




Entity Linking meets Word Sense Disambiguation: a Unified Approach

Andrea Moro, Alessandro Raganato and Roberto Navigli

Transactions of the Association for Computational Linguistics (TACL), 2:231-244, 2014



Word Sense Disambiguation and Entity Linking together!



The authors gratefully acknowledge the support of the ERC Starting Grant MultijEDI No. 259234.

Conference: ACL 2014

The 52nd Annual Meeting of the Association for Computational Linguistics Baltimore, USA, June 22-26 2014

Abstract

Entity Linking (EL) and **Word Sense Disambiguation (WSD)** both address the lexical ambiguity of language. But while the two tasks are pretty similar, they differ in a fundamental aspect: in **EL** the textual mention can be linked to an entity which may or may not contain the exact mention, while in **WSD** there is a perfect match between the word form (better, its lemma) and a suitable sense.

We present a **unified graph-based approach** to EL and WSD based on a **loose identification of candidate meanings** coupled with a **densest subgraph** heuristic which selects high-coherence semantic interpretations.

The Best of Two Worlds

Our main goal is to **bring together** the two worlds of WSD and EL:

1. Keep the **set of candidate meanings** for a given mention as **open as possible**
2. Provide an effective method for **handling this high level of ambiguity**.

A **key assumption** of our work is that the **lexicographic knowledge** used in WSD is also useful to tackle the EL task, and vice versa the **encyclopedic knowledge** utilized in EL helps disambiguate nominal mentions in a WSD setting.

A joint task

Our task is to **disambiguate and link all nominal and named entity mentions** occurring within a text. The linking task is performed by associating each mention with the most suitable entry of a given knowledge base.

