



Fires and blizzards

Syntagmatic cues for event nouns in Italian

Irene Russo * – Tommaso Caselli °

*University of Pisa

° ILC – CNR Pisa

{irusso@ling.unipi.it; tommaso.caselli@ilc.cnr.it}



Aim of our work

- What is an event noun?
 - Previous research (Grimshaw 1990, Zucchi 1993) on the status of event nouns has concentrated on nominalizations
 - The notion of event noun is a scalar concept (Simone, 2008)
- Which factors can potentially determine the emergence of a noun as eventive apart from morphological marking?
 - Measure for eventivity relying on syntagmatic cues



Outline

- Theoretical background (limitations and difficulties)
- Corpus analysis to define our working hypothesis
- Corpus-based measure of eventivity
- Evaluation of its soundness with respect to other speakers' judgments and lexicographic resources



Event noun and morphological derivation

- Morphological derivation is one of the factor that contribute to the individuation of an event noun
- Deverbal nominals such as *translation* are semantically ambiguous (action/event vs. result) (Bisetto & Melloni 2007) :
 - 1a. La traduzione di questo testo è piena di errori.
'The translation of this text is full of mistakes.'
 - 1b. Molte traduzioni sono piene di errori.
'Many translations are full of mistakes'
 - 2. La traduzione è sul tavolo.
'The translation is on the table.'

suffix	nouns	non eventive
-zione	12	9
-mento	12	10
-tura	12	8
-aggio	12	2
	48	29



The role of syntagmatic cues

- Since Grimshaw's (1990) characterization of nominalizations, syntagmatic co-occurrences are cues to disambiguate between an eventive and a non-eventive use of the same noun:

Process:

- 3a. La costruzione del palazzo è durata due anni.
'The building (of the house) took two years'

Result:

- 3b. La costruzione è alta due piani.
'The building is two floors high'



The role of syntagmatic cues

“There are certain nouns that are not verb derivatives, yet behave like nominalised verbs; that is, they can enter container contexts without suggesting suppressed nominals. Fires and blizzards, unlike tables, crystals, or cows, can occur, begin, and end, can be sudden or prolonged, can be watched and observed – they are, in a word, events and not objects.” (Vendler 1967: 141)



The role of syntagmatic cues

Traditional semantic analyses (Gross and Kiefer, 1995) exploit a limited set of tests:

- 4a. The frequent trips were a nuisance.
- 4b. The destruction of the city in only two days appalled every one.
- 4c. During the party, John left.

- 5a. I rarely attend beer festivals. (BNC)
- 5b. He wished to attend a workshop in Hawaii. (BNC)

However, these tests are not fully satisfying because they don't assess the ultimate import of all these syntagmatic cues in determining to what extent a noun denotes an event.

Distributed and contextually derived representations obtained by syntagmatic evidence can be used to make automatic identification of semantically similar words and categorical distinctions among semantic classes (Lin, 1998; Boleda et al. 2004)



The role of syntagmatic cues

noun	frequency	cominciare	continuare	precedente	successivo
guerra	133072	202	246	34	19
trattativa	51517	172	189	27	25
campagna	56937	166	62	40	11
partita	89172	145	15	88	51
campionato	36762	116	11	28	25
avventura	19428	101	35	14	11
attività	84840	62	112	49	16
riunione	52559	55	17	105	112
colloquio	28830	48	36	17	36
crisi	107664	60	33	56	39
battaglia	49687	125	194	4	7
offensiva	10215	21	60	0	5



The role of syntagmatic cues - *cominciare*

avventura	0.519	pestaggio	0.410
trattativa	0.334	montaggio	0.170
campionato	0.316	monitoraggio	0.146
campagna	0.292	operazione	0.137
battaglia	0.252	riciclaggio	0.105
offensiva	0.205	marcatura	0.094
colloquio	0.167	boicottaggio	0.093
partita	0.163	lettura	0.092
guerra	0.152	trasferimento	0.091
riunione	0.105	manifestazione	0.070
attività	0.074	risanamento	0.059
crisi	0.056	sabotaggio	0.057



