

# Qualities of Eventiveness

Sean Monahan, Mary Brunson

Language Computer Corporation

{smonahan, mary}@languagecomputer.com

## Abstract

Events are not a discrete linguistic phenomenon. Different verbal and nominal predicates express different degrees of eventiveness. In this paper we analyze the qualities that contribute to the overall eventiveness of a predicate, that is, what makes a predicate an event. We provide an in-depth analysis of seven key qualities, along with experimental assessments demonstrating their contributions. We posit that these qualities are an important part of a functional working definition of events.

## 1 Introduction

The problem of event extraction is fundamentally challenging because many definitions of “event” exist. Some predicates clearly indicate events, e.g. “I *ran* 5 miles to the store”, while others indicate states, e.g. “He *is* tall”. However, in natural language text, many predicates fall between these two extremes, e.g. “He *runs* frequently”. In order to successfully extract events, resolve event coreference across documents, and reason about the events, we must understand exactly what an event is. In this paper, we propose a series of qualities that contribute to the overall *eventiveness* of a predicate. We define eventiveness as “the degree to which a predicate is like an event?”.

The concept of “event” is not discrete, but exists along several dimensions. We identify seven qualities of predicates that lead readers to more readily consider them to be events. In order to successfully utilize events in end applications, we believe these qualities must be fully understood.

In this paper, we consider the predicate to be the word (e.g. verb or noun) in the sentence that might indicate the existence of an event. This is also referred to as a *trigger* or *anchor* in event extraction.

Each of the predicates in the following examples (indicated by italics) exhibit different degrees of eventiveness.

1. The tremors have re-awakened bitter memories of the Asian *tsunami* that killed 168,000.
2. Indonesia lies in a zone where the plates shift, sometimes generating *tsunamis*.
3. Electricity was cut off to the city, where people fled their homes fearing a *tsunami*.

The first example is most clearly an event, referring to a specific instance of a tsunami. In the second sentence, the nominal predicate “*tsunamis*” refers to a non-specific event that occurs as a result a natural occurrence. In the third, a tsunami has not occurred but is a feared possibility.

Any end application of extracted events must decide which of these predicates to consider as relevant. An application to “map known tsunamis” might only consider the first event as relevant. An application to detect newsworthy or “emerging” events might only consider the third. An application seeking to understand relationships between events could utilize the second example to determine that plates shifting causes tsunamis. In order to facilitate a wide range of applications, all of these predicates should be extracted as “events”, which can then be separated by the qualities they possess. Furthermore, consideration of these qualities should reflect human judgment about events.

In this paper, we discuss the different qualities that contribute to the eventiveness of a predicate. In Section 2, we describe previous work on defining events. In Section 3, we describe the qualities that we consider to be most representative of events. In Section 4, we describe an experiment we conducted to rate these qualities in terms of how they contribute to eventiveness. In Section 5, we conclude with a summary of our theory of events and a description of how this will aid applications in understanding events.

## 2 Related Work

There have been significant efforts to understand the idea of “events” in a variety of different communities, including physics, philosophy, psychology, and both theoretical and computational linguistics. We draw our qualities of eventiveness from across this literature in order to form a more complete view of what an event is.

Quine (1985) considers an event to be a well-individuated physical object which is clearly defined in space and time. This contrasts with the TimeML schema (Pustejovsky et al., 2003), which regards event as a “a cover term for situations that happen or occur”. Lombard (1986) considers an event to necessarily be a change. These definitions by themselves do not sufficiently explain the full boundary between event and non-event, but are useful in informing our qualities.

In addition to TimeML, a pragmatic definition of events was also adopted for ACE (2005). ACE utilized a wide definition of event, though only a small set of event types were annotated, along with their specificity, actuality, and arguments. More recently, TAC KBP (2014) has built on the ACE definition in order to extract event information to incorporate into a knowledge base.

Understanding how events are perceived by individuals has also been researched by psychologists in order to learn how people construct mental models of events. Radvansky and Zacks (2011) investigate the mental representation of an event and how this encompasses the event’s spatiotemporal location, the people and objects involved, and the relations between these elements. A working definition of events should consider these psychological conceptions.

