Unsupervised extraction of semantic relations using discourse cues

Juliette Conrath

IRIT
ANR Polymnie (ANR-12-CORD-0004)

Monday 24th November, 2014

Introduction

From lexical semantics to discourse : Lexical relations as a tool for natural language processing

Goal

Extraction of lexical relations between events

Applications

- Prediction of discourse relations
- Semantic tasks : causal and temporal information extraction...

Introduction

Prediction of discourse relations

- Discourse: A coherent combination of sentences.
- Discourse relations: A description of how two segments of discourse are logically connected to one another.
- Discourse structure: A representation of discourse as segments logically connected by relations.

Classes of relations

- Temporal
- Causal
- Comparison (mainly contrast)
- Expansion (e.g. elaboration or continuation)

Introduction

- Discourse connectives = markers of discourse relations

Explicit relation : causality

Le candidat a démontré tout son savoir-faire lors de la dernière épreuve. **De ce fait**, le jury a été conquis.

Introduction

- Discourse connectives = markers of discourse relations

Explicit relation : causality

Le candidat a démontré tout son savoir-faire lors de la dernière épreuve. **De ce fait**, le jury a été conquis.

- Challenge of discourse relation prediction :
  Relations are often implicit (52% in the PDTB according to Prasad et al. (2008))

Implicit relation : causality

Le candidat a démontré tout son savoir-faire lors de la dernière épreuve. Le jury a été conquis.

Introduction

Building an annotated corpus for discourse relations

Manual method

Too cumbersome : not much available data

Automatic method

- Explicit relations : markers
- Implicit relations ?
  - Sporleder & Lascarides (2008) and Braud & Denis (2013) : create artificially implicit relations by eliminating markers from explicit relations
  - Redundancy hypothesis : the relation is not only supported by the marker, but also by the rest of the context

Introduction

Proposed approach:

- Infer relation through knowledge of the semantic information of the two main verbs

Verbs as relation cues

Le candidat a **démontré** tout son savoir-faire lors de la dernière épreuve. Le jury a été **conquis**.

- Build a large lexical resource linking pairs of verbs to semantic relations:
  LECSIE: Linked Event Collection for Semantic Information Extraction

Introduction

Overview:

1. Collect frequencies of triples in a large corpus:

Verb 1	Verb 2	Relation	Count

2. Compute significance measures of each triple
   $\Rightarrow$ **Result : Large lexical resource LECSIE**

Verb 1	Verb 2	Relation	Significance measures

3. Evaluate the resource :
   - against intuition (intrinsic evaluation)
   - in a relation prediction task (extrinsic evaluation)

Building *LECSIE*
Resources

Corpus - French

frWaC (Baroni et al., 2009), collected on the Web in the .fr domain - 1.6 billion words.

- Newswire
- Blogs
- etc.

Corpus - English

Gigaword (Graff & Cieri, 2003), collected from four distinct international sources of English newswire - 1.8 billion words.

- Agence France Press English Service
- Associated Press Worldstream English Service
- The New York Times Newswire Service
- The Xinhua News Agency English Service

Building *LECSIE*
Resources

Parser - French

BONSAI (Candito et al.,2010)

- Morpho-syntactic labeling : MElt (Denis & Sagot, 2012)
- Syntactic analysis (dependency) : MaltParser (Nivre et al., 2007)