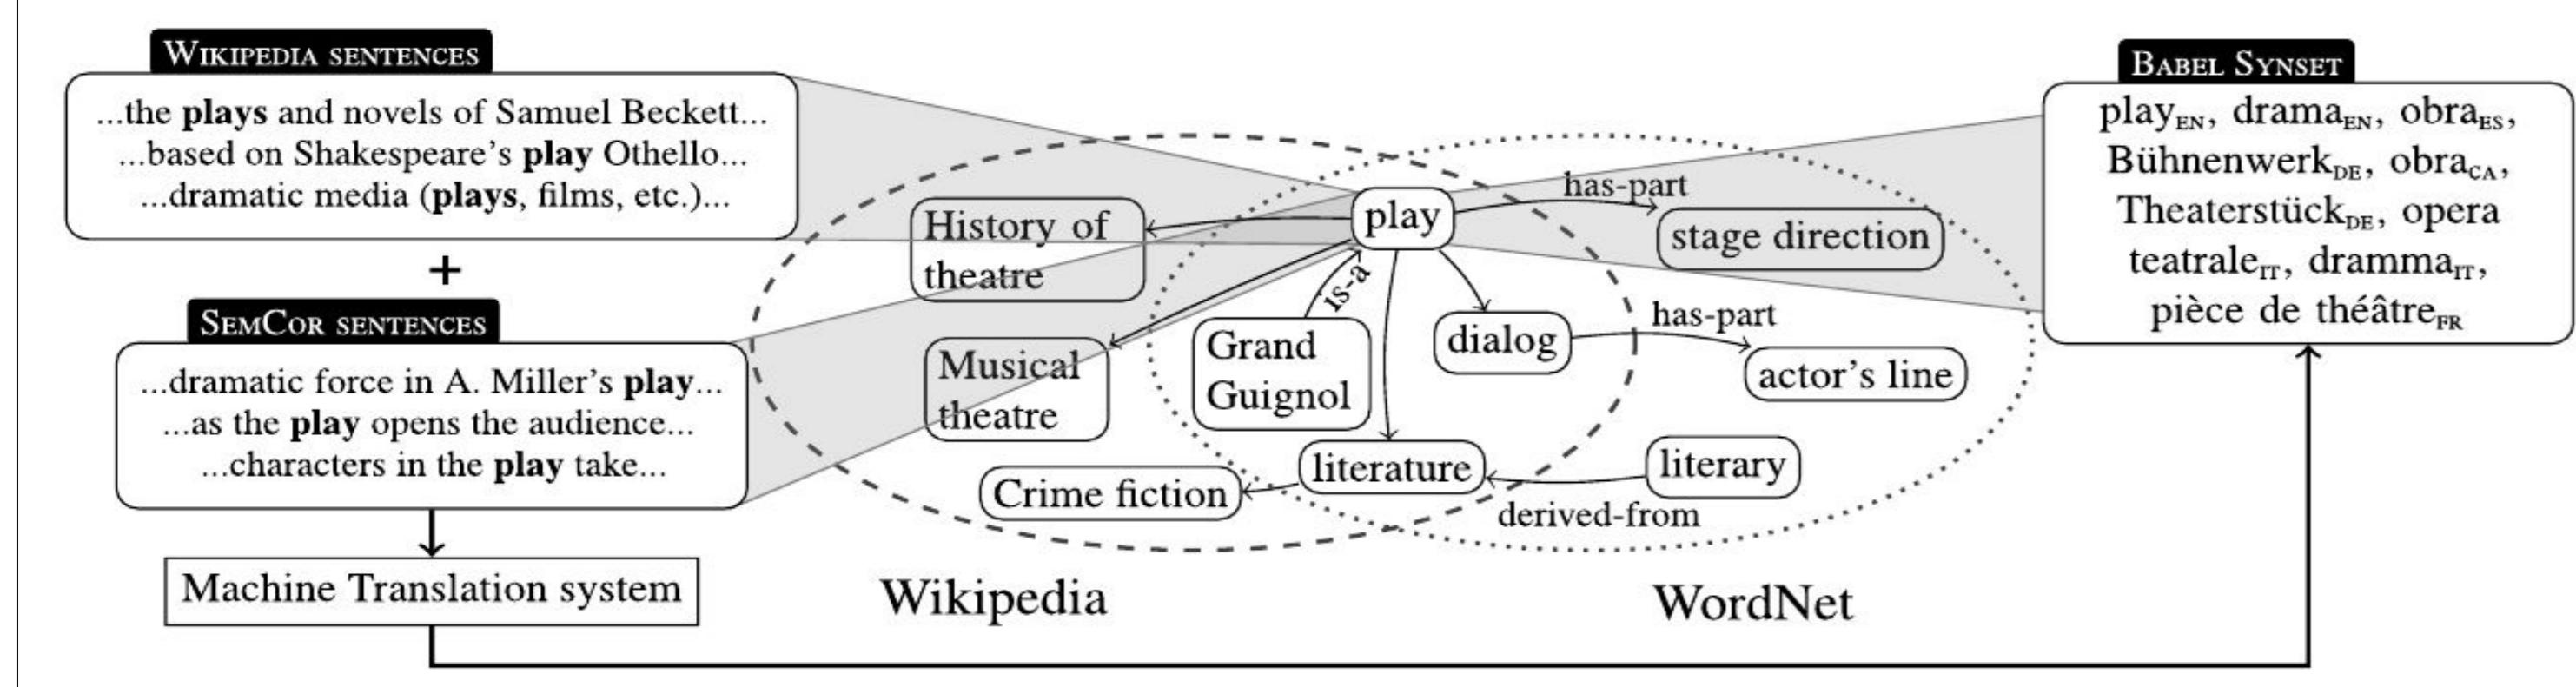
Our definition does not enforce any constraints in terms of what to link, i.e., unlike Wikification and WSD, we can link **overlapping fragments** of text.

Example. Given the text fragment: "Major League Soccer" we identify and disambiguate several different entity and concept mentions:

Major League Soccer major league league soccer

BabelNet (<http://babelnet.org>)

BabelNet is a **multilingual knowledge base** which consists of roughly nine million concepts and named entities together with their lexicalizations in **50 different languages** and 250 million semantic relations. At the core of this resource lies the **integration of encyclopedic**, i.e., from Wikipedia and Wikidata, and **lexicographic knowledge**, i.e., from WordNet, Open Multilingual WordNet Wiktionary and OmegaWiki within a unified, multilingual structured network.



Our Approach

1. Given a lexicalized semantic network, we associate with each vertex, i.e., either concept or named entity, a **semantic signature**, that is, a set of related vertices. This is a preliminary step which needs to be performed only once, independently of the input text.
2. Given a text, we **extract all the linkable fragments** from the text and, for each of them, list the possible meanings according to the semantic network.
3. We create a **graph-based semantic interpretation** of the whole text by linking the candidate meanings of the extracted fragments using the previously-computed semantic signatures. We then extract a **dense subgraph** of this representation and **select the best candidate meaning** for each fragment.