The problem of understanding specific events is closely related to that of *event identity*, which considers whether two events mentioned in text are regarded as the same. Many of the definitions of event identity found in literature (e.g. Bejan and Harabagiu, 2010) were established to facilitate event coreference. Hovy et al. (2013) move beyond exact event coreference to consider the notion of *quasi-identity*. Quasi-identity refers to events which are the same in some respects, but not in others. We believe that definitions of events that restrict certain qualities are not effective for informing the quasi-identity relationship. For example, generic events can inform specific instances of that event type.

In the field of theoretical linguistics, there are many concepts that contribute to the idea of eventiveness, including aktionsart and transitivity. Vendler (1957) introduced the classification of verbs into different aspectual (aktionsart) categories, including accomplishments, achievements, activities, and states. The first three categories all correspond with the idea of events, though to varying degrees. In distinguishing between events and states, Comrie (1976) discusses the important factor that states do not require energy to maintain, while events do. Also, Talmy (2000) and Croft (2012) discuss at length the related notion of force-dynamic relations, which deals with the transmission of force between participants.

Additionally, there exists a significant overlap between the dimensions of grammatical transitivity (as a prototypical notion) and the qualities that define events. The concept of transitivity has been researched extensively within the linguistics community, primarily with the goal of understanding grammatical relationships within clauses.

Hopper and Thompson (1980) propose ten different dimensions intended to measure the notion of transitivity, which the authors define as a property of a clause that communicates how effectively an action is “transferred” or “carried-over” from agent to patient. The more effectively the activity can be carried over, the higher the transitivity.

Although Hopper and Thompson suggest that there is no single semantic notion that encompasses the nature of transitivity, they state that they have considered terms such as ‘activity’ and ‘intensity’, both of which are also relevant to the notion of eventiveness. Tsunoda (1981; 1985) adds several dimensions to the notion of transitivity as a prototype which we believe further support the relationship between transitivity and eventiveness, including genericity, completion, and realization.

Experimental work was conducted by Madnani et al. (2010) to collect information about subjects’ perception of the various transitivity dimensions for given actions. However, the authors do not examine the transitivity dimensions of mode (realis vs. irrealis), agency, and individuation of the object, all of which we regard as also playing a very important part in a predicate’s measure of “eventiveness”. Additionally, they guide the subjects through the process of measuring transitivity by prompting them for specific properties, as opposed to utilizing human intuition.

Quality	Definition
Occurrence	The quality of a predicate that signals that a change in the state of the world has taken place.
Spatio-Temporal Grounding	The degree to which the predicate is grounded in time and space.
Lexical Aspect	The boundedness and duration of the predicate with respect to time.
Agency	The degree to which the main event participant can be regarded as a “causer” or “doer”.
Affectedness	The degree to which the action of the predicate affects the semantic patient.
Actuality	The knowledge of whether the predicate actually took/takes place.
Specificity	The degree to which the predicate refers to a particular instance of an event.

Table 1: Definitions of Event Qualities

### 3 Qualities of Events

Given the wide array of definitions and descriptions of events from different perspectives, we believe that each offers a unique insight into this multi-faceted problem. We seek to identify the qualities of eventiveness and determine which are the most salient. The seven qualities we consider are listed in Table 1. In this section, we provide a detailed definition, examples, and justification as to why each quality is important to eventiveness. For our examples, we consider predicates with explicit textual indicators of the qualities. However, many predicates can possess these qualities independent of textual evidence. Additionally, although every quality is examined in isolation here, the interaction between these qualities is an important consideration. In Section 4, we describe the experiment we conducted in order to demonstrate the extent to which each quality contributes to eventiveness.

