

SENTIWORDNET 3.0: An Enhanced Lexical Resource for Sentiment Analysis and Opinion Mining

Stefano Baccianella, Andrea Esuli, and Fabrizio Sebastiani

Istituto di Scienza e Tecnologie dell'Informazione
Consiglio Nazionale delle Ricerche
Via Giuseppe Moruzzi 1, 56124 Pisa, Italy
E-mail: $\langle\text{firstname.lastname}\rangle@isti.cnr.it$

Abstract

In this work we present SENTIWORDNET 3.0, a lexical resource explicitly devised for supporting sentiment classification and opinion mining applications. SENTIWORDNET 3.0 is an improved version of SENTIWORDNET 1.0, a lexical resource publicly available for research purposes, now currently licensed to more than 300 research groups and used in a variety of research projects worldwide. Both SENTIWORDNET 1.0 and 3.0 are the result of automatically annotating all WORDNET synsets according to their degrees of positivity, negativity, and neutrality. SENTIWORDNET 1.0 and 3.0 differ (a) in the versions of WORDNET which they annotate (WORDNET 2.0 and 3.0, respectively), (b) in the algorithm used for automatically annotating WORDNET, which now includes (additionally to the previous semi-supervised learning step) a random-walk step for refining the scores. We here discuss SENTIWORDNET 3.0, especially focussing on the improvements concerning aspect (b) that it embodies with respect to version 1.0. We also report the results of evaluating SENTIWORDNET 3.0 against a fragment of WORDNET 3.0 manually annotated for positivity, negativity, and neutrality; these results indicate accuracy improvements of about 20% with respect to SENTIWORDNET 1.0.

1. Introduction

In this work we present SENTIWORDNET 3.0, an enhanced lexical resource explicitly devised for supporting sentiment classification and opinion mining applications (Pang and Lee, 2008). SENTIWORDNET 3.0 is an improved version of SENTIWORDNET 1.0 (Esuli and Sebastiani, 2006), a lexical resource publicly available for research purposes, now currently licensed to more than 300 research groups and used in a variety of research projects worldwide.

SENTIWORDNET is the result of the automatic annotation of all the synsets of WORDNET according to the notions of “positivity”, “negativity”, and “neutrality”. Each synset s is associated to three numerical scores $Pos(s)$, $Neg(s)$, and $Obj(s)$ which indicate how positive, negative, and “objective” (i.e., neutral) the terms contained in the synset are. Different senses of the same term may thus have different opinion-related properties. For example, in SENTIWORDNET 1.0 the synset $[estimable(J, 3)]^1$, corresponding to the sense “*may be computed or estimated*” of the adjective *estimable*, has an Obj score of 1.0 (and Pos and Neg scores of 0.0), while the synset $[estimable(J, 1)]$ corresponding to the sense “*deserving of respect or high regard*” has a Pos score of 0.75, a Neg score of 0.0, and an Obj score of 0.25.

Each of the three scores ranges in the interval $[0.0, 1.0]$, and their sum is 1.0 for each synset. This means that a synset may have nonzero scores for all the three categories, which would indicate that the corresponding terms have, in the sense indicated by the synset, each of the three opinion-related properties to a certain degree.

This paper is organized as follows. Section 2. briefly

¹We here adopt the standard convention according to which a term enclosed in square brackets denotes a synset; thus $[poor(J, 7)]$ refers not just to the term *poor* but to the synset consisting of adjectives $\{inadequate(J, 2), poor(J, 7), short(J, 4)\}$.

charts the history of SENTIWORDNET, from its very earliest release to the current version, thus providing context for the following sections. Section 3. examines in detail the algorithm that we have used for generating SENTIWORDNET 3.0, while Section 4. discusses accuracy issues.

SENTIWORDNET 3.0 is freely available for non-profit research purposes from <http://sentiwordnet.isti.cnr.it/>

2. A brief history of SENTIWORDNET

Four different versions of SENTIWORDNET have been discussed in publications:

1. SENTIWORDNET 1.0, presented in (Esuli and Sebastiani, 2006) and publicly made available for research purposes;
2. SENTIWORDNET 1.1, only discussed in a technical report (Esuli and Sebastiani, 2007b) that never reached the publication stage;
3. SENTIWORDNET 2.0, only discussed in the second author’s PhD thesis (Esuli, 2008);
4. SENTIWORDNET 3.0, which is being presented here for the first time.

