

# Event Nominals: Annotation Guidelines and a Manually Annotated Corpus in French

Béatrice Arnulphy, Xavier Tannier, Anne Vilnat

LIMSI-CNRS & Univ. Paris-Sud  
91403 Orsay, France  
firstname.surname@limsi.fr

## Abstract

Within the general purpose of information extraction, detection of event descriptions is an important clue. A word referring to an event is more powerful than a single word, because it implies a location, a time, protagonists (persons, organizations. . .). However, if verbal designations of events are well studied and easier to detect than nominal ones, nominal designations do not claim as much definition effort and resources. In this work, we focus on nominals describing events. As our application domain is information extraction, we follow a named entity approach to describe and annotate events.

In this paper, we present a typology and annotation guidelines for event nominals annotation. We applied them to French newswire articles and produced an annotated corpus. We present observations about the designations used in our manually annotated corpus and the behavior of their triggers. We provide statistics concerning word ambiguity and context of use of event nominals, as well as machine learning experiments showing the difficulty of using lexicons for extracting events.

**Keywords:** event nominals, manual annotation, annotation guidelines

## 1. Introduction

Detection of event descriptions is important in many information extraction applications (*e.g.*, temporal ordering of events on a chronological axis, improving of question-answering systems). If verbal designations are well studied and easier to detect than nominal ones, studies on nominal designations lack both definition effort and resources. In French, only one corpus containing annotated event nominals is available. The FR-TimeBank (Bittar, 2010) was annotated following the TimeML annotation guidelines. These guidelines allow to annotate event nouns, but only to some extent: event nouns are not the main issue in those guidelines.

In our work, we only focus on nominals describing events. We rule out verbs and temporal expressions that are represented by other categories than noun. This work is in line with a named entity approach we follow; in this perspective, the TimeML annotated corpus does not fit our objectives. According to these observations and because we need annotated resources, we annotated our own corpus.

We consider that an event is what happens, corresponding to a change of state. It can be either recurring or unique, predicted or not. It may last a moment or be instantaneous. It can also occur indifferently in the past, the present or the future.

In this article, we first present our annotation guidelines and event typology, according to this definition. We then present our annotated corpus of news articles, and finally present observations about the designations used in the corpus and the behavior of event triggers (words which can be a clue for the detection of the event reading of nouns).


## 2. Annotation Guidelines

Our annotation guidelines detail a typology of events (which is not especially focused on nouns; it could fit to all the designations of events), as well as instructions for

deciding whether a noun or a noun phrase is an event or not. These guidelines are not language-dependent, even if we mostly worked on French when we edited the manual.

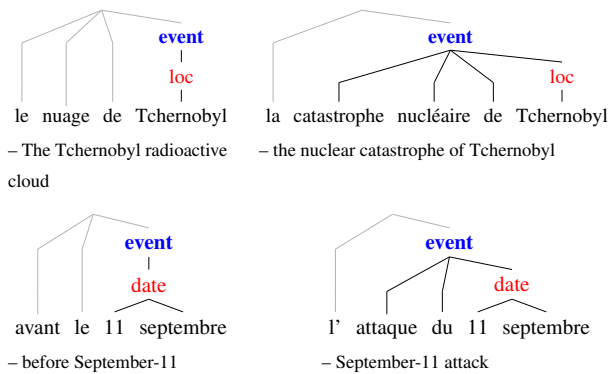
### 2.1. Annotation of Events and other Tags

Our guidelines are complementary to the Quaero project annotation guidelines for named entities (Rosset et al., 2011). The general aim of the Named Entity track from Quaero project is to build a knowledge base which would reference the named entities and the relations between all their occurrences (as an example, “Barack Obama” and “US president” refer to the same person in topical documents). Within the framework of this project, we are interested in the nominal mention of events. Because our work is in line with existing named entities annotations, our event nominals annotation must respect the neighborhood of the other entities which are considered. Our event annotation overlaps with the other entities or between events, as in the following example.



This illustrates that a location can be a part of an event and that an event can be composed of another one (the recurring *festival de Cannes* and the instantiation of this recurring event: the 60th one).

Moreover, other named entity types can be events by the means of metonymy (*e.g.*, location (Lecolle, 2009) or date (Calabrese Steimberg, 2008)).



## 2.2. Typology

Our typology constitutes the main point of the annotation guidelines. For each event, the following information can be filled out: *modality*, *frequency*, *referential time* of this event, as well as whether the event *designation* is from a reported speech or not.

