What a Beautiful Multilingual World: BabelNet 2.0 & Friends!

Roberto Navigli

Dipartimento di Informatica
Sapienza Università di Roma
Linguistic Computing Laboratory
http://lcl.uniroma1.it

TIA 2013, Paris, France

It's all about knowledge!

- Intuitively, we all know what knowledge is…
- …and why we need it

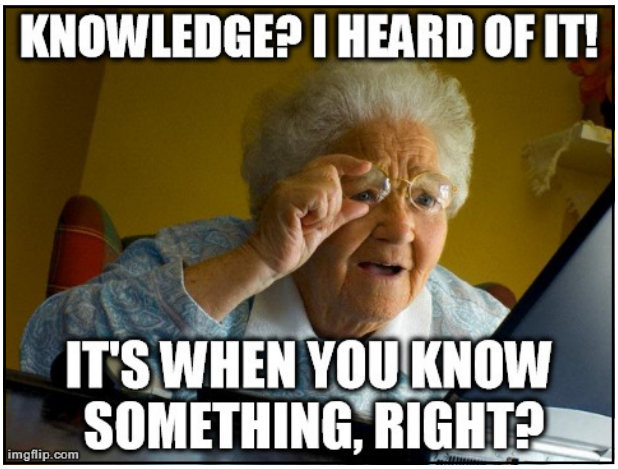

It's all about knowledge!

- But can we expect computers to know?
- Can't computers just use, e.g., statistical techniques?

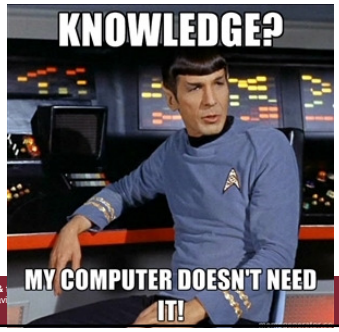

State-of-the-art Machine Translation

EN: I love chocolate, so I bought a bar in the supermarket.

IT: Amo il cioccolato, così ho comprato un bar in un supermercato.

FR: J'aime le chocolat, donc j'ai acheté un bar dans un supermarché.

State-of-the-art Machine Translation

- EN: These are movies in which the music genre, e.g. **rock**, is an important element but not necessarily central to the plot. Examples are Easy Rider (1969), The Graduate (1969), and Saturday Night Fever (1978).

State-of-the-art Machine Translation

- EN: These are movies in which the music genre, e.g. **rock**, is an important element but not necessarily central to the plot. Examples are Easy Rider (1969), The Graduate (1969), and Saturday Night Fever (1978).
- IT: Questi sono i film in cui il genere musicale, ad es **roccia**, è un elemento importante, ma non necessariamente al centro della trama.


State-of-the-art Machine Translation

- EN: Knowledge of the distribution of underground **rock** densities can assist in interpreting subsurface geologic structure and rock type.

Danger here!


State-of-the-art Machine Translation

- EN: Knowledge of the distribution of underground **rock** densities can assist in interpreting subsurface geologic structure and rock type.
- IT: La conoscenza della distribuzione di densità di **rock underground** può aiutare a interpretare in sottosuolo struttura geologica e tipo di roccia.

It's not that the "big data" approach is bad, it's just that mere statistics is not enough

The Knowledge Acquisition Bottleneck

- Knowledge is crucial in language-related research areas
  - Word Sense Disambiguation
  - Named Entity Recognition/Linking
  - Information Extraction
  - **(your favourite area here)**

On a **large scale**, I mean

- However, providing knowledge is **difficult** and **costly**

AKA: The Hamster Wheel

Resources to the rescue!

- Various projects undertaken to make lexical knowledge available in machine readable form
  - WordNet [Fellbaum, 1998]
  - Open Mind Word Expert [Chklovski & Mihalcea, 2002]
  - EuroWordNet [Vossen, 1998]
  - Multilingual Central Repository [Atserias et al. 2004]
  - The WordNetPlus project [Boyd-Graber et al., 2006]
  - OntoNotes [Hovy et al., 2006]
  - Wikipedia (collaborative effort)
  - Wiktionary (collaborative effort)
  - Omega Wiki (collaborative effort)
  - …

Wisdom of the Crowd

But we need an ontology, not just an encyclopedia!

- And, ideally, we need it to be large-scale, wide-coverage

Word Sense Disambiguation in a Nutshell

spring (target word)

"**Spring** water can be found at different altitudes" (context)

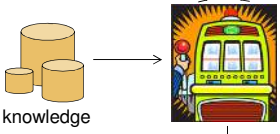

knowledge

WSD system

sense of target word

The Richer, The Better

The Richer, The Better

- Highly-interconnected semantic networks have a great impact on knowledge-based WSD even in a fine-grained setting [Navigli & Lapata, IEEE TPAMI 2010]

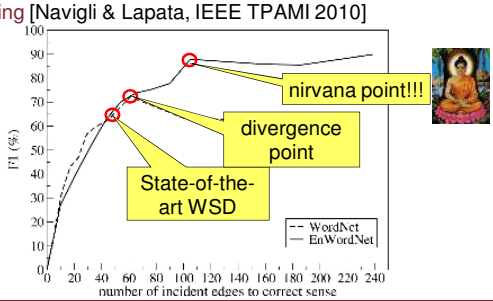

State of the Art "in a nutshell"

- Supervised approaches
  - Require **large amounts** of training data
  - Do **not generalize** across domains and languages

- Knowledge-based approaches have a **higher potential**
  - Lexical knowledge resources **only partly available**

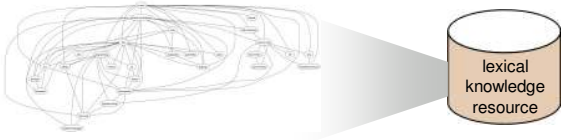

State of the Art "in a nutshell"

- Supervised approaches
  - Require **large amounts** of training data
  - Do **not generalize** across domains and languages

- Knowledge-based approaches have a **higher potential**
  - Lexical knowledge resources **only partly available**
  - Only for **few languages** (e.g. not all 23 EU official languages)
  - **Heterogenous** and with **low coverage**

This is where the ERC (and my project) comes into play

erc MULTIJEDI

Multilingual Joint word sensE DIsambiguation

A 5-year ERC Starting Grant (2011-2016)
on Multilingual Word Sense Disambiguation

Multilingual Joint Word Sense Disambiguation (MultiJEDI)

