:habitual + Most basically, :habitual + indicates aggregation of occurrences of events or :stable - states, implying they are recurrent.²³ :habitual + is used for clauses that contain a lexically dynamic verb and that denote a regular recurrence of an event, whether generic (“*Bears usually eat blueberries*”) or specific (“*I used to make pie daily in summer*”) (Mathew and Katz, 2009). This contrasts with episodic events, which refer to specific finite, irregular events (“*Mary ate oatmeal for breakfast yesterday*”). :habitual + also applies to recurring :stable - states (“*Mary is always in her seat when the bell rings*”). This contrasts with :stable +, which applies to inherent characteristics of kinds or individuals (Friedrich et al., 2016).

4 Challenging Phenomena

Many phenomena related to tense and aspect are noteworthy for their complexity and debated in the literature cited in §2; these phenomena are additionally challenging for annotation frameworks that attempt to distill such phenomena into clean categories and relations (§2.2).

The phenomena discussed here are presented as a result of the challenges our own annotators encountered during initial pilot tasks. In the process of developing the current annotation scheme, we conducted two pilot annotation studies (each using slightly modified annotation roles and guidelines) to validate and update our approach. Two expert annotators individually annotated a total of 200 existing AMRs for *The Little Prince* (the only text written for a non-specialist audience for which AMRs are publicly available), calculated inter-annotator agreement²⁴ at the overall level and for individual AMRs, and then discussed areas of agreement and disagreement. To address disagreements, we have developed a thorough set of

²²The criteria of ‘realized’ alludes to a distinction we explain further in §4 for modality. We recognize that this currently does not allow us to specify the telicity of non-realized events, an issue we will address as we further develop the annotation scheme. Telic events may be formally understood as accomplishments or achievements (Vendler, 1957).

²³The current annotation scheme is unable to capture event quantification intermediary between :habitual + and :ongoing -, such as “*I ate an apple every hour for the whole day*”. We leave this work to future developments of the scheme. The reader is directed to more recent work on such matters for reference (Bunt and Pustejovsky, 2010).

²⁴We do not present these numbers as the scheme has since evolved on the basis of the disagreements.

annotation guidelines with which to provide annotators. Currently, a larger-scale pilot study is under way with non-expert annotators that will allow us to further evaluate these guidelines and the labels presented here. Some key challenging phenomena are enumerated below.

4.1 The perfect

The semantics of the English (and cross-linguistic) perfect is far from straightforward (Anderson, 1982; Portner, 1998; Croft, 2012). We recognize three semantic categories of the perfect:

(i) CONTINUATIVE casts an eventuality as having lasted up to the present moment, without specifying whether it will continue. In the present perfect, these receive time annotation of :time (u / up-to :op1 (n / now)). In the past and future perfects, the time is specified as before or after now, and if there is an appropriate temporal reference point in the AMR due to an adverbial, time up-to that point is also specified. Aspectual annotation depends on the nature of the eventuality: “*The wildfire has spread*” as in (5) is :ongoing ?; “*Mary has lived in London for five years*” is :stable +; and “*I have been visiting the hospital daily*” is :habitual +.

(ii) EXISTENTIAL, emphasizing that something has happened ever, at least once in the past: “*The Orioles have won three World Series championships*”. Modifying our treatment of past time, we annotate this with :time (b / before :mod (e / ever) :op1 (n / now)). Aspect is annotated as usual.

(iii) RECENT, emphasizing that something has recently happened: “*The Orioles have won!*”, “*I’ve decided to stay home tonight.*” Modifying our treatment of past time, we annotate this with :time (b / before :mod (j / just) :op1 (n / now)). Aspect is annotated as usual.

4.2 Quotations

Currently, we annotate tense/aspect in quotations as if they were independent sentences. A deeper representation would relate the time of the quoted content to the time of the containing sentence, which is better left to discourse-level tense and aspect annotation. We also annotate the speech act predicate in the AMR that corresponds to quotation marks, even in the absence of an explicit speech verb.

4.3 Copula Constructions

Copula constructions are annotated following the guidelines for :stable +/- unless exceptional as described here. Presentational copulas (“*Here is the book you loaned me.*”) are only annotated for semantically salient :time. Existential copulas (“*There is a hat on the table.*”) are marked for semantically salient :time and as :stable -. Categorical copulas (“*This is a hat.*”) are annotated as :stable +; :time is only added if semantically salient. Finally, certain uses of the copula, such as in identificational constructions (“*It is/was John who ate the cake.*”), function semantically like auxiliaries to an embedded clause and are left bare in the current annotation scheme.