Since versions 1.1 and 2.0 have not been discussed in widely known formal publications, we here focus on discussing the differences between versions 1.0 and 3.0. The main differences are the following:

1. Version 1.0 (similarly to 1.1 and 2.0) consists of an annotation of the older WORDNET 2.0, while version 3.0 is an annotation of the newer WORDNET 3.0.
2. For SENTIWORDNET 1.0 (and 1.1), automatic annotation was carried out via a weak-supervision, semi-supervised learning algorithm. Conversely, for SEN-

TIWORDNET (2.0 and) 3.0 the results of this semi-supervised learning algorithm are only an intermediate step of the annotation process, since they are fed to an iterative random-walk process that is run to convergence. SENTIWORDNET (2.0 and) 3.0 is the output of the random-walk process after convergence has been reached.

3. Version 1.0 (and 1.1) uses the glosses of WORDNET synsets as semantic representations of the synsets themselves when a semi-supervised text classification process is invoked that classifies the (glosses of the) synsets into categories *Pos*, *Neg* and *Obj*. In version 2.0 this is the first step of the process; in the second step the random-walk process mentioned above uses not the raw glosses, but their automatically sense-disambiguated versions from EXTENDEDWORDNET (Harabagiu et al., 1999). In SENTIWORDNET 3.0 both the semi-supervised learning process (first step) and the random-walk process (second step) use instead the *manually* disambiguated glosses from the Princeton WordNet Gloss Corpus², which we assume to be more accurate than the ones from EXTENDEDWORDNET.

3. Generating SENTIWORDNET 3.0

We here summarize in more detail the automatic annotation process according to which SENTIWORDNET 3.0 is generated. This process consists of two steps, (1) a weak-supervision, semi-supervised learning step, and (2) a random-walk step.

3.1. The semi-supervised learning step

The semi-supervised learning step is identical to the process used for generating SENTIWORDNET 1.0; (Esuli and Sebastiani, 2006) can then be consulted for more details on this process. The step consists in turn of four substeps: (1) seed set expansion, (2) classifier training, (3) synset classification, and (4) classifier combination.

1. In Step (1), two small “seed” sets (one consisting of all the synsets containing 7 “paradigmatically positive” terms, and the other consisting of all the synsets containing 7 “paradigmatically negative” terms (Turney and Littman, 2003)) are automatically expanded by traversing a number of WORDNET binary relations than can be taken to either preserve or invert the *Pos* and *Neg* properties (i.e., connect synsets of a given polarity with other synsets either of the same polarity – e.g., the “also-see” relation – or of the opposite polarity – e.g., the “direct antonymy” relation), and by adding the synsets thus reached to the same seed set (for polarity-preserving relations) or to the other seed set (for polarity-inverting ones). This expansion can be performed with a certain “radius”; i.e., using radius k means adding to the seed sets all the synsets that are within distance k from the members of the original seed sets in the graph collectively resulting from the binary relationships considered.

2. In Step (2), the two sets of synsets generated in the previous step are used, along with another set of synsets assumed to have the *Obj* property, as training sets for training a ternary classifier (i.e. one that needs to classify a synset as *Pos*, *Neg*, or *Obj*). The glosses of the synsets are used by the training module instead of the synsets themselves, which means that the resulting classifier is indeed a gloss (rather than a synset) classifier. SENTIWORDNET 1.0 uses a “bag of words” model, according to which the gloss is represented by the (frequency-weighted) set of words occurring in it. In SENTIWORDNET 3.0 we instead leverage on the manually disambiguated glosses available from the Princeton WordNet Gloss Corpus, according to which a gloss is actually a sequence of WORDNET *synsets*. Our gloss classifiers are thus based on what might be called a “bag of synsets” model.
3. In Step (3) all WORDNET synsets (including those added to the seed sets in Step (2)) are classified as belonging to either *Pos*, *Neg*, or *Obj* via the classifier generated in Step (2).
4. Step (2) can be performed using different values of the radius parameter, and different supervised learning technologies. For reasons explained in detail in (Esuli and Sebastiani, 2006), annotation turns out to be more accurate if, rather than a single ternary classifier, a *committee* of ternary classifiers is generated, each of whose members results from a different combination of choices for these two parameters (radius and learning technology). We have set up our classifier committee as consisting of 8 members, resulting from four different choices of radius ($k \in \{0, 2, 4, 6\}$) and two different choices of learning technology (Rocchio and SVMs). In Step (4) the final *Pos* (resp., *Neg*, *Obj*) value of a given synset is generated as its average *Pos* (resp., *Neg*, *Obj*) value across the eight classifiers in the committee.