Key Objective 1: create **knowledge** for **all languages**

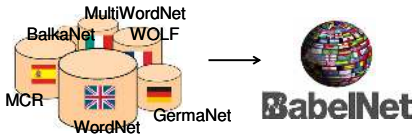

Multilingual Joint Word Sense Disambiguation (MultiJEDI)

Key Objective 2: use **all languages** to disambiguate **one**

The Vision

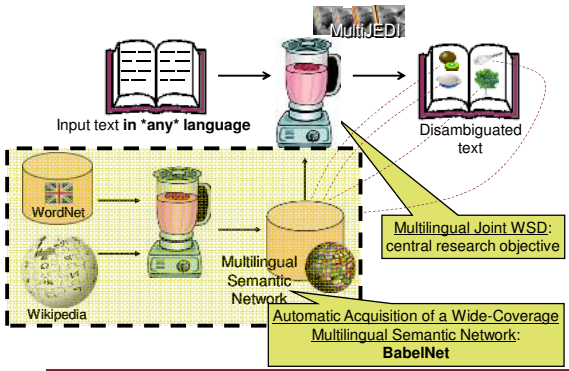

Objective 1: Creating a Multilingual Semantic Network

- Start from two large **complementary** resources:
  - WordNet: full-fledged taxonomy
  - Wikipedia: multilingual and continuously updated

BabelNet [Navigli and Ponzetto, AIJ 2012]

- A wide-coverage multilingual semantic network including both encyclopedic (from Wikipedia) and lexicographic (from WordNet) entries

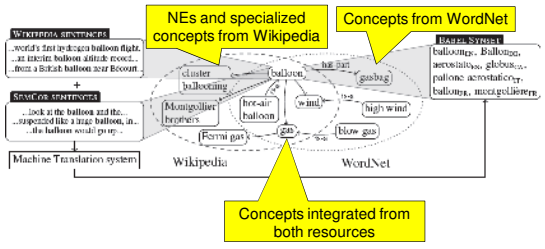

BabelNet integrates the best of both worlds

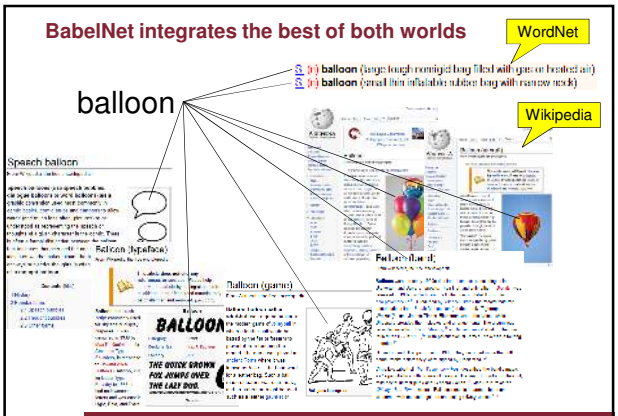

Taming the long tail…

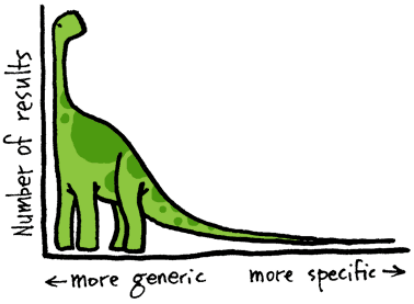

WordNet [Miller et al., 1990; Fellbaum, 1998]

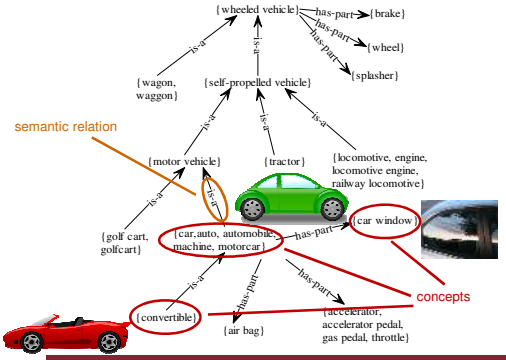

Wikipedia [The Web Community, 2001-today]

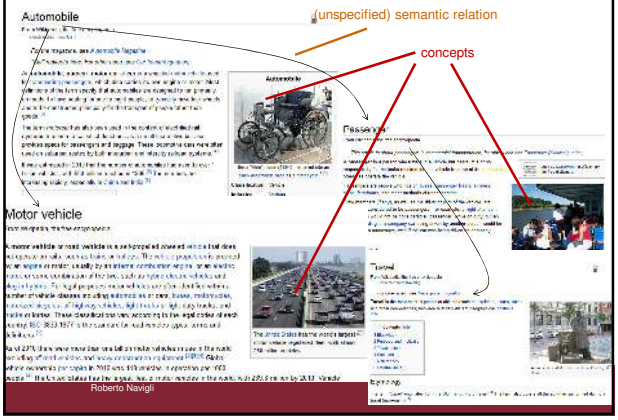

BabelNet: concepts and semantic relations (1)

- Concepts and relations in BabelNet are harvested from **WordNet** and **Wikipedia**:
  - **WordNet:** synsets, lexico-semantic relations → **BabelNet:** concepts, semantic relations
  - **Wikipedia:** pages, hyperlinks → **BabelNet:** concepts, semantic relations

BabelNet: concepts and semantic relations (2)

- We encode knowledge as a labeled directed graph:
  - Each vertex is a Babel synset

balloon$_{EN}$, Ballon$_{DE}$, aerostato$_{ES}$, aerostato$_{IT}$, pallone aerostatico$_{IT}$, mongolfière$_{FR}$

  - Each edge is a semantic relation between synsets:
    - **is-a** (balloon is-a aircraft)
    - **part-of** (gasbag part-of balloon)
    - **instance-of** (Einstein instance-of physicist)
    - …
    - **unspecified/relatedness** (balloon related-to flight)