4.4 Inceptive States

Stative verbs may be coerced to have achievement readings, such as “*He suddenly **knew** the answer*” or “*The students **got bored** in class yesterday.*” In such construals, the verb marks the initial boundary of a state and may be understood as *start/begin to* VERB or *become* VERB (Breu, 1994; Croft, 2012). We mark such inceptive states as :ongoing - (:complete +, if realized), although the frames that indicate the completion of inception vary depending on the construction: *know-01* itself would be marked as :ongoing - :complete + in the first example, while *get-03* would be marked as a special *become* sense of the verb in the second example.²⁵

4.5 Negation

AMR currently marks negation cues as :polarity -, without detailing scope, focus, or negated events. No explicit differentiation is made between syntactic, lexical, and morphological negation. We adapt our current annotation scheme to this practice and annotate sentence-level temporal and aspectual information as if :polarity - were not present. For example: “*I never learned to draw anything*” is annotated

²⁵Cross-linguistic work may elucidate how best to annotate these construals in future work. For example, the inchoative (Finnish, Lithuanian) explicitly marks these and the perfective aspect in Romance changes stative verbs to achievement verbs.

as :polarity - :ongoing - :complete + :time (b / before :op1 (n / now)). Though this will sometimes result in fully labeling unrealized events (contrary to §3.5), AMR’s treatment of negation does not differentiate realized events (“*The patient did not test positive*”) from unrealized events (“*The patient was not tested*”).

4.6 Conditionals and modals

Though space does not permit us to outline the entirety of issues raised by modality and conditionals, it is worth making several observations.

(i) AMR does not canonicalize modal meanings expressed verbally. Thus “*You need to leave*”, “*You are required to leave*”, and “*You must leave*” are essentially paraphrases, but the modality is semantically represented as need-01, require-01, and obligate-01 (the canonical frame for modal auxiliaries of requirement). Though we hope to address this discrepancy in a future proposal, for now we annotate as follows: Canonical modal frames in AMR (obligate-01, possible-01, permit-01, etc.) do not receive annotation; rather, their embedded verbs receive annotation for both tense and aspect following above guidelines. Auxiliaries that are represented lexically in the AMR (such as need-01 and require-01) are annotated for tense and aspect themselves; their embedded verbs receive no annotation.

(ii) AMR does not in general mark a distinction between realis (realized) and irrealis (future, hypothetical, desired, etc.) eventualities, apart from certain expressions of negation (see above). For example, “*I tried to eat*” and “*I managed to eat*” differ only in the predicate frame for the matrix verb (try-01 vs. manage-02); the PropBank frame lexicon does not mark that the ARG1 of manage-02 is realis while the ARG1 of try-01 is irrealis.²⁶ We do not attempt to fully distinguish realis and irrealis eventualities in assigning tense and aspect to modal/conditional eventualities. However, we allude to future work in this area by only marking :complete +/- on realized events. This decision is driven by the idea that a goal-oriented event cannot be assessed for culmination or completion if it has not occurred.

(iii) Modal auxiliaries and conditionals can be used generically (“*I can eat an entire pizza by myself*”; “*If I am really hungry, I get angry*”) and non-generically (“*I couldn’t sleep last night*”; “*If you don’t eat that slice, I will*”). While conditional relations in AMR are typically expressed with :condition, this can be reified to have-condition-91 for adding :habitual + when the conditional as a whole is generic.

(iv) Deontic and related modalities (Palmer, 2001) refer to future but uncertain events. Thus the eventuality in question should be annotated with future time.

5 Conclusions and Future Work

This paper has proposed a framework for enhancing the Abstract Meaning Representation to encode the semantics of tense and aspect in English. As with most complex annotation tasks, many cases are nontrivial to categorize for tense and aspect. Striking a balance between specifying tense and aspect information in detail, while keeping the annotation task realistic and scalable, presents a challenge.

The current framework is coarse-grained and leaves room for further consideration of extensions such as more precise temporal relations between events, and the annotation of tense/aspect phenomena at the discourse level. We plan to examine additional genres, expecting differences in the distribution of temporal and aspectual categories—e.g., we speculate that present time, future time, and progressive aspect will be more common in speech than writing. Finally, since AMR has been adapted to other languages (Xue et al., 2014; Li et al., 2016; Migueles-Abraira et al., 2018), it is worth more fully examining how the expression of tense and aspect in other languages reflects deeper cross-linguistic semantics that AMR could capture.

Acknowledgements

We are grateful to Tim O’Gorman, Lori Levin, James Pustejovsky, Paul Portner, Nancy Chang, and members of the NERT research group at Georgetown for helpful discussions, and to anonymous reviewers for their feedback. Regan and Croft are supported by funding from the Defense Threat Reduction Agency under grant number HDTRA1-15-1-0063. Donatelli’s research was partially sponsored by the

²⁶Though neither embedded verb would receive a label following the current scheme, we provide the example as fodder for future work.