Parser - English

Stanford Parser (Chen & Manning, 2014) - dependency parsing

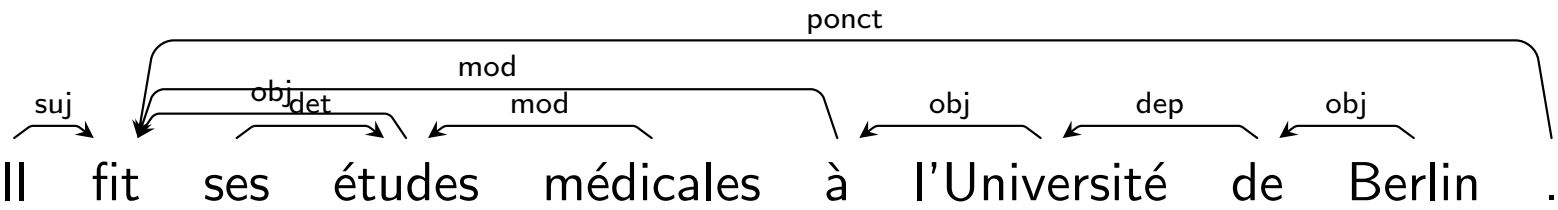

Building *LECSIE*
Resources

Discourse connectives - French

LexConn (Roze et al., 2012): 358 connectives associated with one or several relations - 263 non-ambiguous connectives (only one relation)

Discourse connectives - English

134 connectives collected from the Penn Discourse TreeBank (PDTB) (Prasad et al., 2008) - 104 non-ambiguous connectives

- As a first step, only non-ambiguous connectives are used: direct correspondance between a connective and a relation
- Next step: disambiguation techniques

Example of connective in LexConn

- Lemme : alors_que
- Formes : alors que, alors qu'
- Usage 1 - Relation : contraste
  Par contre, Molière et Shakespeare m'enthousiasmaient, **alors que** *Corneille et même Racine m'ennuyaient profondément.*
- Usage 2 - Relation : background
  Au début de juin, **alors que** *je me trouvais encore à Pau, il m'arriva un jour d'ouvrir ma bible au hasard et d'y lire ces mots...*

Building *LECSIE*
Relation set

Relation	**Connective example**	
	French	**English**
cause	parce que	because
contrast	mais	but
narration	puis	then
continuation	et / encore	and / again
background	alors que	while
temporal localisation	quand	when
detachment	de toute façon	anyway
elaboration	en particulier	in particular
commentary	au fait	by the way
alternation	ou	or
rephrasing	du moins	at least
evidence	effectivement	indeed