BabelNet: objectives

1. Provide a unified resource
   - By establishing an automated mapping between Wikipedia pages and WordNet senses
2. Enable multilinguality
   - By collecting the lexicalizations of concepts in different languages using:

     a) Wikipedia interlanguage links
     b) Statistical Machine Translation

An example of mapping

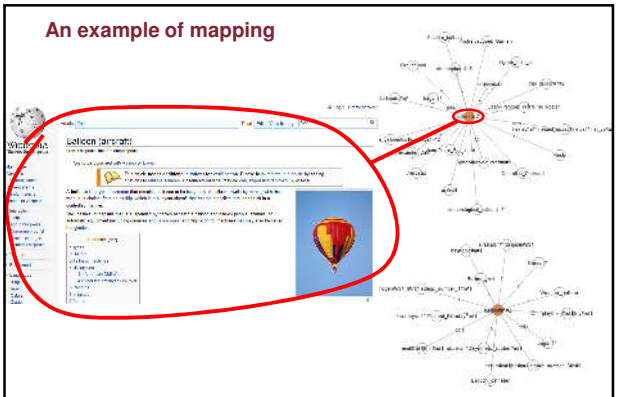

Creation of the Wikipedia disambiguation contexts

ctx(Balloon (aircraft)) = { }

Creation of the Wikipedia disambiguation contexts

ctx(Balloon (aircraft)) = { aircraft }

Creation of the Wikipedia disambiguation contexts

ctx(Balloon (aircraft)) = { aircraft, aerostat, buoyancy, airship, …, gondola }

Creation of the Wikipedia disambiguation contexts

ctx(Balloon (aircraft)) = { aircraft, aerostat, buoyancy, airship, …, gondola, ballooning, hydrogen, aeronautics }

Building BabelNet: Mapping Wikipedia to WordNet

$$\mu(w) = \underset{s \in Senses_{WN}(w)}{\operatorname{argmax}} p(s|w) = \underset{s}{\operatorname{argmax}} \frac{p(s,w)}{p(w)} = \underset{s}{\operatorname{argmax}}\, p(s,w)$$

- Given a Wikipage w and its disambiguation context ctx(w):
  - For each WordNet sense s of w, calculate score(s, w) as follows:

$$score(s,w) = \sum_{cw \in Ctx(w)} \sum_{s' \in Senses_{WN}(cw)} \sum_{p \in paths_{WN}(s,s')} e^{-(length(p)-1)}$$

The Wikipedia page context in the WordNet graph

ctx(Balloon (aircraft)) = { aircraft, aerostat, buoyancy, airship, …, gondola }

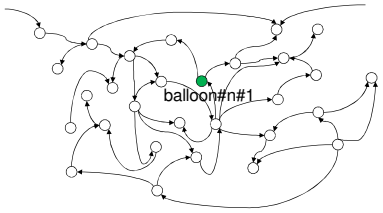

The Wikipedia page context in the WordNet graph

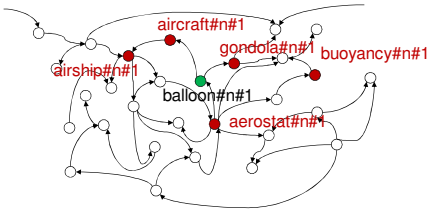

The Wikipedia page context in the WordNet graph

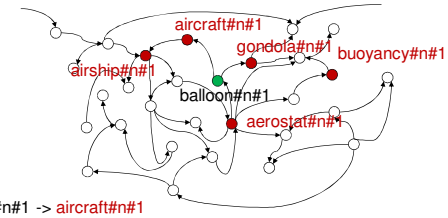

balloon#n#1 -> aircraft#n#1
balloon#n#1 -> aircraft#n#1 -> airship#n#1
balloon#n#1 -> gondola#n#1
balloon#n#1 -> gondola#n#1 -> flight#n#1 -> buoyancy#n#1
balloon#n#1 -> aerostat#n#1

The Wikipedia page context in the WordNet graph

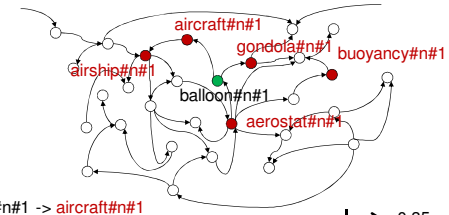

balloon#n#1 -> aircraft#n#1
balloon#n#1 -> aircraft#n#1 -> airship#n#1 ↦ 0.35
balloon#n#1 -> gondola#n#1
balloon#n#1 -> gondola#n#1 -> flight#n#1 -> buoyancy#n#1
balloon#n#1 -> aerostat#n#1

Building BabelNet: Translating Babel synsets

1. Exploiting Wikipedia interlanguage links

Building BabelNet: Translating Babel synsets

2. Filling the lexical translation gaps using a Machine Translation system to translate the English lexicalizations of a concept

- On August 27, 1783 in Paris, Franklin witnessed the world's first hydrogen **{{Balloon (aircraft)|balloon}}** flight.

Google Translate

- Le 27 Août, 1783 à Paris, Franklin vu le premier vol en **ballon** d'hydrogène.

Building BabelNet: Translating Babel synsets

2. Filling the lexical translation gaps using a Machine Translation system to translate the English lexicalizations of a concept

- For each word sense s, we translate:
  - sentences from SemCor (a corpus annotated with WordNet senses) which contain s
  - sentences from Wikipedia linked to the Wikipage of s
- The most frequent translation of s is selected for each target language

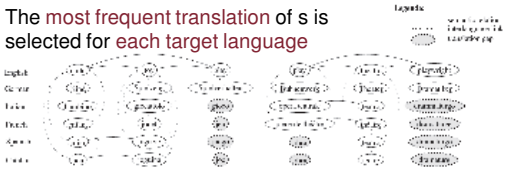

The most frequent translation of a word in a given meaning

left context	term	right context
	wikification	may refer to: the…
geoinformatics services' and '	wikification	of GIS by the masses'
the process may be called	wikification	(as in **...**
which is then called "	wikification	and to the related problem
reason needs copyediting,	wikification	, reduction of POV, work on references
huge amount of cleanup,	wikification	, etc. Version of 12 Nov

The most frequent translation of a word in a given meaning

left context	term	right context
	wikificazione	potrebbe riferirsi a: il…
servizi geoinformatici' e '	wikification	di GIS dalle masse'
il processo chiamato	wikificazione	(come in **...**
che è quindi chiamato	wikificazione	e al problema correlato…
ragione richiede copyediting,	wikification	, riduzione di POV, lavoro su reference
grandi quantità di pulizia,	wikificazione	, ecc. Versione del 12 Novembre

The most frequent translation of a word in a given meaning

left context	term	right context
	wikificazione	potrebbe riferirsi a: il…
servizi geoinformatici' e '	wikification	di GIS dalle masse'
il processo chiamato	wikificazione	(come in **...**
che è quindi chiamato	wikificazione	e al problema correlato…
ragione richiede copyediting,	wikification	, riduzione di POV, lavoro su reference
grandi quantità di pulizia,	wikificazione	, ecc. Versione del 12 Novembre

BabelNet: a multilingual encyclopedic dictionary!

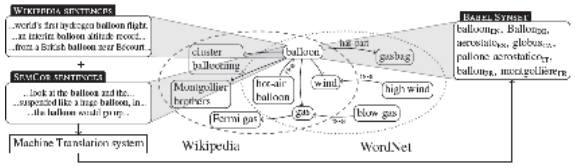

- Available online: http://babelnet.org

For research purposes...