#### 3.1 Occurrence

Occurrence, the idea of something having happened, largely coincides with what we believe to be an event. In fact, the TimeML definition (Pustejovsky et al., 2003) of *event* covers situations that “happen or occur”. We consider this to be equivalent to the idea of “change in the state of the world”, because if the final state is the same as the initial state, then nothing can be said to have happened or occurred. As such, we contend that the greater the degree to which an event can be considered to have “happened” or “occurred”, the greater the amount of eventiveness it will exhibit.

Note that the determination of ‘state’ here goes beyond mere appearances: a person who bounces a ball and catches it appears to be in exactly the same state as before, but in reality, some energy has been expended. Most verbs exhibit the quality of having “occurred”, with the notable exception of statives<sup>1</sup>, which are a fairly lexically con-

<sup>1</sup>Note that TimeML has a special class of events marked as STATE.

strained category (copular verbs, many verbs of cognition, etc.). Thus, for verbal predicates, we can regard verbs that indicate an action rather than a state as having “occurred” and being eventive. In general, the more energy and motion involved in the predicate, the more eventive it is. In the example below, *running* would be considered more eventive than *sitting*.

1. He was *running* on the track. (high energy)
2. He was *sitting* in the chair. (low energy)

For nominal events, the situation is more complicated. We must distinguish the set of nouns that can indicate an event, such as “earthquake”, from the set of nouns which cannot, such as “epicenter”. For deverbal nouns, we also must distinguish between process nouns, such as “the *building* of the house”, and result nouns, as in “the *building* I work in”. In order to distinguish the quality of occurrence, we can use the diagnostic of determining whether the predicate can be appropriately associated with words such as “happened”, “took place”, or “occurred”. For example, a *presentation* event can “occur”, but the physical materials also called the *presentation* cannot be said to have “occurred”.

1. The *presentation* occurred in the boardroom.
2. \*The *presentation* slides occurred.

#### 3.2 Spatio-Temporal Grounding

Spatio-temporal grounding deals with the degree to which an event is able to be “pinpointed” to a particular time and place. We hypothesize that a predicate that is more able to be grounded in time and/or space will be perceived as being more eventive than a predicate which is less able to be grounded spatio-temporally.

Quine (1985) considers events to be individuated by their placement in space and time, which implies that any given event should be able to be associated with both a time and a place. Indeed, the close association of events with their locations and times manifests itself in our ability to refer to well-known events by their time or location,

such as *Chernobyl* or *9/11*. Another consequence of the spatio-temporal grounding of events is that one can refer to events that happen relative to other events, e.g. *before*, *after*, *nearby*. Of the following examples, the last seems most eventive.

1. He *fought* the law.
2. He *fought* the law yesterday.
3. He *fought* the law yesterday in court.

### 3.3 Lexical Aspect

Lexical aspect deals not with *when* a predicate occurs in relation to time (i.e. *tense*), but *how*. It examines, as Comrie (1976) puts it, “the internal temporal constituency of a situation”. This covers both how the event is bounded in time (telicity) and how long it lasts (durativity). A durative event can allow for increased eventiveness in that it allows for more changes in the state of the world simply because it lasts longer. At the same time, many punctual (instantaneous) events have the potential to be very eventive because they can produce large amounts of change in a very short time, therefore producing a more *drastic* change (e.g. an assassination or fatal lightning strike). Thus, both durative and non-durative events seem to be able to contribute to eventiveness in unique ways.

Regarding telicity, we believe that events which are bounded in time (i.e. having endpoints) generally evoke a more pronounced sense of eventiveness because they are more easily distinguishable from the “backdrop” of other occurrences and states. In fact, it is by definition that all events must have a beginning (otherwise, they would not be able to be referred to as “occurrences”), and we believe that event endings or markers of completion move an event even closer to a prototypical notion of “high eventiveness”.

Vendler (1957) categorizes verbs into four categories depending on their durativity and telicity: *state*, *achievement* (telic, punctual), *accomplishment* (telic, durative), *activity* (atelic, durative). Comrie (1976) adds to this the category of *semelfactive* (atelic, punctual). Examples of these categories follow.