The role of syntagmatic cues - *successivo*

riunione	0.213	passaggio	0.334
colloquio	0.125	pestaggio	0.082
campionato	0.069	riapertura	0.062
partita	0.058	dichiarazione	0.054
avventura	0.057	monitoraggio	0.048
offensiva	0.049	montaggio	0.048
trattativa	0.049	elezione	0.047
crisi	0.036	pagamento	0.026
campagna	0.020	sondaggio	0.019
attività	0.019	movimento	0.019
guerra	0.015	fallimento	0.018
battaglia	0.014	fornitura	0.016

A matter of polysemy

- The degree of eventivity for a deverbal noun could intuitively depends on the number of non-eventive senses codified.

noun	Sense s	Ev_senses	cominciare	continuare	precedente	successivo	sum
pestaggio	1	1	0.410	0.493	0	0.082	0.986
monitoraggio	1	1	0.146	0.440	0.048	0.048	0.684
boicottaggio	1	1	0.093	0.422	0	0	0.516
elezione	2	2	0.033	0.004	0.284	0.047	0.370
dichiarazione	3	1	0.021	0.029	0.25	0.054	0.359
operazione	3	3	0.137	0.056	0.101	0.015	0.311
montaggio	2	2	0.170	0	0.024	0.048	0.243
manifestazione	4	4	0.07	0.099	0.041	0.010	0.223
amministrazione	3	1	0.005	0.001	0.199	0.001	0.207
lettura	3	1	0.092	0.076	0.025	0.010	0.204
risanamento	4	4	0.059	0.126	0	0	0.185
produzione	3	2	0.053	0.051	0.055	0.016	0.177
trasferimento	2	2	0.091	0.012	0.012	0.006	0.121
sabotaggio	3	3	0.057	0.057	0	0	0.115



Identification of syntagmatic cues

Verbs for nouns in subject position	Verbs for nouns in DO position	Pre- and post-nominal adjectives	Light verbs
avvenire continuare restare risultare	annullare auspicare cessare evitare ostacolare	anticipato attuale imprevisto odierno successivo	fare dare mettere



How relevant are the syntagmatic cues?

WORKING HYPOTHESIS:

➤ Hoey's (2005) notion of priming: every word is mentally primed for collocational use and our knowledge of it includes its cooccurrences features → A lexical item is primed for eventivity when its cooccurrences with specific triggers are statistically relevant.

Comparison of statistical correlation between:

➤ normalized frequencies of the syntagmatic cues

AND

➤ speakers' judgements on **eventivity**

➤ lexical resources (ItalWordNet and De Mauro Dictionary)



Testing the syntagmatic cues

Test data: 200 nouns randomly extracted from the Italian TreeBank

Two parameters for noun extraction:

- Morphology: morphologically marked nominalizations vs. morphologically unmarked: e.g. *costruzione* [building], *frenata* [braking], *avvio* [start], *disegno* [drawing]
 - Morphological nouns have been extracted by keeping into account the suffix productivity (Gaeta, 2004)
- Non-derivational nouns e.g. *barone* [baron], *guerra* [war]...

NOUN FORMS	# of instances
Morphologically marked nominalizations	56
Morphologically unmarked nominalizations	43
Non-derivational nouns	101



Testing the syntagmatic cues (2)

Noun	Eventivity – syntagmatic cues	Eventivity – speakers' judgements	Eventivity - LRs
biotecnologia [biotechnology]	0	1.0	0
assemblea [meeting/assembly]	0.040	2.857	0.125
relazione [relationship]	0.015	2.857	0.625
aborto [abort]	0.102	5.0	0.433

Testing the syntagmatic cues - Normalized frequencies (3)

A preliminary analysis of the normalized frequencies of the test data (200 nouns):

- an empirical support to the claim of the existence of a *continuum* of eventivity for nouns
- a varied distribution of the nouns along the *continuum* into 3 main groups: (i.) *non eventive*, (ii.) *fuzzy*, (iii.) *eventive*
- the fact that the eventive pole of the continuum is composed by
 - *Purely eventive nouns; e.g. pestaggio [beating], sconfitta [defeat], avvio [start]...*
 - *Dot types of the kind [EVENT • NOT_EVENT] – 45/88 (51.13%); e.g. dichiarazione [declaration], incremento [increase]...*