Semantic Signatures

For each Babel synset, we acquire **strongly related synsets** in 3 steps:

1. We first weight the edges using triangles:

$$weight(v, v') := |\{(v, v', v'') : (v, v'), (v', v''), (v'', v) \in E\}| + 1$$

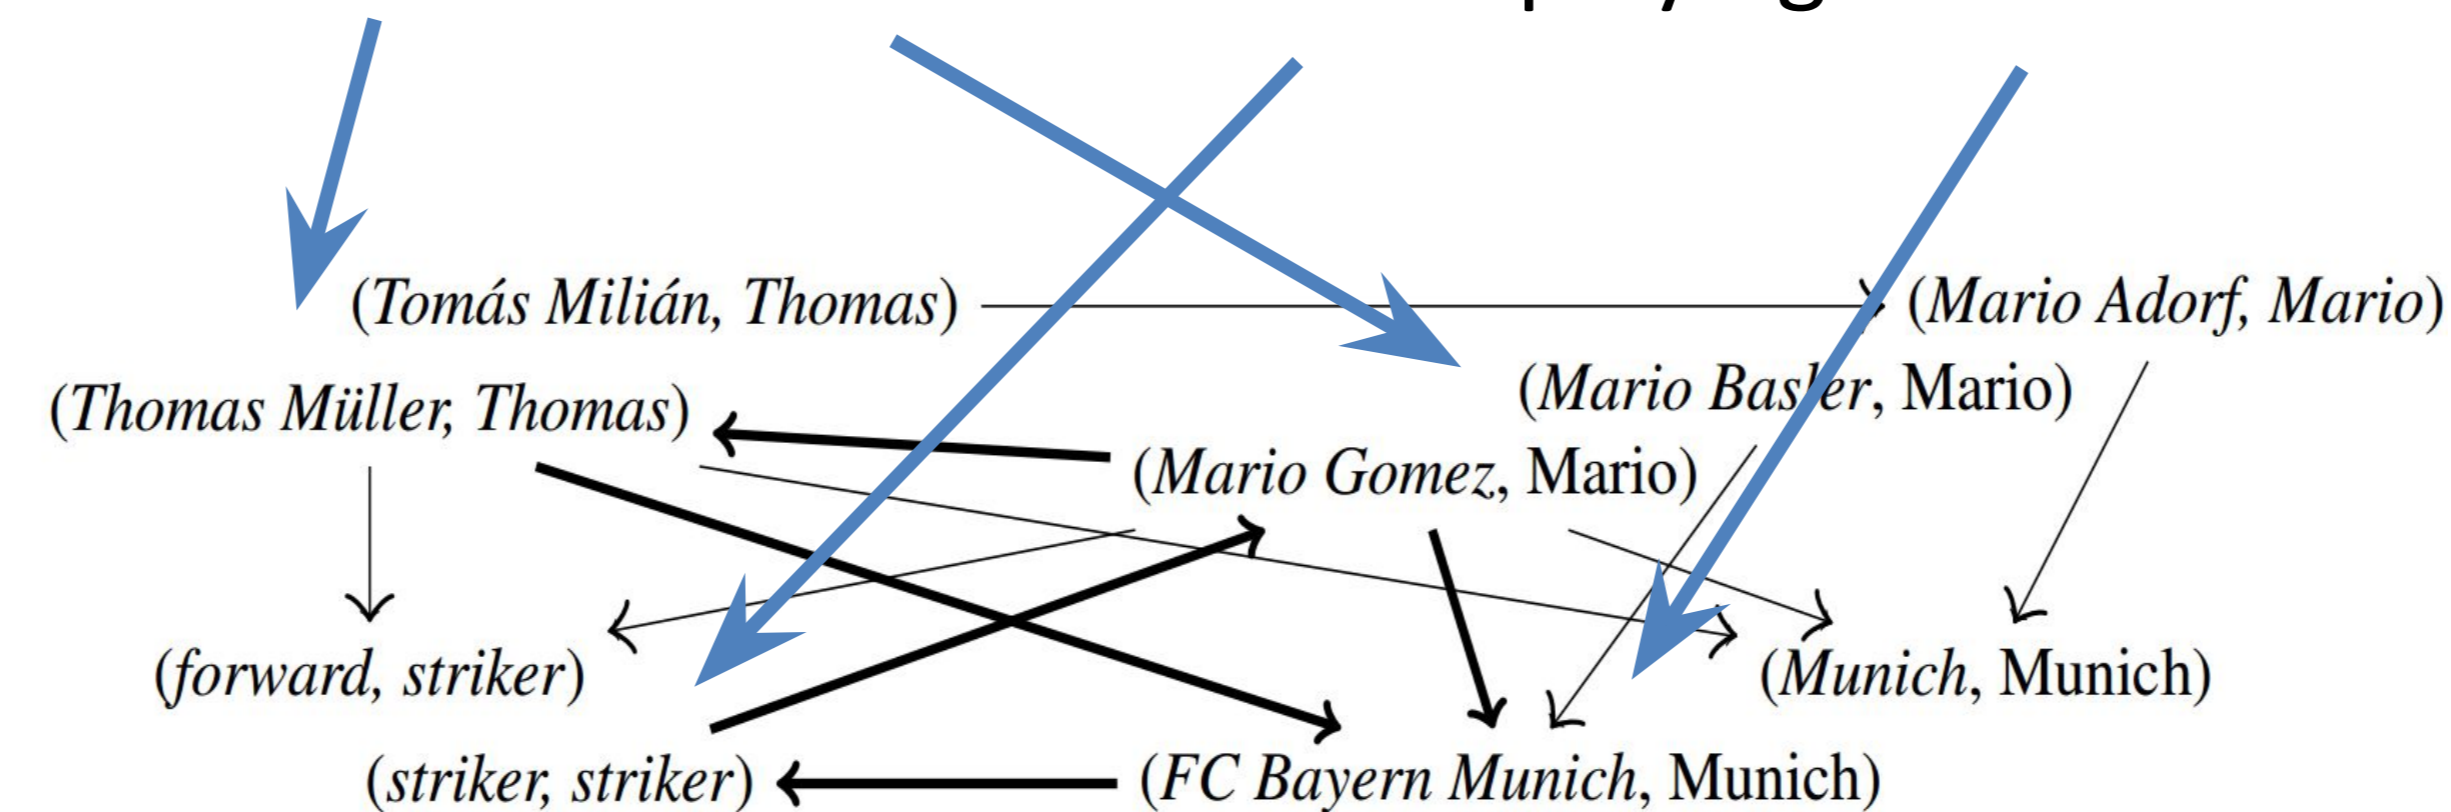
2. We then run a random walk with restart: $P(v'|v) = \frac{weight(v, v')}{\sum_{v'' \in V} weight(v, v')}$

