

Personal event extraction annotation guide

Joshua Eisenberg, PhD
Lead Scientist, NLU
Artie, Inc.
joshua.eisenberg@artie.com

1. Introduction

1.1 Motivation

At Artie we are bringing virtual characters to life. Conversation is central to these interactions. Our virtual characters need to learn about the lives of the people they are talking to. If the virtual character can automatically extract personal events from speech, then the virtual character can reference these events in future conversations. This will help interaction seem more real, and allow

Our goal in this annotation study, is to train computers to identify events that involve the person speaking. Specifically, given a speech utterance transcribed as English text, the computer must identify events that involve the speaker, and encode the semantics of this event into a structured format.

1.2 What is annotation

Annotation is the process of explicitly encoding information about a text that would otherwise remain implicit. For this study, annotation is the record of human judgments identifying where an event is referenced in text. We will be annotating the spans of text that signal the occurrence of an event in the world that the dialogue is taking place, or in a narrative level that is being expressed in the dialogue.

2. Events

In this study, we will annotate dialogue for the occurrence of *events*, and *nonevents*. The following is our definition of event:

Event: an action or state of being occurring in the reality¹ depicted by a span of text.

According to the annotation guide for TimeML (the Time markup language) the definition of *event* is:

We consider “events” a cover term for situations that happen or occur. Events can be punctual ... or last for a period of time.... We also consider as events those predicates describing states or circumstances in which something obtains or holds true....

¹ By reality, I’m referring to the reality the speaker is in, or a narrative level they are depicting. I use the term reality, since a speaker might be talking about a made-up reality, or even a different narrative level.

A rule of thumb for deciding whether something is an event: does the event in questions advance the plot forward? Anything from lighting a cigarette to a character saying something can be considered an event. Plot events are colloquially considered to be the most important events in a story (like the main character defeating evil), however here we dictate that all actions, regardless of how mundane they seem, to be considered an event. In addition to actions which drive a plot forward, we consider states of being as events. States detail the specifics of a situation. States are events since they portray the current (even past or future See section 2.2.2 for more details about states).

Regarding the grammar of events: events are typically portrayed through verbs; events represent action and so do verbs. However, not all events are represented via verbs. Also, not all verbs represent events. It is confusing, but we will try to make this more concrete through the examples in this guide.

2.1 Recording annotations

You will record your annotations directly into text files. Events are surrounded by brackets. If an event is made up of multiple words, surround the full multi-word expression with brackets. In section 2.2 you can see many examples of the annotation syntax.

2.2 Examples

2.2.1 Getting acquainted with the annotation syntax

Chloe: I **{took}** some of the really important things to me, **{put}** them in a blanket, and then **{held}** onto the blanket every night, **{hoping}** that it would **{turn}** small with me.

Ira Glass: So what would be **{wrapped up}** in the blanket?

Chloe: My stuffed animal, Boo, my cat stuffed animal, one of my books, and a couple other stuffed animals.

Ira Glass: And were you **{hoping}** that this would happen, or were you **{fearing}** that this would **{happen}**, that you would **{get small}**?

Chloe: I **{was}** both.

Excerpt 1: From [episode 687 of *This American Life*](#)

Some observations to note from **Excerpt 1:**

- There are no events annotated the second time Chloe speaks. This is since the event is referenced in the previous utterance by Ira. When Chloe speaks, she is listing the things she wraps up in her blanket. The list of items isn't a reference to the event happening, just details about what is being affected.
- There is an event annotated the last time Chloe speaks, since she confirms her internal state, via a reference to Ira's suggestions in the previous utterance.

2.2.2 Annotating states

Expressing a state, or a change of state, should be annotated as an event. These types of events are less obvious to identify than events that represent actions (usually represented by dynamic verbs). A state explains the details of a situation. This could be a mental state which can be represented with words like *feel*, *realize*, *love*, *forgot*, *know*. States can also represent physical situations, which represent how things look, quantities of things, or the quality of things.

2.2.2.1 Animacy

In this study, we are only annotating states that express information about animate objects. We will NOT be annotating states about inanimate objects. A definition of animacy, by [Labiba Jahan in 2018](#) is:

Animacy is the characteristic of being able to independently carry out actions (e.g., movement, communication, etc.). For example, a person or a bird is animate because they move or communicate under their own power. On the other hand, a chair or a book is inanimate because they do not perform any kind of independent action.