BabelNet 2.0 is online: http://babelnet.org

BabelNet knows terminology

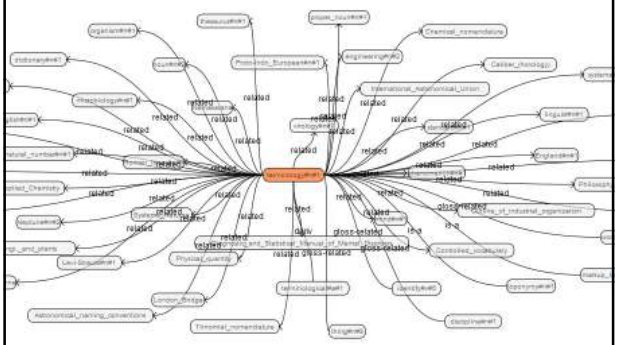

BabelNet knows Paris 13!

BabelNet knows Paris 13!

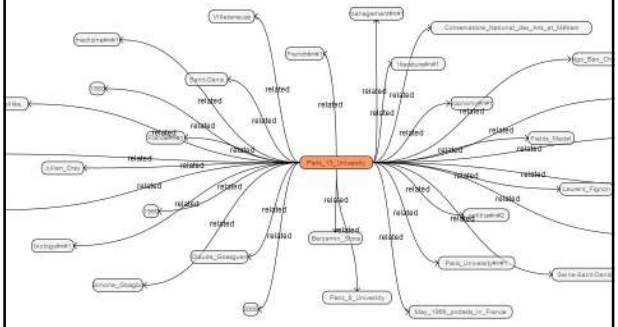

The BabelNet API

```
BabelNet bn = BabelNet.getInstance();
System.out.println("SYNSETS WITH English word: \"bank\": ");
List<BabelSynset> synsets = bn.getSynsets(Language.EN, "bank");
for (BabelSynset synset : synsets)
{
    System.out.print("  =>[  - synset.getId() + "] SOURCE:  - synset.getSource() +
                     "; NN SYNSET:  " + synset.getWordNetOffsets() + ";\n  " +
                     "  NATIV LEMMA:     synset.getNativeLemma()) + ";\n  SENSES (Lemma): [ ");
    for (BabelSense sense : synset.getSenses(Language.EN))
        System.out.print(sense.toString()+" ");
    System.out.println("]\n  -----");
    Map<IRelation,  List<BabelSynset> relatedSynsets = synset.getRelatedMap();
    for (IRelation relationType : relatedSynsets.keySet())
    {
        List<BabelSynset> relationSynsets = relatedSynsets.get(relationType);
        for (BabelSynset relationSynset : relationSynsets)
        {
            System.out.println("    - RELATION: " + relationType.getSynset() +
                               " + relationSynset.getId() +
                               " + relationSynset.toString(Language.EN));
        }
    }
    System.out.println("");
}
```

Retrieve all synsets with the English lemma "bank"

Print information about each synset

Get the (relation, synsets) map of the synset neighbours

Get the synsets related by a given relation type

Print the information of each related synset

BabelNet goes at a faster pace than I can cope with

Anatomy of BabelNet 2.0

Previous version had 6!

- **50** languages covered (including Latin!)
- List at http://babelnet.org/stats.jsp

Anatomy of BabelNet 2.0

- **50** languages covered (including Latin!)
- **9.3M** Babel synsets (concepts and named entities)
- **50M** word senses
- **262M** semantic relations (28 edges per synset on avg.)
- **7.7M** synset-associated images
- **18M** textual definitions

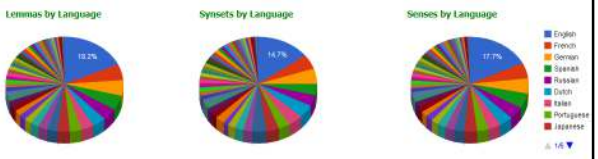

Anatomy of BabelNet 2.0

- **50** languages covered (including Latin!)
- Integrates:
  - WordNet 3.0
  - Wikipedia (2012 dump)
  - OmegaWiki: a collaborative multilingual dictionary
  - Open Multilingual WordNet [Bond and Foster, 2013]
- Translations for all open-class parts of speech

WordNet+Open Multilingual WordNet+Wikipedia+…

+OmegaWiki+automatic translations…

+textual definitions

+Wikipedia categories

+images

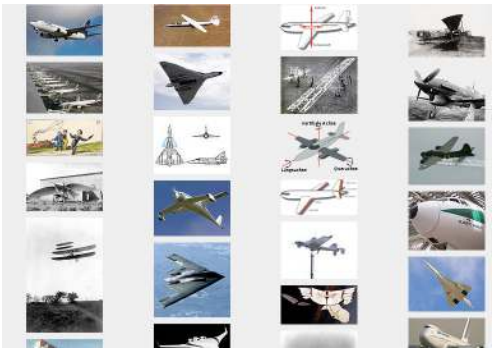

Evaluations: I have to go fast here!

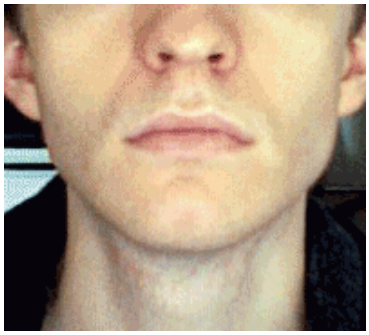

WordNet-Wikipedia mapping accuracy

- Overall quality of the mapping: ~**84%**
  - On a random sample of 1k Wikipages
  - Note: this concerns only those 50k synsets in the intersection
- Quality of the mapping of frequent words: ~**91%**

Evaluation of BabelNet against gold standard resources

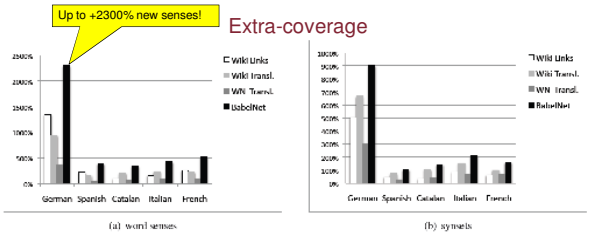

Coarse-grained Word Sense Disambiguation with BabelNet

Resource	Algorithm	Nouns only $P/R/F_1$	All words $P/R/F_1$
WordNet	Degree	80.1	79.7
	PLength	80.2	79.8
	SProbability	79.8	79.3
	PageRank	79.9	79.4
BabelNet	Degree	**84.7**	**82.3**
	PLength	**85.4**	**82.7**
	SProbability	**84.6**	**82.1**
	PageRank	82.3	80.1
	SUSSX-FR	81.1	77.0
	TreeMatch	N/A	73.6
	NUS-PT	82.3	**82.5**
	SSI	**84.1**	**83.2**
	MFS BL	77.4	78.9
	Random BL	63.5	62.7

State of the art results!