1. He is *building* a house. (telic, durative)
2. He is *swimming*. (atelic, durative)
3. He *shot* the man. (telic, punctual)
4. He is *knocking* on the door. (atelic, punctual)

### 3.4 Agency

Agency deals with the amount of control and volition involved in an event. We regard agency as a measure of the degree to which a participant willfully executes an action and maintains control over it. As such, we assert that the greater the degree of agency attributed to the causer or performer of an predicate, the higher the eventiveness that the predicate will display. Involved in this idea are the related notions of frequency/normalcy of occurrence and causality. Consider the following.

1. The wine *aged* in a barrel. (no agent)
2. The vintner *aged* the wine in a barrel. (agent: *vintner*)

The presence of the agent causes the second predicate to seem more eventive than the first. The first implies a natural process. The second implies a volitional effort on the part of the vintner (the agent) to cause the wine to undergo this process in a particular location, likely with some control over when the aging would begin and end before being bottled. The relevance of these predicates to many applications is dependent on the existence of the agent.

Dowty (1991) lists prototypical characteristics of high and low agency. For high agency, he lists volition, sentience, effect upon another participant, and self-produced mobility. For low agency, he lists internal change, incremental theme (when something incrementally disappears or is used up), and movement induced by another participant.

We can describe “natural processes” as those occurrences which come about as a result of actions whose main participants are characterized by low agency. Most natural processes, such as “aging”, are not considered very eventive. However, distinct from these are certain natural *occurrences* that do involve movement and great effect on the world (such as earthquakes, lightning, and landslides).

We also hypothesize that the frequency or “normalcy” of predicates is related to the degree to which they are perceived as eventive. The growth of grass (low agency) is an extremely frequent and “normal” type of process (and thus should be seen as less eventive), whereas an earthquake (higher agency) is a much rarer occurrence and should therefore be seen as more eventive. Note that this factor is also highly relevant to the “newsworthiness” of the predicate.

Agency is also intricately linked to causality, since prototypical agents often cause a change of state in patients. In many cases, the agent of a particular event can itself be characterized as an event (e.g. “The *earthquake* caused three buildings to crumble”). In this example, the earthquake - while formally the agent of the “crumble” event - is itself considered to be an event.

### 3.5 Affectedness

Affectedness is the degree to which an event affects its participants, most importantly the participant in the semantic patient role of the predicate. We generally hypothesize that the more affected a patient is by the event it is a participant of, the greater the eventiveness of that predicate.

The actual manifestation of the notion of “affectedness” can take a variety of forms. First, we posit that an event can affect its patient to a greater extent if the patient is more animate. To this end, we consider a general animacy hierarchy that is a modification of the hierarchy proposed by Silverstein (1976): *Human Proper Noun* > *Human Common Noun* > *Animate Noun* > *Inanimate Noun*, e.g. *Sheila* > *woman* > *bear* > *rock*.

Second, we suggest that an event can affect its patient to a greater extent if the action that is taking place is more severe or extreme. For example, we would consider “He *killed* the man” to be more eventive than “He *wounded* the man”, simply because of the longer-lasting effect of “kill”.

Both of these notions are grounded in Hopper and Thompson’s (1980) transitivity dimensions of *Individuation of O* and *Affectedness of O*, where O generally represents the semantic patient. They contend that a particular action is able to be “more effectively transferred” to a highly individuated patient (one that is a proper noun, human or animate, concrete, singular, count, and referential/definite) than to a patient that is low in individuation (one that is common, inanimate, abstract, plural, mass, and non-referential). We believe that eventiveness has a direct correlation with patient individuation in all dimensions but one: the singular vs. plural distinction. We contend that all other things being equal, the broader the semantic patient role is, the greater the overall effect of the event (e.g. *He killed five men* as opposed to *He killed one man*), and therefore the greater the eventiveness.

1. He *punched* some pillows. (low individuation)

2. He *punched* his brother. (high individuation)
3. He *bruised* the man’s leg. (low affectedness)
4. He *broke* the man’s leg. (high affectedness)

Tsunoda (1981) notes that this affectedness is independent of the amount of agency the agent possesses: a person killed by a stray bullet is just as affected as a person who is intentionally killed.