The distinction between animate and inanimate objects is made because animate characters are more closely tied to being involved in driving a plot forward. We are more interested in capturing events which depict deliberate actions, than details describing the characteristics of inanimate objects. **However, we will still annotate inanimate objects involved non-stative events**, for example:

Right before my eyes, a boulder **{came smashing down}** from above me, just narrowly **{missing}** my car.

Above, the boulder is in animate object, but it was a participant in the event of falling from the sky, and the event of *missing* the speaker's car. Below are two examples of states of inanimate objects which do not get annotated:

The mountain was steep.

The path looked dangerous.

The first example represents a state of being in the adjective *steep*, but it is describing a mountain, which is inanimate, so it does not get annotated. The second example follows similarly, since the inanimate *path* is having its state described.

In terms of deciding what actual words that should be annotated for states: the word that explains the state or how the state is changing gets annotated. For the purpose of this annotation study, we will annotate the phrase that gives information as to what the current state is, or how it is changing. Consider the following:

We are **{running out}** of bread.

Snoopy is {**hungry**}.

In the first example, *running out* is annotated, since it represents what is happening. *We* are starting to experience a shortage of bread, which is a state. The word *are* represents that *we* are experiencing something, but *are* does not get annotated, since it only signals that a state is going to be expressed. The same reasoning applied for the second example. The event annotated is *hungry*, which is an adjective. In this sentence it represents the state of Snoopy. The word *is* signifies that a state is going to be detailed. We are more interested in annotating the type of state occurring, not the fact that a state will be expressed, hence we are more interested in annotating words that give more detail than *is* or *are*.

She was so {**angry**}.

The above example is meant to illustrate annotating the adjective describing what specific state of being experienced, instead of the verb, *was*, that signals that *she* was experiencing a state. Let's take another look at annotating animate and inanimate states.

A: What type of pizza did he {**choose**}?

B: It was mushroom and olives.

In the above example, in the response by **B** nothing is annotated. This is because *it* refers to the inanimate pizza, and we don't annotate the states of inanimate objects. In this situation, the state is the pizza having the toppings of *mushrooms and olives*. In another pizza focused example:

I {**ordered**} pizza at 11:45 PM. It was a bad {**mistake**}

In the above example, *was* does not get annotated, but *mistake* does. The *was* is referring to the event of ordering pizza in the previous sentence. However, the event's state being *mistake* does get marked, since it is classifying the event in the previous sentence as a mistake, updating the state of being. Additionally, *mistake* gives more information than *was*. Mistake changes how the event in the previous sentence gets interpreted, by updating the state. The usage of *mistake* is describing the state of an event, not an inanimate object.

Another type of state is a [copula](#). Copula's are annotated if they give information about animate objects, but not inanimate objects.

He {is} the president.

1

{We're} the chefs.

2.2.3 Negated or modal events

Events that did not happen should still be annotated.

Snoopy did not {score} a {goal}.

In the example above, *score* and *goal* are both annotated as events, even though the goal was not scored. It is just as important to know that an event didn't happen, as it is to know if it did happen.

I would of {bought} the guitar, if I {had} enough money.

Above, *bought* is annotated as an event even though it did not happen.

2.2.4 Tense of events

Most events will be referred to in the past tense, and some even in the present. However, some events can be referred to in the future tense; about what a person anticipates happening in the future, or actions the speaker intends on taking.

I'm going to {buy} a new pair of shoes.

In the above example, only *buy* is annotated, since going just refers to the fact that the event of buying might happen in the future.

2.2.5 Nested Events

The band is {planning} to {play} a cover by the Talking Heads.

There can be multiple events in the same sentence. In the example above, the event of *planning*, is happening now, and the event of *playing* will potentially happen in the future.

Phish {played} a song called *You Enjoy Myself*, where the guitarist and bassist are doing synchronized {jumps} on a trampoline, while {playing} their instruments and {dancing}.

Above is a series of events, occurring simultaneously. They are each annotated as they add information about what is happening. Note, that we do not annotate *doing* here, since it does not add any more information than the events *playing* and *dancing*. Here the word *doing* is more structural than informative. Verbs like doing can be annotated as event though, as in this example:

He finally {did} it.