### 2.2.1. Modality

- **Factual event**, an event that actually happens (in the past, the present or the future).

(1) Laurent Cantet and his movie “Entre les murs” won **<EVENT> 2008 Cannes festival Golden Palm </EVENT>**

- **Hypothetical event**, an event that (from the speaker’s point of view) can be supposedly past (examples 2 and 3), or expected to happen in the future (example 4).

(2) Firefighters may have died during **<EVENT> duty </EVENT>**.

(3) Land use may have been responsible for the **<EVENT> 1930s dust bowl </EVENT>**.

(4) The **<EVENT> summit </EVENT>** should take place in December in Prague

- **Nonfactual event**, an event that never happened

(5) The **<EVENT> summit </EVENT>** has been canceled.

(6) The **<EVENT> alleged attack </EVENT>** had resulted in a lightning strike of the train line.

- **Abstract event**, a general event without any reference to a particular instantiation.

(7) The **<EVENT> crisis </EVENT>** follows an excessive confidence period.

### 2.2.2. Frequency

- **Unique event**, an event that only occurs once.

(8) The book relates the story of **<EVENT> Gandhi’s assassination </EVENT>**.

- **Recurring event**, an event that can occur several times, often periodically.

(9) The **<EVENT> Olympic Games </EVENT>**

(10) During his reign, a **<EVENT> sale of slaves </EVENT>** was held every month

- **Instantiation of a recurring event.**

(11) The **<EVENT> Atlanta 1996 Olympics Games </EVENT>**

(12) **<EVENT> 17 years of Cannes Film Festival </EVENT>**

### 2.2.3. Time

Events are anchored to time, they must be characterized by the time of their achievement. We annotated them according to the utterance time (even in case of reported speech).

- **Before**

(13) They resulted in a **<EVENT> failure </EVENT>**.

(14) Men of his generation did the **<EVENT> Second World War </EVENT>**.

- **Now**

(15) The **<EVENT> Paralympic Games </EVENT>** are always held on the sidelines of the **<EVENT> Olympic Games </EVENT>**.

- **After**

(16) The **<EVENT> France-China summit </EVENT>** will be held in December in Prague.

## 2.3. Event or not Event: not trivial

Two major problems deserve specific instructions: disambiguation and boundaries.

Concerning the disambiguation between eventive or non-eventive reading of noun phrases (NPs), the guide contains advices such as the followings:

- Imagine some non-ambiguous appropriate substitutes for the noun. If no substitution can be done, then we consider it as an event. This proved to be very effective.

In the following example 17, “*treaty*” cannot be replaced in the text by the object “piece of paper”, but rather by the action/event “modification”.

(17) He appeared to accept **<EVENT> a border demarcation treaty </EVENT>**.

In example 18, “*triathlon*” can be replaced by “*race*”, while in example 19 it is rather replaced by “*bike*”.

(18) France is one of the first **<EVENT> triathlon </EVENT>** organizers.

(19) Triathlon **<EVENT> races </EVENT>** vary in distance.

- Take inspiration from examples of eventive and non-eventive uses of the same word, that can be found in dictionaries, together with their proper definition.

- Remember that the items of enumerations are often (but not always) of the same class.

(20) The police presence is reinforced to be use for “<EVENT> *special operations* </EVENT>”, <EVENT> *arrests* </EVENT> and <EVENT> *targeted murders* </EVENT>

In example 21, one of the items enumerated is surely an event (“*police violence*” refers to the acts of violence of policemen, in context). “*Racism*” is definitely not an event. “*Unhealthy gear*” is ambiguous: As it is a chain of events, we annotate it as an event. Two out of three phrases are events in this enumeration.

(21) Since 1995, these children of East London continue to denounce racism, <EVENT> *police violence* </EVENT> and <EVENT> *unhealthy gear of suburban sets* </EVENT>.

- When the decision seems impossible, prefer not to annotate as event.

Delimiting the event boundaries is a difficult issue and the guidelines provide instructions to solve this problem. We decided to annotate according to syntax. We keep in the event tag only the nominal dependencies such as adjectives and also spatial and temporal complements, but we do not include the relative or infinitive clauses.

#### 2.4. Events and Existing Named Entities

Some event designations are very specific and specifically interesting. By the way of what happen, existing named entities (which already refer to another type of entity) become the name of an event. Locations and dates are particularly concerned.