Our experiment in this study tests primarily for individuation, and further testing is required to specifically examine Hopper and Thompson’s *affectedness of O* dimension. Additionally, future studies could examine Tsunoda’s (1981) claim that resultative predicates (e.g. *break*, *kill*) generally encode higher transitivity than non-resultative predicates (e.g. *hit*, *shoot*). We believe that such predicates should exhibit higher eventiveness because they lexically explicate the change in the world that has taken place as a result of an action. Similarly, future experiments could consider not only the patient, but also how the agent and/or other participants are affected by the action.

### 3.6 Specificity

Specificity can be defined as the degree to which a predicate refers to a particular instance (or instances) of an event, where that event must be well-grounded in time and space and well-individuated from other events. We believe that as specificity of a predicate increases, eventiveness increases as well. Thus, specific events should have higher eventiveness than habitual events (ones that recur but do not have a well-defined spatio-temporal location and/or number of occurrences), and generic events (where no specific instance is in focus).

While both habitual and generic predicates are less eventive, they differ in several ways. Habitual events typically imply that instances of the event have occurred, but with no specific information about these occurrences, whereas generic predicates refer to events that are treated more as general classes of occurrences in the world rather than individuated events. The following examples illustrate this quality.

1. The chicken laid an egg on Tuesday. (specific)
2. The chicken lays two eggs a week. (habitual)
3. Chickens lay eggs when fertile. (generic)

As noted in the example in the introduction, habitual and generic events are of great value for acquiring world knowledge that can apply to specific

instances of those events. We consider this task to be very similar to detecting when two events share quasi-identity (Hovy et al., 2013). In the above example, the habitual event gives a likely next date for egg laying, and the generic event gives us the knowledge that the chicken is fertile and therefore able to lay eggs.

We also hypothesize that the more specific the event (e.g. *lays an egg on Tuesdays and Fridays* rather than *lays two eggs a week*), the more eventive the predicate will seem. This intuition connects with recent research into detecting the difference between habitual and specific events (Mathew and Katz, 2009), where the existence of semantic arguments to the predicate contributes to specificity. Often, arguments missing from generic events would display other properties of eventiveness (e.g., agency or spatio-temporal grounding) if they were present.

### 3.7 Actuality

Actuality refers to whether an action is realis or irrealis, that is, whether or not it actually occurs. We regard actualized (realis) predicates as exhibiting a higher eventiveness than unactualized (irrealis) predicates, as the former present actual changes in the state of the world, whereas the latter posit only potential or hypothetical changes.

The notion of whether or not a predicate is actualized corresponds to the “Effectiveness Condition” parameter of *realization* (Tsunoda, 1981) in transitivity theory. A predicate’s fulfillment of the Effectiveness Condition generally correlates to a greater “completeness” of lexical meaning and also corresponds to a higher degree of affectedness of the patient.

There are a wide variety of contexts in which irrealis predicates can occur; among these, predicates may be modified by epistemic modality (*might have*), deontic modality (*hopes, orders, promises*), abilities (*is able to*), and negative polarity (*didn’t*). We also consider future tense events to be irrealis, as by definition they have not yet occurred. It should be noted that epistemic events exist between realis and irrealis, and may exhibit more eventiveness than other forms of modality.

Within the class of negative events, we can contrast simple negation events (events modified by negators such as *no* and *not*) with avoided or prevented events. Avoided events involve a conscious decision (thus requiring agency) to not perpetrate

the event. Prevented events, on the other hand, involve an external agent preventing the event from occurring. In general, the act of preventing an event from occurring is itself an event.