3. The semantic signature of a vertex v is the set of vertices visited at least μ times during the random walk.

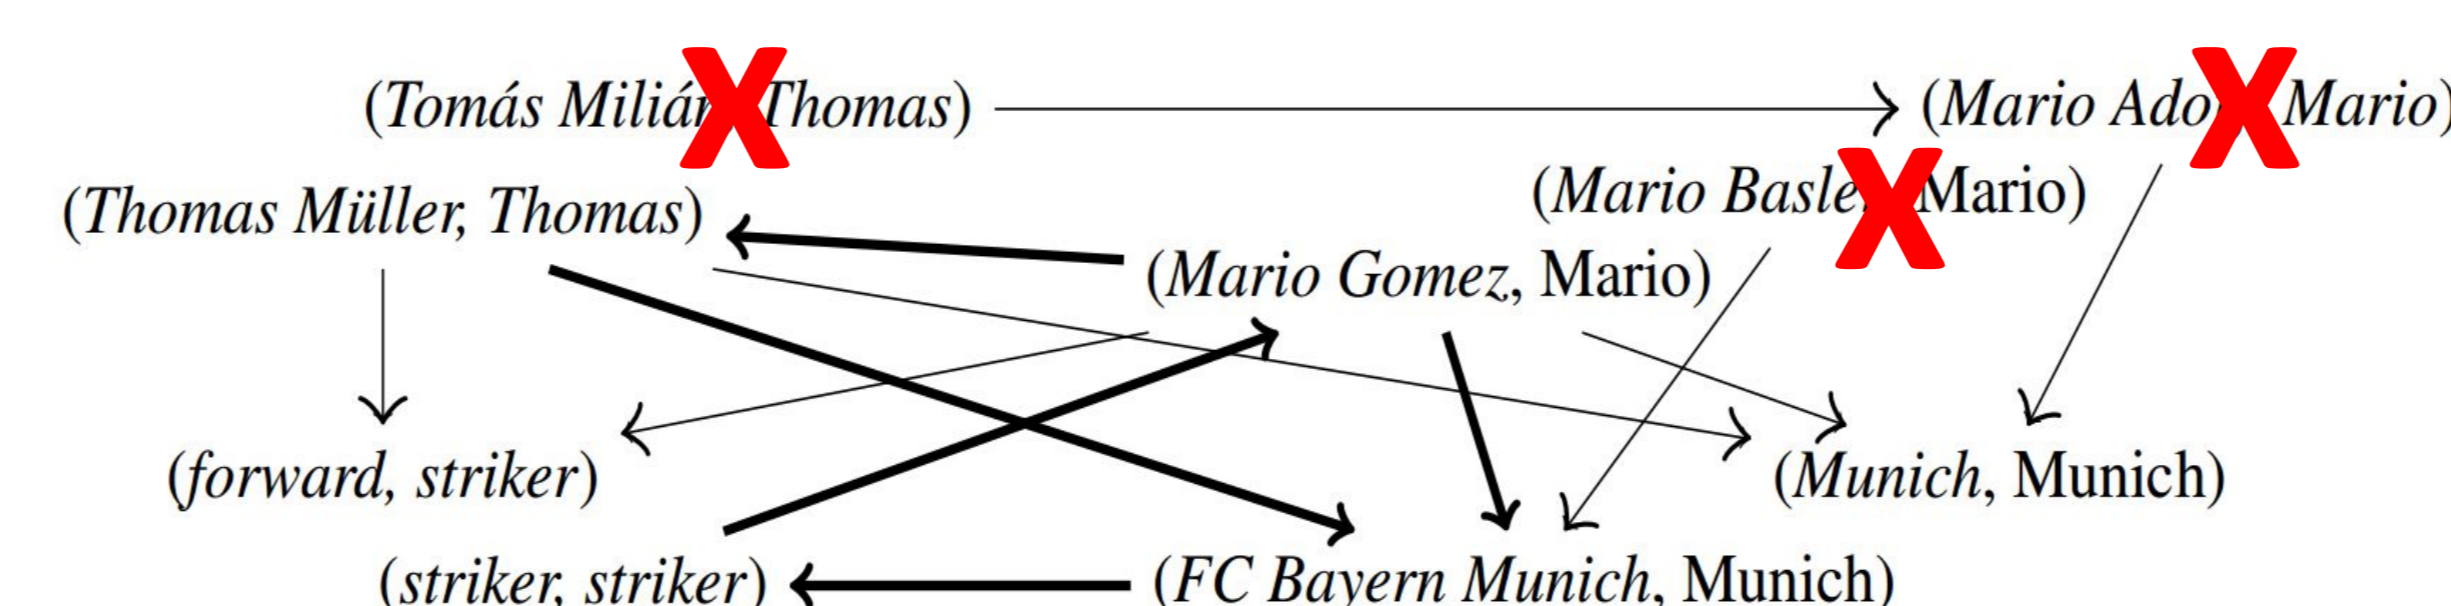
Example:

SemSign(striker) = {offside, soccer player, ..., athlete}

Thomas and Mario are strikers playing in Munich



We then extract a dense subgraph by iteratively removing loosely connected candidate meanings



Experimental Evaluations

Our experiments on **six gold-standard datasets** show the **state-of-the-art performance** of our approach as well as its **robustness across languages**.

System	Sens3		Sem07		SemEval-2013 English		French		German		Italian		Spanish	
	WN	BN	WN	Wiki	BN	Wiki	BN	Wiki	BN	Wiki	BN	Wiki	BN	
Babelfy	68.3	62.7	65.9	87.4	69.2	71.6	*56.9	81.6	69.4	84.3	66.6	83.8	69.5	