Testing the syntagmatic cues - Computing the Spearman coefficient (4)

Correlation with the syntagmatic cues	Spearman coefficient value
Speakers' judgments	.731
ItalWordNet	.516
De Mauro Dictionary	.607

→ high and statistically relevant correlation with the speakers judgments

→ speakers' judgments correlate more with non-morphologically derived nouns (.542)

→ good values also with the lexical resources BUT less relevant due to the incompleteness and low frequent eventive readings



Identifying a measure for eventivity

On the basis of the analyses presented we have identified a statistical measure of eventivity based on the normalized frequencies

$$\text{Eventivity} \geq 0.01$$

- How reliable is the measure?
- Can it be used to tag event nouns in corpora?



Testing the eventivity measure

Annotation experiment on a corpus of Italian

→ 149 newspaper articles (Italian TreeBank and PAROLE)

→ 63.397 tokens

→ 18.308 tokens labelled as NOUNS

6 human annotators – BA students in Computational Linguistics

Use of the TimeML specifications (Pustejovsky et al 2003) adapted to Italian



Testing the eventivity measure (2)

- 4369 noun tokens have been annotated as events
- Annotation performance:
 - P&R = 0.87
 - kappa = 0.86

Creation of 2 stop-word lists from SIMPLE/CLIPS:

- Always eventive items: nouns belonging to semantic type PHENOMENON
- Items never eventive: noun belonging to semantic types TIME and AMOUNT



Testing the eventivity measure (3)

Lemmatization of the nouns in the corpus:

- 811 lemmas not annotated (**non-eventive nouns**)
- 485 lemmas annotated as eventive and on which annotators agree (**eventive nouns**)
- 78 lemmas on which annotators disagree (**disagreement**)

Noun Subgroup above the eventivity threshold	Tagged items
Non-eventive	105 (13%)
Eventive	288 (60%)
Disagreement	22 (29%)

Testing the eventivity measure (4)

Noun subgroup and morphology	Average eventivity measure	Average eventivity from adjectives	Average eventivity from verbs	Average eventivity from light verbs
non-eventive_MORPH	0.013	0.011	0.002	0
non-eventive_NOTMORPH	0.005	0.004	0.003	0.0009
eventive_MORPH	0.019	0.014	0.004	0.002
eventive_NOTMORPH	0.014	0.009	0.003	0
disagreement_MORPH	0.60	0.53	0.062	0.0004
disagreement_NOTMORPH	0.011	0.007	0.001	0.0009

- the existence of a corresponding verb is a relevant cue
- adjectives represent the most salient cue due to their frequency
- disagreement nouns are always above the eventivity threshold;
instances of dot types not unambiguously detected



A further experiment

- Is the measure reliable with respect to speakers' intuitions for randomly chosen nouns?
 - more informants, with no linguistic background
 - nouns randomly selected from CLR
 - dataset of 80 nouns, but evaluated subset of just 20 nouns
- A lower correlation (.358 vs. .731)
 - each subject has provided intuitions just for a subset of 20 nouns, and no one has performed a comparative ranking
 - nouns selected from lexicographic resources are eventive or non-eventive in a clear way
 - cues need to be weighted ("la palude inizia"). The measure is reliable but should be improved



Conclusion and Future Work

Main contribution of this work: identification of a reliable threshold of eventivity (EVENTIVITY MEASURE) for nouns

- The role of the eventivity measure is two-folded:
- Theoretical approaches: membership of a lexical item to a wider semantic class is a matter of degree
→ *continuum* of eventivity
 - The eventivity measure can be used to discover instances of dot types or coercions
 - Weight of the components of a dot type → disagreement nouns above the threshold
- Practical applications: statistical measure to automatically annotate corpora (KYOTO Project)



Conclusion and Future Work (2)

What is the role of morphology? Should it be employed as a feature of its own in the development of the eventivity measure?

Through syntagmatic co-occurrences is possible to find out common nouns occasionally coerced to an event reading (Simone 2008). Could it be possible to elaborate a data-based typology of the coercion phenomena?

What is the role of adjectives? What is the interplay between the higher frequency of adjectives as type selector and their role in coercions?



Thank you