Current state of the art

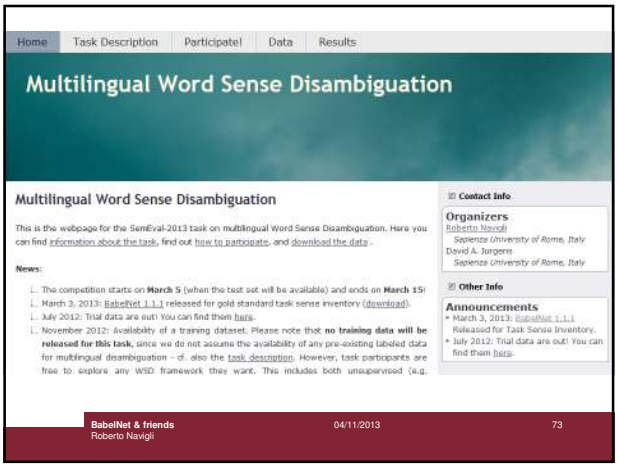

Key fact!

Annotating with BabelNet: all in one!

- Annotating with **BabelNet** implies annotating with WordNet and Wikipedia
- (now also OmegaWiki and Open Multilingual WordNet!)

Dataset for Multilingual Word Sense Disambiguation

Language	Instances	Single-words	Multiword expressions	Named Entities
		BabelNet		
English	1931	1604	127	200
French	1656	1389	89	176
German	1467	1267	21	176
Italian	1706	1454	211	41
Spanish	1481	1103	129	249
		Wikipedia		
English	1242	945	102	195
French	1039	790	72	175
German	1156	957	21	176
Italian	1977	869	85	41
Spanish	1103	758	107	248
		WordNet		
English	1644	1502	85	57

We are not alone in the (resource) universe!

- DBPedia [Bizer et al. 2009] - a resource obtained from structured information in Wikipedia
  - «Describes 3.77M things»
  - Core of the Linked Open Data Cloud
- YAGO [Suchanek et al. 2007]
  - «Contains 10M entities and 120M facts about these entities»
  - Links Wikipedia categories to WordNet synsets
- MENTA [de Melo and Weikum, 2010]
  - A «multilingual taxonomy with 5.4M entities»
- WikiNet [Nastase and Strube, 2013]
  - Semantic network connecting Wikipedia entities
  - «3M concepts and 38+M relations»
- Freebase (http://freebase.com): collaborative effort
  - Structured data; started from Wikipedia, MusicBrainz, ChefMoz, etc.

So where is the novelty?

- We provide a **unified**, **integrated** inventory and network for **both** word senses and named entities

Now in the Linked Open Data cloud…

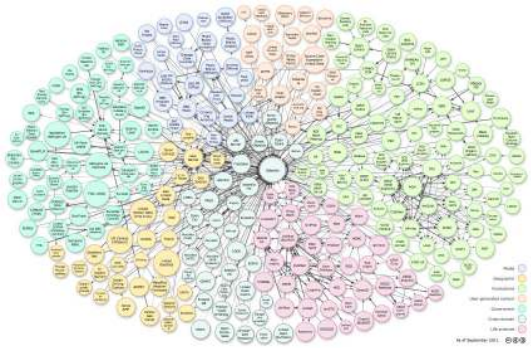

Actually, in the *Linguistic* Linked Open Data cloud…

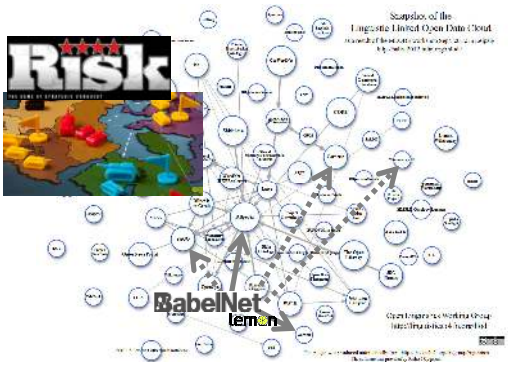

SPred: Semantic Predicates on a large scale [Flati & Navigli, ACL 2013]

- Choose a set of semantic classes
- Given a lexical predicate

$$w_1\ w_2\ \ldots\ w_i\ ^*\ w_{i+1}\ \ldots\ w_n$$

  identify the semantic class distribution for *
- Classify new filling arguments

The idea in a nutshell: start from text (1/6)

cup of *

…pounded millet mixed with two	cups of	butter, every day…
…senior citizens would linger over	cups of	coffee and exchange news…
...began over a	cup of	yogurt and a broken printer that …
…in the quarterfinal	cup of	Yugoslavia as a renowned team…
…which is equivalent to 3	cups of	regular coffee, he will be able to…
He drank several more	cups of	wine and started behaving wildly…
…or fruit, as well as small	cups of	kosher wine or other beverages…
…two Euroleague titles, two	Cups of	Italy, one Korać Cup and one…

The idea in a nutshell: focus on lexical predicates (2/6)

cup of *

…pounded millet mixed with two	**cups of**	butter, every day…
…senior citizens would linger over	**cups of**	coffee and exchange news…
...began over a	**cup of**	yogurt and a broken printer that …
…in the quarterfinal	**cup of**	Yugoslavia as a renowned team…
…which is equivalent to 3	**cups of**	regular coffee, he will be able to…
He drank several more	**cups of**	wine and started behaving wildly…
…or fruit, as well as small	**cups of**	kosher wine or other beverages…
…two Euroleague titles, two	**Cups of**	Italy, one Korać Cup and one…

The idea in a nutshell: focus on filling arguments (3/6)

cup of *

…pounded millet mixed with two	**cups of**	**butter**, every day…
…senior citizens would linger over	**cups of**	**coffee** and exchange news…
...began over a	**cup of**	**yogurt** and a broken printer that …
…in the quarterfinal	**cup of**	**Yugoslavia** as a renowned team…
…which is equivalent to 3	**cups of**	regular **coffee**, he will be able to…
He drank several more	**cups of**	**wine** and started behaving wildly…
…or fruit, as well as small	**cups of**	**kosher wine** or other beverages…
…two Euroleague titles, two	**Cups of**	**Italy**, one Korać Cup and one…

The idea in a nutshell: gather all the filling arguments (4/6)

cup of *

The idea in a nutshell: group similar arguments (5/6)

cup of *

The idea in a nutshell: associate semantic labels (6/6)

cup of *

cup of *

Next feature in BabelNet!