Building *LECSIE*
Method

1. Search through corpus for connectives
2. Intra-sentential search

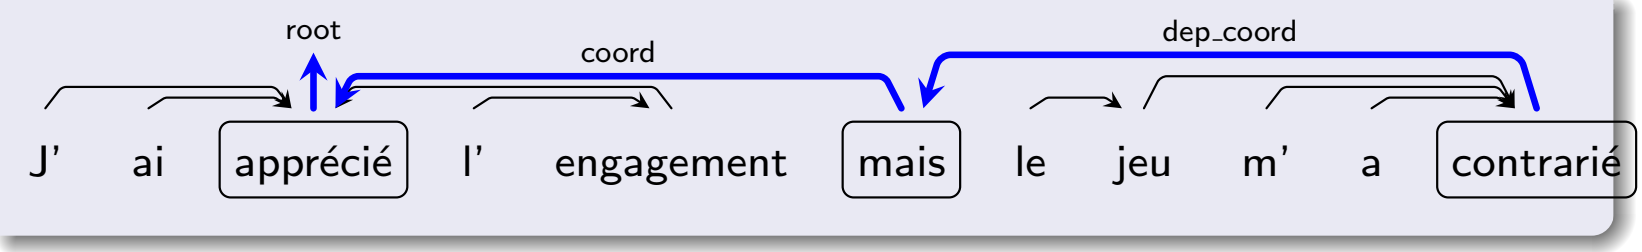

3. Inter-sentential search

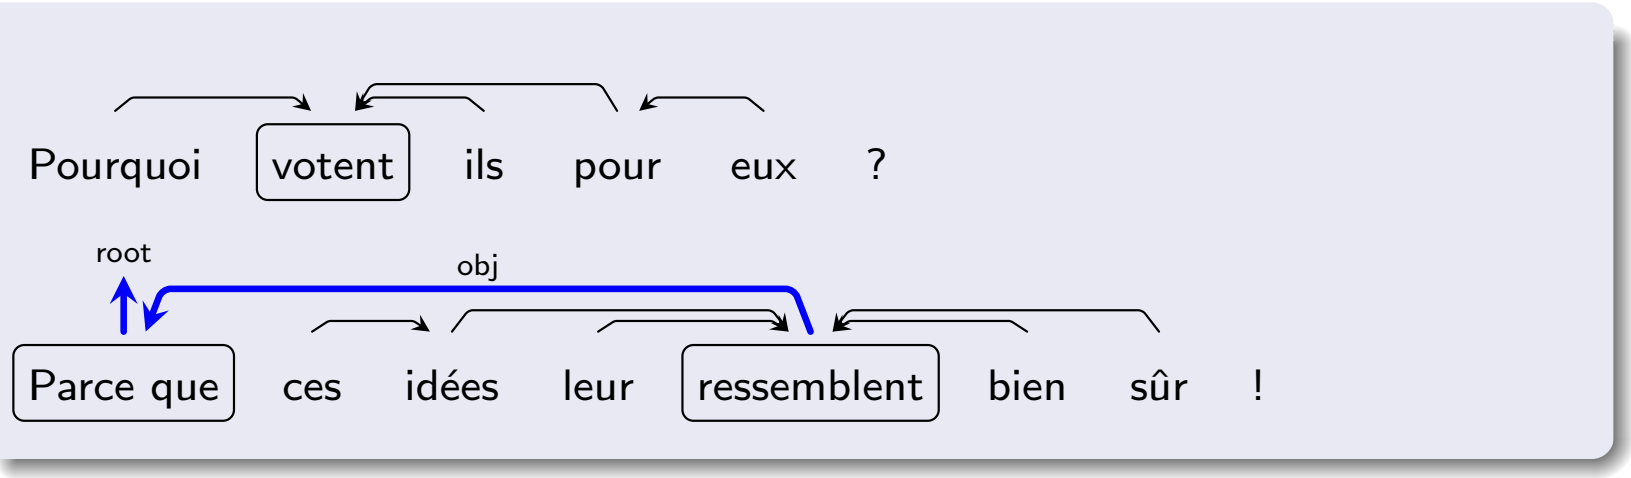

Building *LECSIE*
Method

3. Context
   - support or modal verb, negation, reflexive particle, passive voice
   - idiomatic usage of a preposition using the Dicovalence resource (**tenir de** vs **tenir à**)
   - verbal locutions using the LEFFF resource (**prendre garde**, **faire référence**)
   - filter out most common verbs (14)

Results of this search in frWac

2 million occurrences of triples (verb 1, verb 2, relation)

Results of this search in Gigaword

55 million occurrences of triples (verb 1, verb 2, relation)

Building *LECSIE*

4. Resulting knowledge bases
   - LECSIE-fr: more than 1 million distinct triple types
   - LECSIE-eng: more than 4.6 million distinct triple types

Relation	**Distribution**	
	LECSIE-fr	**LECSIE-eng**
contrast	50,104%	26,236%
cause	33,108%	30,33%
continuation	8,243%	9,549%
narration	6,362%	–
background	1,853%	–
temporal localisation	0.177%	28,932%
detachment	0,149%	–
elaboration	0.002%	0,820%
alternation	0.002%	4,124%

Association measures

- Most common measure: Pointwise Mutual Information (PMI) - adapted to 3 variables

$$PMI = \log\left(\frac{P(V_1, V_2, R)}{P(V_1) \times P(V_2) \times P(R)}\right)$$

Idea: Compare the probability of occurence of two lexical items with a particular semantic relation to the probability of occurrence of each item independantly.