1. He *bought* a new car. (realis)
2. He might *buy* a new car. (future)
3. He might have *bought* a new car. (epistemic)
4. He is able to *buy* a new car. (ability)
5. He wants to *buy* a new car. (deontic)
6. He was prevented from *buying* a new car. (negative, prevention)
7. He did not *buy* a new car. (negative, simple negation)

Typically, systems which utilize events concentrate on realis events only; however, when dealing with events across documents, the information associated with irrealis predicates is very useful for establishing quasi-identity relationships. There are several motivating examples of unactualized event types that are necessary for deeper understanding of events. If a crime occurs, for instance, a particular suspect’s *ability* to commit that crime becomes relevant. Likewise, if some order is given to perform an action, and the action later occurs, the quasi-identity relationship between the “director” and the action is immediately relevant.

## 4 Experiment

In order to perform a concrete analysis of the qualities of eventiveness in the real world, we undertook a small experiment in which human participants rated the eventiveness of different predicates in context. We hypothesize that a predicate with an explicit indicator of one of these qualities would be considered more eventive than a similar predicate without that indicator.

### 4.1 Methodology

For each quality, we created one sentence with and one without explicit evidence of that quality. The two sentences utilize the same predicate and differ only in their expression of the quality of interest. For example, “He *graduated* college” possesses the positive actuality quality, while “He *promised* to *graduate* college” does not. This allows us to compare the ratings for these pairs of sentences.

The sentences were placed into example groups consisting of a pseudo-random sampling of the sentences, enforced to only have one instance of a predicate within each group. Each example group

### Mechanical Turk Instructions

**Directions:** Please rate the following words in terms of whether they indicate an event in the context of the given sentence located above each word. A rating of '5' means that it is very much an event, and a '1' rating means that it is not at all an event. Read the definition/examples below carefully before beginning.

**Definition:** An event is a cover term for situations that happen or occur. Events can be punctual (instantaneous) or last for a period of time.

#### Examples:

1. I am **building** a new house. (*building* is an event)
2. I like the Empire State **Building**. (*Building* is not an event, but an object)
3. Robert **grew** to be tall. (*grew* is an event)
4. Robert **is** tall. (*is* is not an event)

#### Question Prompt:

How much like an event does this word seem?

1 (not at all) 2 (slightly) 3 (moderately) 4 (fairly) 5 (very)

Table 2: Annotation Instructions

consisted of eight example sentences, with a total of nine example groups.

We collected the eventiveness ratings from participants on Amazon Mechanical Turk, who rated each predicate in the example group on an integer scale from one to five. We collected 50 ratings for each sentence, and participants were allowed to complete multiple example groups. Overall, we had 76 unique participants, who completed an average of 5.9 example groups each. The participants spent an average of 9 seconds rating the predicate in each sentence.

We also included a variety of control “non-events”, which included result nouns as well as statives. These exhibited statistically lower eventiveness than any of the non-control predicates.

#### 4.1.1 Instructions

We provided instructions to each participant as shown in Table 2. These instructions contain a succinct definition of an event, utilizing the TimeML terminology (Pustejovsky et al., 2003). Additionally, we provided four example sentences, two illustrating events and two illustrating non-events. One of the non-events was a stative (“is”), and the other was a result noun (“Building”). These examples illustrate that not all verbs indicate events, and that words like “building” can be events in some contexts but not others.

#### 4.2 Analysis

For our analysis, we examined the mean, variance, and ranking of the eventiveness ratings provided for each predicate by the participants. We compared pairs of sentences<sup>2</sup> based on the probability that a randomly chosen rating for the sentence with the quality would be higher than a randomly chosen rating for the sentence without the quality.

<sup>2</sup>A complete list is available by request.

The statistical significance of this probability can be assessed using a Wilcoxin-Mann-Whitney test.

For example, “He *played* piano” has a mean eventiveness rating of  $\bar{x}=4.56, \sigma=.80$ , and “He is able to *play* piano” has  $\bar{x}=3.82, \sigma=1.35$ . A random rating for *played* is 66.3% more likely to be higher than one for the ability *play*. This difference is statistically significant assuming an acceptable type-I error rate of .05%.

We present in Table 3 results for the pairs of sentences testing each quality with their probabilities. The > indicates the hypothesis that one value of the quality is more eventive than the other. The \* indicates statistical significance.