IMS	71.2	63.3	65.7	-	-	-	-	-	-	-	-	-	-	
UKB w2w	*65.3	*56.0	61.3	-	60.8	-	60.8	-	66.2	-	67.3	-	70.0	
UMCC-DLSI	-	-	64.7	54.8	68.5	*60.5	60.5	*58.1	62.8	*58.3	65.8	*61.0	71.0	
DAEBAK!	-	-	-	-	60.4	-	53.8	-	59.1	-	*61.3	-	60.0	
GETALP-BN	-	-	51.4	-	58.3	-	48.3	-	52.3	-	52.8	-	57.8	
MFS	70.3	65.8	*63.0	*80.3	*66.5	69.4	45.3	83.1	*67.4	82.3	57.5	82.4	*64.4	
Babelfy unif. weights	67.0	65.2	65.0	87.0	68.5	71.9	57.2	81.2	69.8	83.7	66.8	83.8	70.8	
Babelfy w/o dens. sub.	68.3	63.3	65.4	87.3	68.7	71.6	57.0	81.7	69.1	84.4	66.5	83.9	69.5	
Babelfy only concepts	68.2	62.7	65.5	83.0	68.7	70.2	56.6	79.3	69.3	83.0	66.3	84.0	69.7	
Babelfy on sentences	66.0	65.2	63.5	84.0	67.1	70.7	53.6	82.3	68.1	83.8	64.2	83.5	68.7	