Variants of PMI

- normalized PMI
- discounted PMI (Lin & Pantel, 2002)
- local PMI (Evert, 2005)

Association measures

Other measures

- Measure inspired by Do et al. (2011), defined for causal relations: $\mathbf{U_{do}}$
- **Specificity**, inspired by Mirroshandel et al. (2013)
- Measure defined for our study, to evaluate the contribution of each element to the significance :

$$\mathbf{W}_{combined}(V_1, V_2, R) = \frac{1}{3}(w_{V_1} + w_{V_2} + w_R)$$

where : $w_{V_1} = \frac{P(V_1,V_2,R)}{\max_i(P(V_i,V_2,R))}$, w_{V_2} and w_R similarly defined.

Association measures

Contents of LECSIE: examples

Verb 1	attendre	aider	offenser
Verb 2	déguster	gagner	rire
Relation	continuation	cause	contrast
Count	1047	336	114
PMI	7.1069	4.5154	6.7140
Normalized PMI	0.9363	0.5174	0.6845
Discounted PMI	7.0944	4.5005	6.6320
Specificity	0.7255	0.4739	0.5037
U_Do	7.2614	0.8741	0.3910
W_Combined	1.0	1.0	0.9971

Evaluations

Intrinsic evaluation: validation of *LECSIE* as a semantic resource

- Out of context evaluation
- In context evaluation

Extrinsic evaluation: prediction of implicit discourse relations

- Evaluation of coverage on Annodis
- Including the association measures as features for automatic prediction

Intrinsic evaluation : out of context

- Considered relations : **Cause**, **Contrast**, **Narration**
- **100 verbpairs** selected for each relation, in similar proportions of good and bad scores in LECSIE
- Task: judge whether **a pair can be linked by the considered relation**
- Task accomplished by 3 expert annotators followed by an adjudication of the results

Examples of pairs to be evaluated:

Cause	Contrast	Narration
dérouler / prévoir	travailler / expliquer	partir / ne pas oublier
découvrir / ne pas connaître	essayer / ne pas marcher	apparaître / dresser
lancer / penser	augmenter / baisser	acheter / lire

Intrinsic evaluation : out of context

Results: Inter-annotator agreement

Annotators	Cause	Contrast	Narration
1/2	0.16	0.55	0.43
1/3	0.22	0.57	0.46
2/3	0.13	0.56	0.37
mean kappa	0.17	0.56	0.42

Intrinsic evaluation : out of context

MannWhitney-U tests: Are the measures **statistically discriminative** between contrastive and non-contrastive pairs ?

Measure	p value
specificity	2.5e-11
U_do	2.9e-11
normalized PMI	1.28e-10
discounted PMI	1.96e-10
PMI	1.86e-10
W_combined	4.93e-10
local PMI	4.95e-08
inter-sentential count	0.000904
intra-sentential count	0.0721
raw count	0.116

Intrinsic evaluation : in context

- Considered relations : **Cause**, **Contrast**, **Narration**
- 5 pairs selected for each relation, 40 contexts per pair : **600 contexts** to be annotated
- Task : judge for each context whether **the pair can be linked by the considered relation**
- Task accomplished by 3 expert annotators followed by an adjudication of the results

Examples for the pair promettre / élire in a causality relation

- En 1986, Oscar Arias Sánchez est **élu** en **promettant** d'apporter la paix en Amérique centrale.
- Les discussions **promettent** d'être vives sur le sujet. Cela ne devrait pas empêcher pourtant François Chérèque d'être **réélu** sans problème à la tête de la CFDT.