Quality	Result	Prob
Occurrence	Verb > Noun	0.604*
Occurrence	High Energy > Low Energy	0.686*
Spatial	Grounded > Not	0.526
Temporal	Grounded > Not	0.509
Agency	Agency > No Agency	0.641*
Aspect	Atelic Durative > Telic Punctual	0.628*
Aspect	Telic Durative > Atelic Durative	0.471
Affectedness	Individuated > Not	0.505
Actuality	Actual > Ability	0.663*
Actuality	Actual > Epistemic Modality	0.646*
Actuality	Actual > Volitive Modality	0.664*
Actuality	Actual > Commissive Modality	0.620*
Actuality	Actual > Directive Modality	0.642*
Actuality	Actual > Polarity	0.681*
Actuality	Past Tense > Future	0.635*
Actuality	Present Tense > Future	0.626*
Specificity	Specific > Habitual	0.667*
Specificity	Specific > Generic	0.546

Table 3: Results of Eventiveness for Qualities

#### 4.3 Discussion of Results

As shown in Table 3, many of the factors that have been identified in various theoretical descriptions of eventiveness can be shown experimentally to affect people’s perception of the eventiveness of a predicate in a sentence. Below, we discuss the positive results, where our hypotheses were confirmed, as well as the negative results.

For occurrence, agency, and actuality, we found strong evidence that these qualities contribute to eventiveness. For example, “The *attack* happened at dawn” was less eventive than “They *attacked* at dawn”, “The fire *started*” was less eventive than “He *started* the fire”, and “He hopes to *graduate college*” was less eventive than “He *graduated college*”. For actuality, realis predicates were always more eventive than irrealis predicates. An ANOVA test indicated no significant difference between the different forms of irrealis (modality, negation, etc.).

Results for the other qualities were slightly more mixed. For aspect, we found that activities were more eventive than achievements, but contrary to expectation, accomplishments were not more eventive than activities. For specificity, there was a clear distinction between specific and habitual predicates, but no distinction between specific and generic predicates. Our example of a generic predicate, “Football fans watch the Superbowl” could be considered either a generic event or a present tense description, and this might have confused the results. Also, since the definition provided for *event* gave only singular event examples, this may have biased the results in this case.

For spatio-temporal grounding, there was no significant effect. We believe that this is due to the implicit eventive nature of some verbs. We analyzed the predicate “fought”, which was equally eventive with and without a specified time or location. However, such a verb does not require explicit grounding; the reader can assume that any given fight happens at a specific time and location. For affectedness, our examples utilized a verb that is always highly indicative of affectedness and did not adequately capture a good distinction between high and low affectedness. We believe that future experiments can control for these kinds of cases and that example predicates can be found that will isolate the specific qualities.

Another concern is that our design only explicitly tested a single predicate for each quality. However, the nature of the predicates and the sentences we used allowed for post-hoc analysis of the qualities that existed across more than two sentences. Empirical testing showed the same pattern of results across predicates for these qualities.

Overall, the experimental results are extremely interesting in their congruence with the literature on events, but further research is required to deter-

mine the exact contribution of each quality. The current experimental design lacks sufficient power to reliably rank the qualities due to contrast effects within example groups. It is likely that the ordering/grouping of the examples affected the rating of individual examples. In future studies, we plan to control for these effects by controlling the ordering of the examples given to each individual.

## 5 Conclusion

Working definitions of events are often ill-defined and difficult to apply. We have laid out a series of qualities which contribute to the overall eventiveness of a predicate in a sentence. Our findings indicate that the degree to which a predicate is considered an event is a function of these qualities. Evidence for these qualities was validated using participant ratings of predicates.

When developing annotated corpora of events, the decision of whether or not to consider an individual predicate as an event is difficult. Understanding the qualities of eventiveness can explain why one predicate seems less eventive than another (e.g. irrealis, generic).