$wine_n^1$ $coffee_n^1$ $herb_n^2$ $water_n^1$ ...

Wine	Coffee	Earl Grey tea	Water
Sack	Turkish coffee	Green tea	Seawater
White wine	Drip coffee	Indian tea	...
Red wine	Espresso	Black tea	
Claret	Cappuccino	Tea	
Kosher wine	Caffè latte	...	
Madeira wine	Decaffeinated coffee		
Wine in China	...		
...			

Classes sorted by relevance!

Another Plan for Knowledge Acquisition: Ontology Learning

- Knowledge acquisition without relying on hand-crafted knowledge
- Deals with technical domains
- Current approaches have limits:
  - Kozareva & Hovy [2010]: lexico-syntactic patterns + pruning based on the longest path
  - Yang & Callan [2009]: based on incremental clusters of terms
    - 82% F1 for small-scale WordNet is-a sub-hierarchies (39 terms on average)
    - 61% F1 on part-of sub-hierarchies
  - Snow et al. [2006]: hyponym acquisition based on a probabilistic model
    - 58% P, 21% R
- These approaches have **not** been shown to be able to extract large specialized domain ontologies

OntoLearn Reloaded [Navigli, Faralli & Velardi, IJCAI 2011; Computational Linguistics 2013]

Unlike other approaches, we learn both concepts and relations entirely from scratch for **any** domain of interest

OntoLearn Reloaded: Workflow

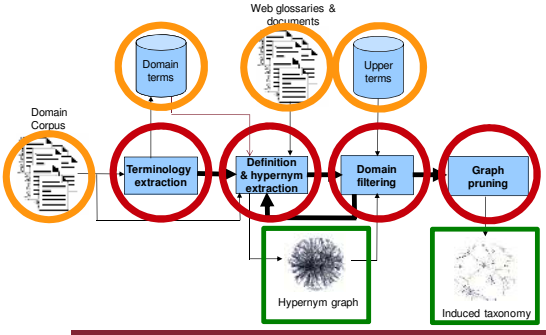

OntoLearn Reloaded: Workflow

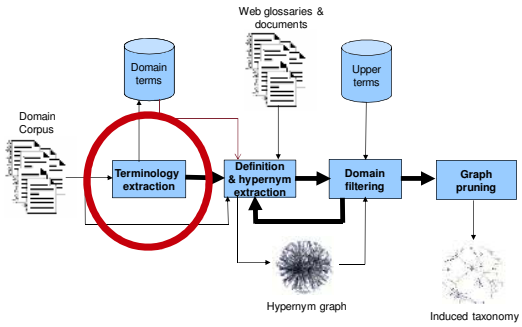

Terminology Extraction

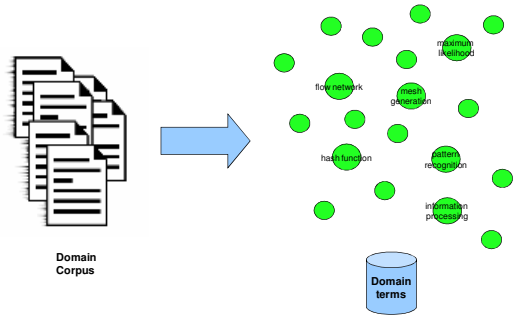

OntoLearn Reloaded: Workflow

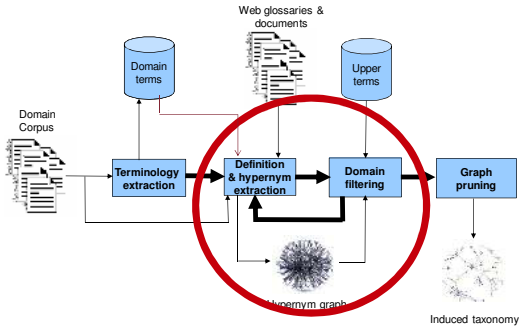

Definition & Hypernym Extraction + Domain Filtering

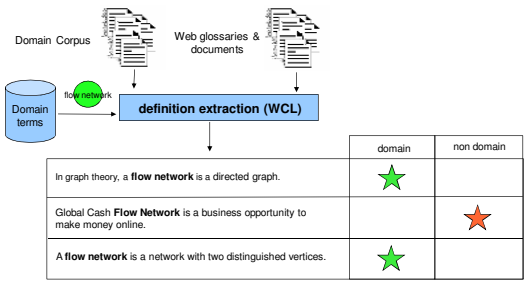

Definition & Hypernym Extraction + Domain Filtering

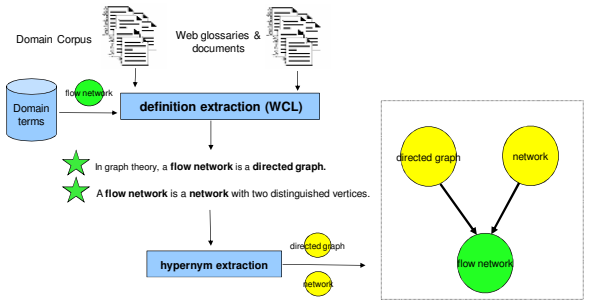

Definition & Hypernym Extraction + Domain Filtering

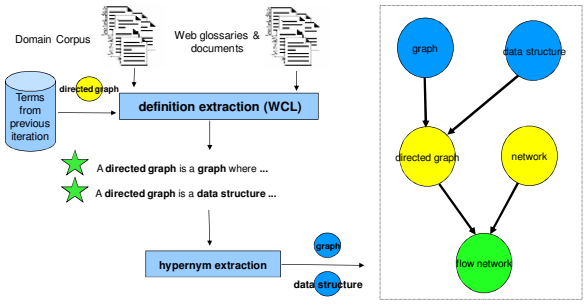

Hypernym Extraction Algorithm: Word Class Lattices

Based on Word-Class Lattices (WCLs), i.e. lattice-based definition models learned by means of a greedy definition alignment algorithm [Navigli & Velardi, ACL 2010]

- Determine whether a sentence is definitional
- If so, return the hypernym(s) of the defined term