Army Research Laboratory and was accomplished under Cooperative Agreement Number W911NF-18-2-0066. The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the Army Research Laboratory or the U.S. Government. The U.S. Government is authorized to reproduce and distribute reprints for Government purposes notwithstanding and copyright notation herein.

References

- Dorit Abusch. 1985. *On verbs and time*. Doctoral dissertation. University of Massachusetts, Amherst.
- James F. Allen, Mary Swift, and Will de Beaumont. 2008. Deep semantic analysis of text. In *Proc. of the 2008 Conference on Semantics in Text Processing*, STEP '08, pages 343–354, Stroudsburg, PA, USA.
- James F. Allen. 1984. Towards a general theory of action and time. *Artificial Intelligence*, 23:123–54.
- Lloyd B. Anderson. 1982. The ‘perfect’ as a universal and as a language-specific category. In Paul J. Hopper, editor, *Tense-aspect: Between semantics and pragmatics*, pages 227–264. John Benjamins, Amsterdam.
- Emmon Bach. 1986. The algebra of events. *Linguistics and philosophy*, 9(1):5–16.
- Laura Banarescu, Claire Bonial, Shu Cai, Madalina Georgescu, Kira Griffitt, Ulf Hermjakob, Kevin Knight, Philipp Koehn, Martha Palmer, and Nathan Schneider. 2013. Abstract Meaning Representation for sembanking. In *Proc. of the 7th Linguistic Annotation Workshop and Interoperability with Discourse*, pages 178–186, Sofia, Bulgaria, August.
- Walter Breu. 1994. Interactions between lexical, temporal and aspectual meanings. *Studies in Language. International Journal sponsored by the Foundation “Foundations of Language”*, 18(1):23–44.
- Harry Bunt and James Pustejovsky. 2010. Annotating temporal and event quantification. In *Proc. of 5th ISA Workshop*.
- Nancy Chang, Daniel Gildea, and Srinu Narayanan. 1998. A dynamic model of aspectual composition. In Morton Ann Gernsbacher and Sharon J. Derry, editors, *Proc. of CogSci*, pages 226–231, Madison, WI, USA, August.
- Nancy Chang. 1997. A motor-and image-schematic analysis of aspectual composition. Technical Report TR-97-034, International Computer Science Institute, Berkeley, CA, September.
- Bernard Comrie. 1976. *Aspect*. Cambridge University Press, Cambridge MA.
- Bernard Comrie. 1985. *Tense*. Cambridge University Press, Cambridge MA.
- William Croft. 2012. *Verbs: Aspect and Causal Structure*. Oxford University Press, Oxford, UK, March.
- David R. Dowty. 1986. The effects of aspectual class on the temporal structure of discourse: semantics or pragmatics? *Linguistics and Philosophy*, 9(1):37–61.
- Annemarie Friedrich and Alexis Palmer. 2014. Automatic prediction of aspectual class of verbs in context. In *Proc. of ACL*, Baltimore, USA.
- Annemarie Friedrich, Alexis Palmer, and Manfred Pinkal. 2016. Situation entity types: Automatic classification of clause-level aspect. In *Proc. of ACL*, pages 1757–1768, Berlin, Germany, August.
- Erhard Hinrichs. 1986. Temporal anaphora in discourses of English. *Linguistics and philosophy*, 9(1):63–82.
- Wolfgang Klein. 1994. *Time in language*. Routledge, London.
- Ronald W. Langacker. 1982. Remarks on English aspect. In Paul J. Hopper, editor, *Tense-aspect: Between semantics and pragmatics*, pages 265–304. John Benjamins, Amsterdam.
- Ronald W Langacker. 1987. *Foundations of cognitive grammar: Theoretical prerequisites*, volume 1. Stanford University Press.
- Beth Levin. 1993. *English verb classes and alternations: A preliminary investigation*. University of Chicago Press.