3.2. The random-walk step

The random-walk step consists of viewing WORDNET 3.0 as a graph, and running an iterative, “random-walk” process in which the $Pos(s)$ and $Neg(s)$ (and, consequently, $Obj(s)$) values, starting from those determined in the previous step, possibly change at each iteration. The random-walk step terminates when the iterative process has converged.

The graph used by the random-walk step is the one implicitly determined on WORDNET by the *definiens-definiendum* binary relationship; in other words, we assume the existence of a directed link from synset s_1 to synset s_2 if and only if s_1 (the *definiens*) occurs in the gloss of synset s_2 (the *definiendum*). The basic intuition here is that, if most of the terms that are being used to define a given term are positive (resp., negative), then there is a high probability that the term being defined is positive (resp., negative) too. In other words, positivity and negativity are seen as “flowing through the graph”, from the terms used in the definitions to the terms being defined.

²<http://wordnet.princeton.edu/glosstag.shtml>

However, it should be observed that, in “regular” WORDNET, the *definiendum* is a *synset* while the *definiens* is a non-disambiguated term, since glosses are sequences of non-disambiguated terms. In order to carry out the random-walk step, we need the glosses to be disambiguated against WORDNET itself, i.e., we need them to be sequences of WORDNET synsets. While for carrying out the random-walk step for SENTIWORDNET 2.0 we had used the automatically disambiguated glosses provided by EXTENDED-WORDNET (Harabagiu et al., 1999), for SENTIWORDNET 3.0 we use the manually disambiguated glosses available from the above-mentioned Princeton WordNet Gloss Corpus.

The mathematics behind the random-walk step is fully described in (Esuli and Sebastiani, 2007a), to which the interested reader is then referred for details. In that paper, the random-walk model we use here is referred to as “the inverse flow model”.

Two different random-walk processes are executed for the positivity and negativity dimensions, respectively, of SENTIWORDNET, producing two different rankings of the WORDNET synsets. However, the actual numerical values returned by the random-walk process are unfit to be used as the final *Pos* and *Neg* scores, since they are all too small (the synset top-ranked for positivity would obtain a *Pos* score of $8.3 * 10^{-6}$); as a result, even the top-ranked positive synsets would turn out to be overwhelmingly neutral and only feebly positive. Since, as we have observed, both the positivity and negativity scores resulting from the semi-supervised learning step follow a power law distribution (i.e., very few synsets have a very high *Pos* (resp., *Neg*) score, while very many synsets are mostly neutral), we have thus fit these scores with a function of the form $F_{Pos}(x) = a_1 x^{b_1}$ (resp., $F_{Neg}(x) = a_2 x^{b_2}$), thus determining the a_1 and b_1 (resp., a_2 and b_2) values that best fit the actual distribution of values. The final SENTIWORDNET 3.0 *Pos*(s) (resp., *Neg*) values are then determined by applying the resulting function $F_{Pos}(x) = a_1 x^{b_1}$ (resp., $F_{Neg}(x) = a_2 x^{b_2}$) to the ranking by positivity (resp., by negativity) produced by the random-walk process.

Obj(S) values are then assigned so as to make the three values sum up to one. In the case in which $Pos(s) + Neg(s) > 1$ we have normalized the two values to sum up to one and we have set $Obj(s) = 0^3$.

As an example, Table 1 reports the 10 top-ranked positive synsets and the 10 top-ranked negative synsets in SENTIWORDNET 3.0.

4. Evaluating SENTIWORDNET 3.0

For evaluating the accuracy of SENTIWORDNET 3.0 we follow the methodology discussed in (Esuli, 2008). This consists in comparing a small, manually annotated subset of WORDNET against the automatic annotations of the same synsets as from SENTIWORDNET.

4.1. Micro-WN(Op)-3.0

In (Esuli, 2008), SENTIWORDNET 1.0, 1.1 and 2.0 were evaluated on Micro-WN(Op) (Cerini et al., 2007), a carefully balanced set of 1,105 WORDNET synsets manually

annotated according to their degrees of positivity, negativity, and neutrality.