In example 22, “Chernobyl” refers to the event, but in example 23 we talk about the disaster which happened in the location of the nuclear central, called Chernobyl.

(22) <EVENT> *Chernobyl* </EVENT> fallout

(23) <EVENT> *Chernobyl disaster* </EVENT>

In example 24, “September-11” refers to the event, but in example 25 it refers to the date of the event, inside the longer expression which name the event.

(24) Costs raised since <EVENT> *September-11* </EVENT>

(25) <EVENT> *September-11 attacks* </EVENT>

Linguistic tests allow us to differentiate date and events (Ehrmann and Hagège, 2009). For example, introduce the phrase “*the event that happened on*” before the date will help to say if the date is only a date or if this date has an eventual reading. If the new sentence is a paraphrase of the first one, then it is an event. Otherwise, it is a common date.

(26) Costs raised since <EVENT> *September-11* </EVENT>  
= Costs raised since *the event that happened on September-11*  
→ the date “September-11” is an event.

(27) <EVENT> *September-11 attacks* </EVENT>  
\* *the event that happened on 11-September attacks*  
→ the date “September-11” is not an event, it is a date included in the event name “September-11 attacks”.

#### 2.5. More Examples

(28) At least the last <EVENT> *G20* </EVENT> did not try to save the world. (metonymic use, “G20 summit”).

(29) G20 members are: [...] (here, “G20” is the organization, not the event)

(30) These <EVENT> *three hours of music* </EVENT> were amazing.

(31) I thought it was like a second <EVENT> *Pearl Harbor* </EVENT>.

(32) He blamed them for deadlock in peace talks (“deadlock” is not an event but a state).

### 3. Corpus

We chose to annotate a corpus of news, because of their high density of nominal events. Journalists name the events: when an event becomes important in a community, it is in the news. Journalists give names to the events, which often freeze until they become idioms (Tannier et al., 2012). We annotated 192 French news articles from *Le Monde* and *L’Est Républicain* (47,646 words), for a total of 1844 events (see Table 1). As a comparison, there are 3695 event nouns in IT-TimeBank (Russo et al., 2011), 1,579 in the corpus from (Creswell et al., 2006) and 663 in FR-TimeBank (Bittar, 2010). English TimeBank 1.2 (Pustejovsky et al., 2003) contains 1,792 non-stative nominal events.

	Le Monde	FR-TimeBank	Total
Texts	83	109	192
Words	31,449	16,197	47,646
<b>Events</b>	<b>1107</b>	<b>737</b>	<b>1844</b>

Table 1: Our manually-annotated corpus

Only considering the heads of the event NPs and only for the event tag (no subtype of the typology), and following the guidelines, the two annotators (the authors of the guidelines) obtained a good agreement ( $\text{Kappa}^1 = 0.808$ ). The corpora was revised on all the documents commonly annotated by the two annotators.

**Comparison with FR-TimeBank.** Among our manually annotated corpus, 109 texts are common with FR-TimeBank (Bittar, 2010). The TimeML aims followed by Bittar and ours are different: TimeML is mostly dedicated to verbal events and such FR-TimeBank is annotated with more temporal expressions: prepositions, verbs, dates, and

<sup>1</sup>We used the Cohen’s Kappa coefficient (Cohen, 1960). This measure compares the agreement against what might be expected by chance. According to Landis and Koch (Landis and Koch, 1977), from 0.6 to 0.8 is what we consider a good agreement. Up to 0.8 is a very good agreement.

temporal and aspectual links between the tagged entities, etc. Nouns are not their principal interest. The result of the annotation is therefore definitely different.

For event nominals annotation, the most important differences between our annotation and the FR-TimeBank can be summarized in: we do not tag the name of states, Bittar does; we tag the named entities (such as a name of war “Arab-Israeli war”), he does not. Though based on heads of event NPs only, the two annotations seem quite similar given to the good inter-annotator agreement ( $\text{Kappa} = 0.704$ ).

#### 4. Observations

We used the whole manually-annotated corpus to compute statistics and to make observations about the words describing nominal events. Our corpus is thus composed of 725 different occurrences of head nouns among a total of 1844 annotated events. We also noticed that 269 of these triggers occur only once in the entire corpus.