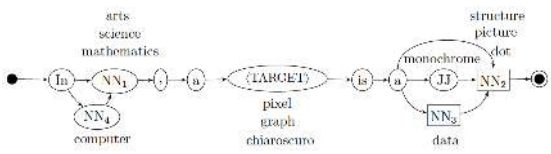

The iterative growth of the hypernym graph

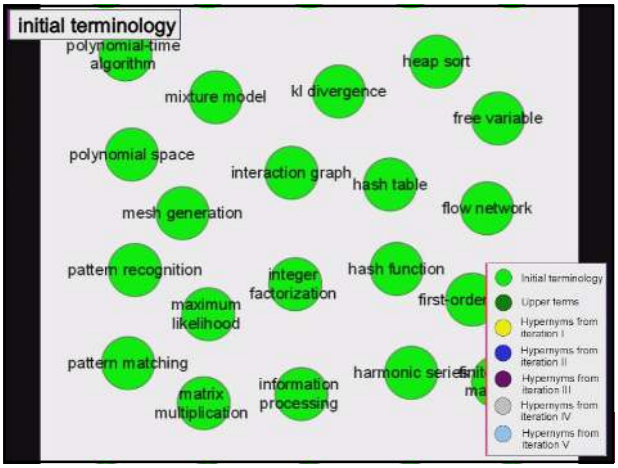
initial terminology
polynomial-time algorithm
heap sort
mixture model
kl divergence
free variable
polynomial space
interaction graph
hash table
mesh generation
flow network
pattern recognition
integer factorization
hash function
maximum likelihood
pattern matching
information processing
matrix multiplication
Initial terminology
Upper terms
Hypernyms from iteration I
Hypernyms from iteration II
Hypernyms from iteration III
Hypernyms from iteration IV
Hypernyms from iteration V

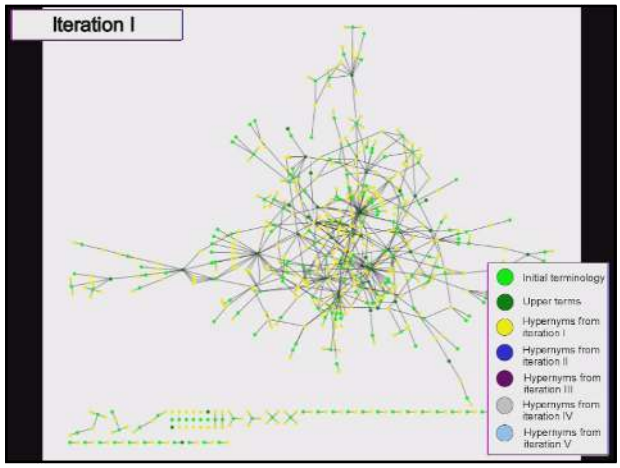
Iteration I
Initial terminology
Upper terms
Hypernyms from iteration I
Hypernyms from iteration II
Hypernyms from iteration III
Hypernyms from iteration IV
Hypernyms from iteration V

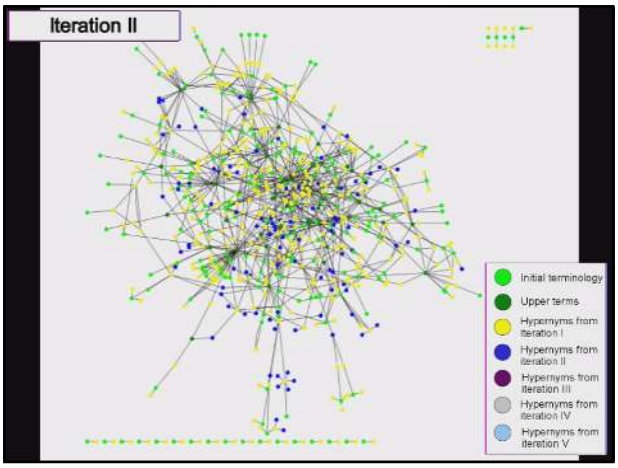
Iteration II
Initial terminology
Upper terms
Hypernyms from iteration I
Hypernyms from iteration II
Hypernyms from iteration III
Hypernyms from iteration IV
Hypernyms from iteration V

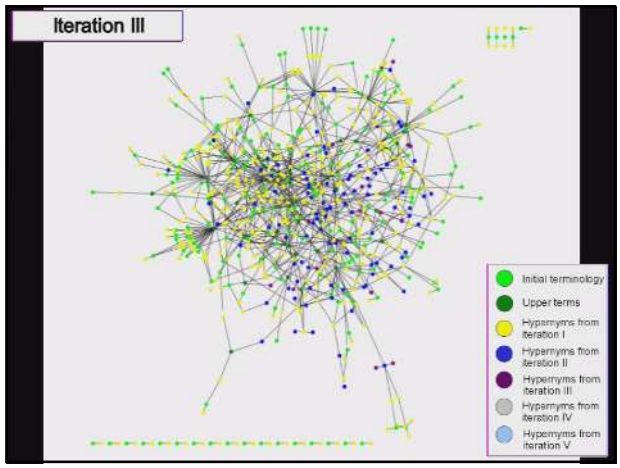
Iteration III
Initial terminology
Upper terms
Hypernyms from iteration I
Hypernyms from iteration II
Hypernyms from iteration III
Hypernyms from iteration IV
Hypernyms from iteration V

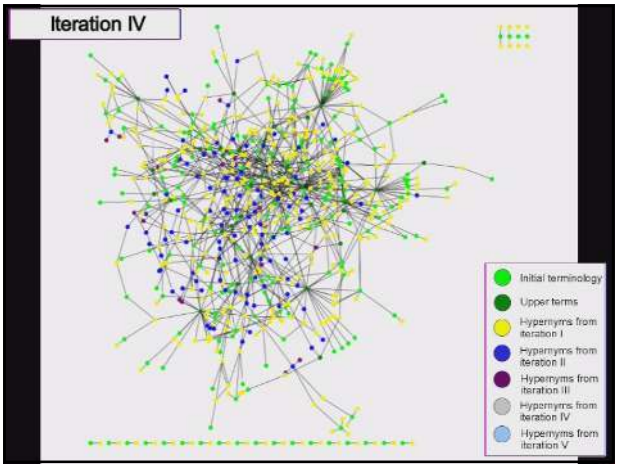
Iteration IV
Initial terminology
Upper terms
Hypernyms from iteration I
Hypernyms from iteration II
Hypernyms from iteration III
Hypernyms from iteration IV
Hypernyms from iteration V