Intrinsic evaluation : in context
Results

		Human association ratio
Cause		
inviter	souhaiter	12.8%
promettre	élire	25.6%
aimer	trouver	38.5%
bénéficier	créer	51.3%
aider	gagner	53.8%
Contrast		
proposer	refuser	59.0%
augmenter	diminuer	64.1%
tenter	échouer	64.1%
gagner	perdre	71.8%
autoriser	interdire	74.4%
Narration		
parler	réfléchir	42.5%
acheter	essayer	70.0%
atteindre	traverser	77.5%
commencer	finir	80.0%
envoyer	transmettre	82.5%

Intrinsic evaluation : in context

Evaluation of our measures : **Pearson correlation** between the association ratio and each association measure for the 15 pairs considered

	All instances	Implicit instances only
normalized PMI	0.749	0.806
specificity	0.747	0.760
$\mathbf{W_{combined}}$	0.720	0.738
discounted PMI	0.716	0.761
PMI	0.709	0.756
local PMI	0.434	0.553
U_{do}	0.376	0.499
raw frequency	0.170	0.242

10 best triples according to the best measures

Verb 1	**Verb 2**	**Relation**
abandonner	mener	background
ne pas s'arrêter	rouler	narration
donner satisfaction sur	réélire	continuation
emporter	ne pas cesser	summary
emprunter	assurer	cause
ne pas manquer	prolonger	detachment
ratifier	trembler	background
avoir honte	faire pitié	cause
avoir droit	cotiser pour	temp. loc.
ne pas représenter	stéréotyper	temp. loc.

10 best triples according to the best measures

Verb 1	**Verb 2**	**Relation**
authorize	reproduce	alternation
concern	dissatisfy	alternation
damage	destroy	alternation
capsize	drown	temporal
comment	decline	continuation
forecast	record	cause
cover	pursue	alternation
desire	publish	cause
hate	love	alternation
acquit	convict	contrast

Extrinsic evaluation

Available resource - French

ANNODIS: 3355 discourse relations annotated on a French corpus of 660 000 words (Afantenos et al., 2012)

Available resource - English

Penn Discourse TreeBank (PDTB): 40600 relations annotated over the 1 million word Wall Street Journal Corpus (Webber et al., 2008)

Extrinsic evaluation
Coverage

Out of the 3355 pairs of units of discourse linked by a relation in Annodis, only 427 include verbs in both units (many nominal sentences or very frequent verbs)

Elements of coverage

- Proportion of verbpairs linked by a relation in Annodis also found in LECSIE
- Proportion of verbpairs linked by a relation in Annodis also found in LECSIE **with the same relation (triplets)**
- Proportion of verbpairs **implicitely** linked by a relation in Annodis also found in LECSIE
- Proportion of verbpairs **implicitely** linked by a relation in Annodis also found in LECSIE **with the same relation (triplets)**

Extrinsic evaluation
Coverage

Results (in % except first line)

	global	narration	cause	contrast	elab.	cont.	BG	other
Number of Annodis triples	427	73	67	41	96	92	24	16
Annodis pairs ∈ LECSIE	**68.9**	71.2	70.8	78.0	68.3	61.9	74.1	62.5
Annodis triples ∈ LECSIE	**26.5**	34.2	50.0	70.7	0.0	20.6	11.1	0.0
Implicit Annodis pairs	83.4	71.2	79.2	**36.6**	**99.0**	94.8	88.9	100.0
Implicit Annodis pairs ∈ LECSIE	**56.9**	52.1	54.2	**31.7**	**67.3**	58.8	66.7	62.5
Implicit Annodis triples ∈ LECSIE	**17.7**	24.7	40.3	31.7	0.0	19.6	11.1	0.0

Extrincic evaluation
Application

Goal: evaluation of the impact of LECSIE on the task of discourse relation prediction

Discourse relation prediction

- Extraction of features from each pair of discourse units:
  - number of tokens in the unit
  - syntactic category of the head token
  - tense agreement between head verbs
  - etc.
- Maximum entropy model (Muller et al., 2012)

- Include association measures as **additional features** for instances containing verbs included in LECSIE
- Annodis: too few instances for significant results

Perspectives

Evaluation of the english version of LECSIE

- Intrinsic evaluation: in context
- Extrinsic evaluation: Coverage on the PDTB + Relation prediction task

Shared arguments

Coreference analysis: Stanford solver (Lee et al., 2011)

- Example : *X push Y / Y fall*