##### 4.1. Behavior of the Event Denominations

Among the nouns that appear more than once in the corpus, only 31% denote events every time they occur. The others have a quite regular distribution. As an example, 56 nouns have an eventive reading in less than 20% of their occurrences, and 129 in less than 50%. Figure 1 gives more information on the proportion of events triggered by each noun in the list, while Table 2 provides a few relevant examples.




Figure 1: Progression of the number of eventive head nouns according to the rate of occurrences of these nouns that have an eventive reading. For example, 29 nouns have an eventive reading in less than 10% of their occurrences while 312 nouns have an eventive reading in less than 100% of their occurrences.

##### 4.2. Pluralization and determiners

As Russo et al. (2011) did on an Italian corpus, we attempted to check two general assumptions of the literature:

- Plural instances rarely have an eventive reading.
- Definite articles generally involve an eventive reading.

For French, we lead to the same conclusion than Russo et al. Indeed, these general assumptions appear not to be verified in our corpus analysis, as shown by Tables 3 and 4.

Example	Translation	Rate
disparition	disappearance	100%
meurtre	murder	100%
démission	resignation	100%
campagne	campaign/country	88.0%
peine	punishment/sadness	88.2%
vote	vote	80.0%
commentaire	comment	66.7%
bombe	bomb	50.0%
signe	sign	44.4%
mort	death/dead	37.5%
prix	price/award	22.2%
conseil	advice/council	10.7%

Table 2: Examples of nouns having (sometimes or always) an eventive reading, together with the rate of their eventive reading in the corpus

The proportions of each class are in fact quite the same for events and non-events. The only notable difference is the frequent use of indefinite articles for events, rather than definite articles, which is quite contrary to the general intuition.

	Event nouns	All nouns
Singular	80.1%	83.4%
Plural	19.9%	16.6%

Table 3: Rates of singular and plural occurrences

	Event nouns	All nouns
Definite article	27.9%	19.9%
Indefinite article	14.3%	6.2%
Demonstrative	4.0%	1.7%
Possessive	6.1%	3.3%

Table 4: Rates of different types of determiners introducing nouns

##### 4.3. Using event head nouns as lexicon

In order to show the difficulty of using lexicons in event extraction, we built a lexicon from the event head nouns annotated in the corpus. Two lists have thus been extracted:

1. The list of nouns that have an eventive reading in 100% of their occurrences in the corpus ( $LEX_{sure\_only}$ ).
2. The list of nouns that have an eventive reading in at least one of their occurrences in the corpus ( $LEX_{all}$ ).

###### 4.3.1. On the whole corpus

We performed two first runs on the whole corpus, applying  $LEX_{sure\_only}$  and  $LEX_{all}$  on the texts. The way lexicons are built ensures a precision of 1 for  $LEX_{sure\_only}$  and a recall of 1 for  $LEX_{all}$ . All results are given in Table 5.

	Precision	Recall	F-measure
<i>LEX<sub>sure_only</sub></i>			
Le Monde	1.000	0.434	0.605
FR-TimeBank	1.000	0.403	0.574
Total	<b>1.000</b>	<b>0.421</b>	<b>0.593</b>
<i>LEX<sub>all</sub></i>			
Le Monde	0.322	1.000	0.487
FR-TimeBank	0.385	1.000	0.556
Total	<b>0.345</b>	<b>1.000</b>	<b>0.513</b>

Table 5: Results with lexicons built on the entire corpus

These values should not be taken as real test results, as the same corpus was used for extracting the lexicons and obtaining the event annotations. They rather show the best values that could theoretically be obtained when using “perfect” lexicons only.

However, these statistics, as well as the rate of event trigger occurrences that really have an eventive reading (Figure 1), confirm that lexicons can be useful, but are far from enough to extract event nominals. Indeed, the recall for *LEX<sub>sure\_only</sub>* is only 42%, which means that 58% of the occurrences are ambiguous. A similar conclusion will be drawn in next Section.

#### 4.3.2. On a test corpus

The operation leading to the lexicons *LEX<sub>sure\_only</sub>* and *LEX<sub>all</sub>* has also been achieved on a development corpus made of 75% of the whole corpus, leading to two new lists *LEX<sub>sure\_only</sub><sup>dev</sup>* and *LEX<sub>all</sub><sup>dev</sup>*. These lexicons have then been applied on the 25% remaining documents, considered as the test corpus.