Instead of deciding each predicate on the basis of the individual qualities being exhibited, annotation specifications should consider how these qualities interact. Drawing an explicit boundary between events and non-events can cause information contained in the non-events to be lost for reasoning. Along the same lines, event extraction capabilities could be greatly improved by the labelling of these qualities on annotated corpora. This would enable event extraction to preserve the fine-grained distinctions between events that are shown to be relevant to human understanding.

In this study, we gave examples of how predicates with lesser eventiveness can provide valuable insight into problems such as event coreference and quasi-identity resolution. These qualities of eventiveness can serve to inform future research into those areas, providing a deeper understanding of the meaning of event coreference. While different applications have different needs, understanding the qualities that contribute to eventiveness will enable applications to more intelligently utilize event information.

## 6 Acknowledgements

This work was sponsored in part by the Air Force Research Laboratory (AFRL).



## References

- ACE. 2005. In *ACE (Automatic Content Extraction) English Annotation Guidelines for Events Version 5.4.3 2005.07.01*.
- Cosmin Bejan and Sanda Harabagiu. 2010. Unsupervised event coreference resolution with rich linguistic features. In *Proceedings of the 48th Annual Meeting of the Association for Computational Linguistics*, pages 1412–1422, Uppsala, Sweden, July. Association for Computational Linguistics.
- Bernard Comrie. 1976. *Aspect: An introduction to the study of verbal aspect and related problems*. Cambridge University Press.
- William Croft. 2012. In *Verbs: Aspect and Causal Structure*, Oxford. Oxford University Press.
- David Dowty. 1991. Thematic proto-roles and argument selection. In *Language*, Vol. 67, No. 3., pages 547–619, September.
- Paul J. Hopper and Sandra A. Thompson. 1980. Transitivity in grammar and discourse. In *Language* 56 (2), pages 251–299, June.
- Eduard Hovy, Teruko Mitamura, Felisa Verdejo, Jun Araki, and Andrew Philpot. 2013. Events are not simple: Identity, non-identity, and quasi-identity. In *Proceedings of the NAACL HLT 2013*.
- TAC KBP. 2014.  
<http://www.nist.gov/tac/2014/kbp/event/>.
- Lawrence B. Lombard. 1986. In *Events: A Metaphysical Study*, London. Routledge and Kegan Paul.
- Nitin Madnani, Jordan Boyd-Graber, and Philip Resnik. 2010. Measuring transitivity using untrained annotators. In *Creating Speech and Language Data With Amazon's Mechanical Turk*, Los Angeles, CA.
- Thomas A. Mathew and E. Graham Katz. 2009. Supervised categorization of habitual versus episodic sentences. Dissertation. Georgetown University.
- James Pustejovsky, Jose Castano, Bob Ingria, Roser Sauri, Rob Gaizauskas, Andrea Setzer, and Graham Katz. 2003. Timeml: Robust specification of event and temporal expressions in text. In *Proceedings of the Fifth International Workshop on Computational Semantics (IWCS)*.
- W. V. O. Quine. 1985. Events and reification. In *E. LePore and B. P. McLaughlin, eds., Actions and Events: Perspectives on the philosophy of Donald Davidson*, pages 162–171, Oxford: Blackwell.
- Gabriel A. Radvansky and Jeffrey M. Zacks. 2011. Event perception. pages 608–620.
- Michael Silverstein. 1976. Hierarchy of features and ergativity. In *Grammatical Categories in Australian Languages*, pages 112–171, Canberra. Australian Institute of Aboriginal Studies.
- Leonard Talmy. 2000. Force dynamics in language and cognition. In *Toward a Cognitive Semantics - Vol. 1*, Cambridge, Mass. The MIT Press.
- Tasaku Tsunoda. 1981. Split case-marking patterns in verb-types and tense/aspect/mood. In *Linguistics* 19, no. 5-6, pages 389–438.
- Tasaku Tsunoda. 1985. Remarks on transitivity. In *Journal of Linguistics* 21, pages 385–396.
- Zeno Vendler. 1957. Verbs and times. In *The Philosophical Review*. Vol. 66 No. 2, pages 143–160, April.