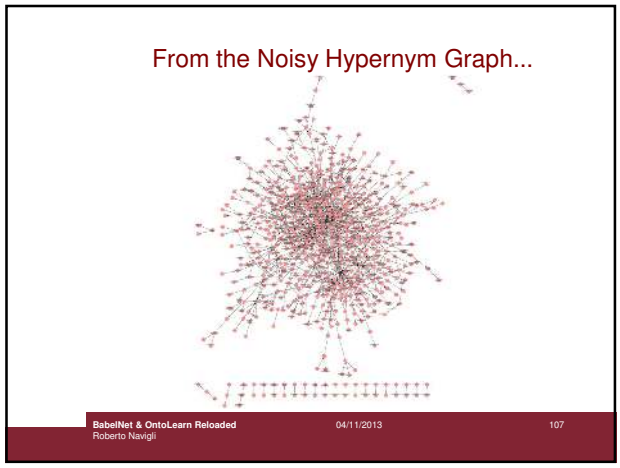
From the Noisy Hypernym Graph...
BabelNet & OntoLearn Reloaded
Roberto Navigli
04/11/2013
107

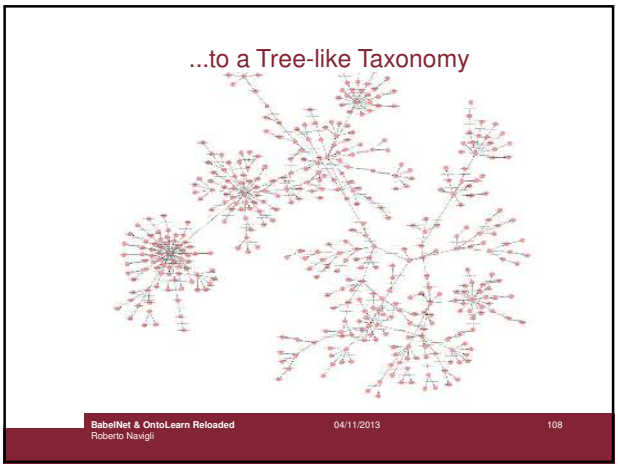
...to a Tree-like Taxonomy
BabelNet & OntoLearn Reloaded
Roberto Navigli
04/11/2013
108

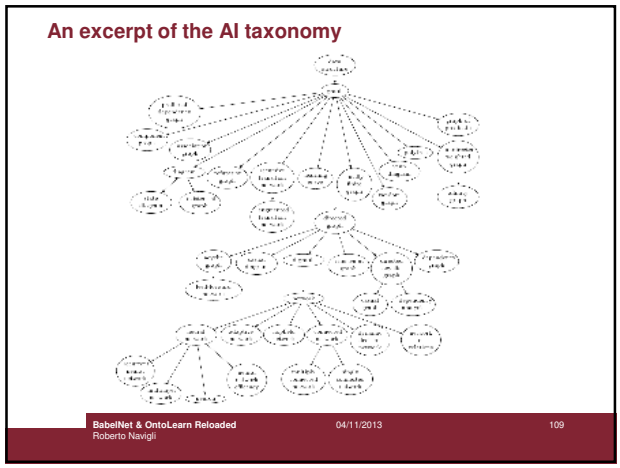
An excerpt of the AI taxonomy
BabelNet & OntoLearn Reloaded
Roberto Navigli
04/11/2013
109

Semantically-Enhanced Open Information Extraction [Moro and Navigli, IJCAI 2013]

Semantic Network	Semantic Relations	Open set of relations
ReVerb	✗	✓
YAGO2	✓	✗
WikiNet	✓	✗
WiSeNet	✓	✓

A. Fader, S. Soderland and O. Etzioni. *Identifying Relations for Open Information Extraction*. In Proc. of EMNLP, 2011.
J. Hoffart et al. *YAGO2: A spatially and temporally enhanced knowledge base from Wikipedia*. Journal of Artificial Intelligence, 2012.
V. Nastase and M. Strube. *Transforming Wikipedia into a large scale multilingual concept network*. Journal of Artificial Intelligence, 2012.
A. Moro and R. Navigli. *WiSeNet: building a wikipedia-based semantic network with ontologized relations*. In Proc. of CIKM, 2012.

Semantically-Enhanced Open Information Extraction [Moro and Navigli, IJCAI 2013]

- Moving to semantically enhanced OIE:

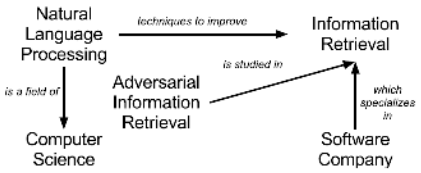

Semantically-Enhanced Open Information Extraction [Moro and Navigli, IJCAI 2013]

- Moving to semantically enhanced OIE:

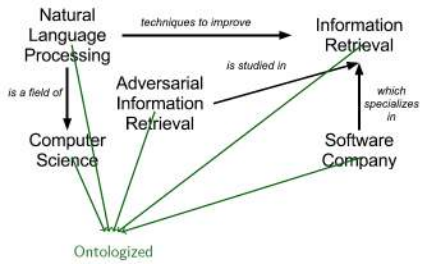

Semantically-Enhanced Open Information Extraction [Moro and Navigli, IJCAI 2013]

- Moving to semantically enhanced OIE:

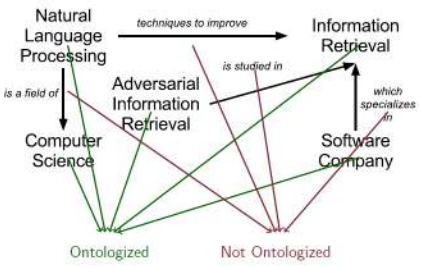

Ontologizing the relation strings

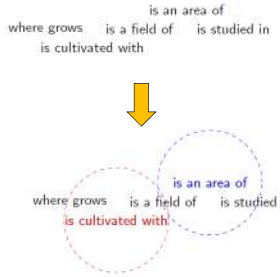

Relation Synset: {where grows, is cultivated with, is a field of}
Relation Synset: {is studied in, is an area of, is a field of}

Improving traditional IE with BabelNet [Moro et al., ISWC 2013]

Google Focused Research Awards

- Using rich semantics increases precision while keeping recall high

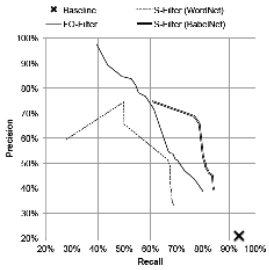

Thanks or…

merci (grazie)

Google Focused Research Awards


SAPIENZA
UNIVERSITÀ DI ROMA
Roberto Navigli
Linguistic Computing Laboratory
http://lcl.uniroma1.it
Linguistic
Computing
Laboratory