Results are presented in Table 6. As expected, all values decrease in comparison with the whole corpus. Moreover the drop is important (between 5 and more than 30 points), showing that the lexicons built on the development corpus can hardly be considered as representative. This can have two different interpretations: either the corpus is too small, or too many nouns can express an event in a very specific context, so that building a “complete” lexicon is illusory (recall that 31% of the events in the entire corpus have only one occurrence).

	Precision	Recall	F-measure
<i>LEX<sub>sure_only</sub><sup>dev</sup></i>			
Le Monde	0.797	0.268	0.401
FR-TimeBank	0.776	0.254	0.383
Total	<b>0.789</b>	<b>0.263</b>	<b>0.395</b>
<i>LEX<sub>all</sub><sup>dev</sup></i>			
Le Monde	0.287	0.672	0.402
FR-TimeBank	0.323	0.683	0.439
Total	<b>0.299</b>	<b>0.676</b>	<b>0.415</b>

Table 6: Results on a test corpus with lexicons built on the development corpus

## 5. Conclusion

In this article, we present our definition of event and place our work among the TimeML ones. We propose a typology concerning events taking into account modality, frequency and referential time. This typology is the basis for our annotation guide for nominal events, which we wrote in order to develop the resource we need for automatic extraction of event nominals. We also present our manually annotated corpus of French event nominal designation. Using the corpus, we were able to make interesting observations on the behavior of event nominals and the way they are used in news articles. Our manually annotated corpus is distributable upon request.

Such a corpus is a good resource for the evaluation of an automatic event nominals extraction system on French. Indeed, we are working in this way.

## 6. Acknowledgement

This work has been partially funded by OSEO, French State agency for innovation and research, as part of the Quaero program.

## 7. References

- André Bittar. 2010. *Building a TimeBank for French: A Reference Corpus Annotated According to the ISO-TimeML Standard*. Ph.D. thesis, Université Paris Diderot.
- Laura Calabrese Steimberg. 2008. Les héméronymes. Ces évènements qui font date, ces dates qui deviennent évènements. *Mots. Les langages du politique*, 3:115–128.
- Jacob Cohen. 1960. A Coefficient of Agreement for Nominal Scales. *Educational and Psychological Measurement*, 20(1):37–46.
- Cassandra Creswell, Matthew J. Beal, John Chen, Thomas L. Cornell, Lars Nilsson, and Rohini K. Srihari. 2006. Automatically Extracting Nominal Mentions of Events with a Bootstrapped Probabilistic Classifier. In *Proceedings of the COLING/ACL 2006 Main Conference*, pages 168–175, Sydney, Australia, jul. Association for Computational Linguistics.
- Maud Ehrmann and Caroline Hagège. 2009. Proposition de caractérisation et de typage des expressions temporelles en contexte. In *16ème*, Senlis, France, June. ATALA.
- J R Landis and G G Koch. 1977. The Measurement of Observer Agreement for Categorical Data. *Biometrics*, 33.
- Michelle Lecolle. 2009. Éléments pour la caractérisation des toponymes en emploi événementiel. In Ivan Evrard, Michel Pierrard, Laurence Rosier, and Dan Van Raemdonck, editors, *Les sens en marge - Représentations linguistiques et observables discursifs : actes du colloque international de Bruxelles, 3-5 novembre 2005*, pages 29–43. L’Harmattan, Paris, nov.
- James Pustejovsky, José Castaño, Robert Ingria, Roser Saurí, Robert Gaizauskas, Andrea Setzer, and Graham Katz. 2003. TimeML: Robust Specification of Event and Temporal Expressions in Text. In *IWCS-5, Fifth International Workshop on Computational Semantics*.

- Sophie Rosset, Cyril Grouin, and Pierre Zweigenbaum, 2011. *Entités Nommées Structurées : guide d'annotation Quaero*. LIMSI-CNRS, Orsay, France. <http://www.quaero.org/media/files/bibliographie/quaero-guide-annotation-2011.pdf>.
- Irene Russo, Tommaso Caselli, and Francesco Rubino. 2011. Recognizing deverbal events in context. In *Proceedings of 12th International Conference on Intelligent Text Processing and Computational Linguistics (CI-Ling 2011), poster session*, Tokyo, Japan, feb. Springer.
- Xavier Tannier, Véronique Moriceau, Béatrice Arnulphy, and Ruixin He. 2012. Evolution of Event Designation in Media: Preliminary Study. In *Proceedings of the Eighth International Language Resources and Evaluation (LREC'2012)*, Istanbul, Turkey, May.