- Bin Li, Yuan Wen, Lijun Bu, Weiguang Qu, and Nianwen Xue. 2016. Annotating The Little Prince with Chinese AMRs. In *Proc. of LAW X – the 10th Linguistic Annotation Workshop*, pages 7–15, Berlin, Germany, August.
- Thomas A. Mathew and E. Graham Katz. 2009. Supervised categorization for habitual versus episodic sentences. In *Sixth Midwest Computational Linguistics Colloquium*, Bloomington, Indiana.
- Christian Matthiessen and John A. Bateman. 1991. Text generation and systemic-functional linguistics: Experiences from english and japanese.
- Noelia Migueles-Abraira, Rodrigo Agerri, and Arantza Diaz de Ilarraza. 2018. Annotating Abstract Meaning Representations for Spanish. In Nicoletta Calzolari, Khalid Choukri, Christopher Cieri, Thierry Declerck, Sara Goggi, Koiti Hasida, Hitoshi Isahara, Bente Maegaard, Joseph Mariani, H el ene Mazo, Asuncion Moreno, Jan Odijk, Stelios Piperidis, and Takenobu Tokunaga, editors, *Proc. of LREC*, pages 3074–3078, Miyazaki, Japan, May.
- Marc Moens and Mark Steedman. 1988. Temporal ontology and temporal reference. *Computational Linguistics*, 14(2):15–28.
- Nasrin Mostafazadeh, Alyson Grealish, Nathanael Chambers, James Allen, and Lucy Vanderwende. 2016. CaTeRS: Causal and temporal relation scheme for semantic annotation of event structures. In *Proc. of the Fourth Workshop on Events*, pages 51–61, San Diego, California, June. Association for Computational Linguistics.
- Tim O’Gorman, Kristin Wright-Bettner, and Martha Palmer. 2016. Richer Event Description: Integrating event coreference with temporal, causal and bridging annotation. In *Proc. of the 2nd Workshop on Computing News Storylines*, pages 47–56, Austin, Texas, USA, November.
- Tim O’Gorman, Michael Regan, Kira Griffitt, Ulf Hermjakob, Kevin Knight, and Martha Palmer. 2018. AMR beyond the sentence: the Multi-sentence AMR corpus. In *Proc. of COLING*, Santa Fe, New Mexico, USA, August.
- Martha Palmer, Daniel Gildea, and Paul Kingsbury. 2005. The Proposition Bank: An annotated corpus of semantic roles. *Computational Linguistics*, 31(1):71–106, March.
- Frank Robert Palmer. 2001. *Mood and modality*. Cambridge University Press.
- Barbara H Partee. 1999. Nominal and temporal semantic structure: Aspect and quantification. *Prague Linguistic Circle Papers: Travaux du cercle linguistique de Prague nouvelle s erie*, 3:91.
- Paul Portner. 1998. The progressive in modal semantics. *Language*, 74(4):760, December.
- James Pustejovsky, Jos e M. Casta no, Robert Ingria, Roser Saur , Robert J. Gaizauskas, Andrea Setzer, Graham Katz, and Dragomir R. Radev. 2003. TimeML: Robust specification of event and temporal expressions in text. In *IWCS-5, Fifth International Workshop on Computational Semantics*, Tilburg, Netherlands, January.
- James Pustejovsky, Harry Bunt, and Annie Zaenen. 2017. Designing annotation schemes: From theory to model. In *Handbook of Linguistic Annotation*, pages 21–72. Springer.
- James Pustejovsky. 2017. ISO-Space: Annotating static and dynamic spatial information. In *Handbook of Linguistic Annotation*, pages 989–1024. Springer.
- Roi Reichart and Ari Rappoport. 2010. Tense sense disambiguation: A new syntactic polysemy task. In *Proc. of the 2010 Conference on Empirical Methods in Natural Language Processing*, pages 325–334, Cambridge, MA, October.
- Hans Reichenbach. 1947. *Elements of symbolic logic*. The Free Press, New York.
- Susan Rothstein. 2008. Telicity, atomicity and the Vendler classification of verbs. *Theoretical and Crosslinguistic Approaches to Aspect*, pages 43–77.
- Carlota S. Smith. 2008. Time with and without tense. In *Time and Modality*, Studies in Natural Language and Linguistic Theory, pages 227–249. Springer, Dordrecht.
- Zeno Vendler. 1957. Verbs and times. *The Philosophical Review*, 66:143–60.
- Nianwen Xue, Ondr ej Bojar, Jan Haji , Martha Palmer, Zde ka Ure ov, and Xiuhong Zhang. 2014. Not an interlingua, but close: comparison of English AMRs to Chinese and Czech. In Nicoletta Calzolari, Khalid Choukri, Thierry Declerck, Hrafn Loftsson, Bente Maegaard, Joseph Mariani, Asuncion Moreno, Jan Odijk, and Stelios Piperidis, editors, *Proc. of LREC*, pages 1765–1772, Reykjav k, Iceland, May.

Yuchen Zhang and Nianwen Xue. 2014. Automatic inference of the tense of Chinese events using implicit linguistic information. In *Proc. of EMNLP*, pages 1902–1911, Doha, Qatar, October.