2.2.6 Annotating questions

Q: Where do you {live}?

A: California.

In this annotation study we will annotate events in questions. For the above example we only annotate the word *live* as an event in the question. We do not annotate *California* since this is just a detail, it is where the person answering lives. If the person answering the question doesn't rephrase the answer in a way that explicitly states the event, then there may not be anything to annotate in the answer. This might not seem natural, since the specifics of the event are not stated in the question. However, we are thinking about the details expressed in the answer as referring back to the specific event in the answer. Hence, we will annotate the event in the answer.

2.2.7 Time expressions

Time expressions are not necessarily annotated. We annotate time expressions when they refer to when the event occurred or the duration of the event. However, we do not annotate phrases that give the frequency of how often the events occur.

I {**danced**} at the club {**until 5 AM**}.

Danielle {**brushes**} her teeth twice a day.

Above, *danced* is annotated as an *event*. The phrase *until 5 AM* is annotated as an event since it refers to the span of time leading up to 5 AM. It is a durational event. In the second example, *twice a day* is not annotated, since it just expressed how often Danielle brushes her teeth.

2.2.8 Repetition, fragments, interrupted speech, introductions

If a speaker repeats a word or phrase, and that word is signaling an event, annotate both occurrences of the word or phrase.

I {**walked**}--- {**walked**} to the market and {**bought**} some brown sugar.

If a speaker is interrupted by themselves annotate both times the event is expressed. This is shown in the example above. Similarly, if an event is interrupted, and repeated in more than one person's utterance, each mention should be annotated. In the dialogue below, the event *smoked* is annotated in both utterances.

A: When he was young he {**smoked**}—

B: --{**Smoked**} cigars.

Next, we will discuss phrases that appear at the beginning of utterances. We will not annotate events expressed in introduction phrases. Here is a not complete list of these introduction phrases:

- Let's see...
- Look,
- I mean,

- You know,

2.2.9 Imperatives

We will annotate imperatives. Imperatives give a command, they propose action or refer to a desired event. Some examples of imperative annotation is below:

Please {**clean**} the floor.

It would {**delight**} me if you would {**dance**} with me.

In the first utterance the imperative *clean* is annotated. In the second utterance, the imperative *dance* is also annotated. Both imperatives introduce an event that the speaker wants to happen.

2.2.10 He said, she said, and like...

If a speaker announces someone is speaking, even themselves, that is an event. The text that represents the speaking action occurring gets marked as an event, but the actual words that are spoken are not marked (unless distinct events are mentioned in the quoted speech).

My mom {**called**} me and {**said**}, “You should {**call**} me more often.”

In the above example, *said* is marked as an event. The full span of what the mother said, is not annotated, even though the act of speaking is an event. The event of speaking is encoded in the annotation for *said*. The other annotations are referring to other events.

My girlfriend {**came over**} and she was {**like**} I {**had**} the weirdest {**day**}.

In the above example, *like* is marked as an event since it portrays the event of the girlfriend speaking. Not all *likes* get annotated, but this one does, since it tells the listener that the girlfriend is speaking.

2.2.11 Events as nouns

We {**biked**} to the {**concert**}.

Although *concert* is a noun, it is an event, because it is something that happened.

2.2.12 Phrasal verbs

When annotating events expressed with phrasal verbs, annotate the full phrase:

I'm {**looking forward to**} the new episode of *Avenue 5*.

In the above example, the full phrase of *looking forward to* is annotated, not just the first word.

3. Concluding remarks

I will be updating this annotation guide after we both annotate our training materials. Once we put these guidelines into practice, I'm sure we will find even more situations we should make rules for. For now this is a living document.

References

Jahan, L., Chauhan, G., & Finlayson, M. (2018, August). [A new approach to animacy detection](#). In Proceedings of the 27th International Conference on Computational Linguistics (pp. 1-12).

Saurí, Roser & Moszkowicz, Jessica & Knippen, Bob & Gaizauskas, Rob & Setzer, Andrea & Pustejovsky, James. (2006). TimeML Annotation Guidelines Version 1.2.1.
https://www.researchgate.net/publication/248737128_TimeML_Annotation_Guidelines_Version_121

Episode 687 of This American Life
<https://www.thisamericanlife.org/687/transcript>