Table 1: F1 scores (percentages) of the participating systems of SemEval-2013 task 12 together with MFS, UKB w2w, IMS, our system and its ablated versions on the Senseval-3, SemEval-2013 task 17 and SemEval-2013 datasets. The first system which has a statistically significant difference from the top system is marked with * ($\chi^2, p < 0.05$).

System	F1
(Ponzetto and Navigli, 2010)	85.5
Babelfy	84.6
UoR-SSI	84.1
UKB w2w	83.6
NUS-PT	*82.3
MFS	77.4
Babelfy unif. weights	85.7
Babelfy w/o dens. sub.	84.9
Babelfy only concepts	85.3
Babelfy on sentences	82.3

Table 2: F1 score (percentages) on the SemEval-2007 task 7. The first system which has a statistically significant difference from the top system is marked with * ($\chi^2, p < 0.05$).

System	KORE50	CoNLL
Babelfy	71.5	82.1
KORE-LSH-G	64.6	81.8
KORE	63.9	*80.7
MW	*57.6	82.3
Tagme	56.3	70.1
KPCS	55.6	82.2
KORE-LSH-F	53.2	81.2
UKB w2w (on BabelNet)	52.1	71.8
Illinois Wikifier	41.7	72.4
DBpedia Spotlight	35.4	34.0
Babelfy unif. weights	69.4	81.7
Babelfy w/o dens. sub.	62.5	78.1
Babelfy only NE	68.1	78.8

Table 3: Accuracy (percentages) of state-of-the-art EL systems and our system on KORE50 and AIDA-CoNLL. The first system with a statistically significant difference from the top system is marked with * ($\chi^2, p < 0.05$).

Conclusions

1. We presented a novel, integrated **state-of-the-art approach** to **Entity Linking** and **Word Sense Disambiguation**;
2. Our graph-based approach exploits the semantic network structure to its advantage: two key features of BabelNet, that is, its **multilinguality** and its **integration of lexicographic and encyclopedic knowledge**, make it possible to run our general, unified approach on the two tasks of Entity Linking and WSD in any of the languages covered by the semantic network;
3. At the core of our approach lies the **effective treatment of the high degree of ambiguity** of partial textual mentions by means of a 2-approximation algorithm for the densest subgraph problem;
4. We have developed a Web service which is available online at **babelfy.org**

References

- [1] Eneko Agirre, Oier Lopez de Lacalle, and Aitor Soroa. 2014. *Random walks for knowledge-based word sense disambiguation*. Computational Linguistics, 40(1):57-84.
- [2] Marco Cornolti, Paolo Ferragina, and Massimiliano Ciaramita. 2013. *A framework for benchmarking entity-annotation systems*. In Proc. of WWW, pages 249-260.
- [3] J. Hoffart, S. Seufert, D. B. Nguyen, M. Theobald, and G. Weikum. 2012. *KORE: keyphrase overlap relatedness for entity disambiguation*. In Proc. of CIKM, pages 545-554.
- [4] Roberto Navigli. 2009. *Word sense disambiguation: A survey*. ACM Comput. Surv., 41(2):1-69.
- [5] Roberto Navigli and Simone P. Ponzetto. 2012. *BabelNet: The automatic construction, evaluation and application of a wide-coverage multilingual semantic network*. Artif. Intel., 193:217-250
- [6] Zhi Zhong and Hwee Tou Ng. 2010. *It makes sense: A wide-coverage word sense disambiguation system for free text*. In Proc. of ACL (Demo), pages 78-83.