Micro-WN(Op) consists of 1,105 synsets manually annotated by a group of five human annotators (hereafter called J1, ..., J5); each synset s is assigned three scores $Pos(s)$, $Neg(s)$, and $Obj(s)$, with the three scores summing up to 1. Synsets 1-110 (here collectively called Micro-WN(Op)(1)) were tagged by all the annotators working together, so as to develop a common understanding of the semantics of the three categories; then, J1, J2 and J3 independently tagged all of synsets 111–606 (Micro-WN(Op)(2)), while J4 and J5 independently tagged all of synsets 607–1105 (Micro-WN(Op)(3)). Our evaluation is performed on the union of synsets composing Micro-WN(Op)(2) and Micro-WN(Op)(3). It is noteworthy that Micro-WN(Op) as a whole, and each of its subsets, are representative of the distribution of parts of speech in WORDNET: this means that, e.g., if $x\%$ of WORDNET synsets are nouns, also $x\%$ of Micro-WN(Op) synsets are nouns. Moreover, this property also holds for each single part Micro-WN(Op)(x) of Micro-WN(Op).

As for the evaluation of SENTIWORDNET 3.0, it should be noted that Micro-WN(Op) is the annotation of a subset of WORDNET 2.0, and cannot be directly used for evaluating SENTIWORDNET 3.0, which consists of an annotation of WORDNET 3.0. Deciding which WORDNET 3.0 synset corresponds to a given synset in Micro-WN(Op) cannot be determined with certainty, and may even be considered an ill-posed question. In fact, several of the synsets in Micro-WN(Op) do not exist any longer in WORDNET 3.0, at least in the same form. For example, the synset [good (A, 22)] does no longer exist, while the synset {gloomy (A, 2), drab (A, 3), dreary (A, 1), dingy (A, 3), sorry (A, 6), dismal (A, 1)} now contains not only all of these words (although with different sense numbers) but also blue (A, 3), dark (A, 9), disconsolate (A, 2), and grim (A, 6).

As a result, we decided to develop an automatic mapping method that, given a synset s in WORDNET 2.0, identifies its analogue in WORDNET 3.0. We then took all of the WORDNET 2.0 synsets in Micro-WN(Op), identified their WORDNET 3.0 analogues, assigned them the same $Pos(s)$, $Neg(s)$, and $Obj(s)$ as in the original Micro-WN(Op) synset, and used the resulting 1,105 annotated WORDNET 3.0 synsets as the gold standard against which to evaluate SENTIWORDNET 3.0.

Our synset mapping method is based on the combination of three mapping strategies, which we apply in this order:

1. **WORDNET sense mappings:** We first use the sense mappings between WORDNET 2.0 and 3.0 available at <http://wordnetcode.princeton.edu/3.0/WNsnsmap-3.0.tar.gz>. These mappings were derived automatically using a number of heuristics, and are unfortunately limited to nouns and verbs only. Each mapping has a confidence value associated to it, ranging from 0 (lowest confidence) to 100 (highest confidence). The majority of mappings have a 100 confidence score associated to them. As recom-

³This happened only for 16 synsets.

Table 1: The 10 top-ranked positive synsets and the 10 top-ranked negative synsets in SENTIWORDNET 3.0.

Rank	Positive	Negative
1	good#n#2 goodness#n#2	abject#a#2
2	better_off#a#1	deplorable#a#1 distressing#a#2 lamentable#a#1 pitiful#a#2 sad#a#3 sorry#a#2
3	divine#a#6 elysian#a#2 inspired#a#1	bad#a#10 unfit#a#3 unsound#a#5
4	good_enough#a#1	scrimy#a#1
5	solid#a#1	cheapjack#a#1 shoddy#a#1 tawdry#a#2
6	superb#a#2	unfortunate#a#3
7	good#a#3	inauspicious#a#1 unfortunate#a#2
8	goody-goody#a#1	unfortunate#a#1
9	amiable#a#1 good-humored#a#1 good- humoured#a#1	dispossessed#a#1 homeless#a#2 roof- less#a#2
10	gainly#a#1	hapless#a#1 miserable#a#2 misfortu- nate#a#1 pathetic#a#1 piteous#a#1 pitiable#a#2 pitiful#a#3 poor#a#1 wretched#a#5

mended in the documentation associated to the mappings, we have used only the highest-valued mappings (those with a 100 score), ignoring the others. Heuristics used for the determination of mappings include the comparison of sense keys, similarity of synset terms, and relative tree location (comparison of hypernyms). By using these mappings we have mapped 269 Micro-WN(Op) synsets to WORDNET 3.0.

- Synset term matching:** If a Micro-WN(Op) synset s_i (that has not already been mapped in the previous step) contains exactly the same terms of a WORDNET 3.0 synset s_j , and such set of terms appears only in one synset in both WORDNET 2.0 and 3.0, we consider s_i and s_j to represent the same concept.
- Gloss similarity:** For each Micro-WN(Op) synset s_i that has not been mapped by the previous two methods, we compute the similarity between its gloss and the glosses of all WORDNET 3.0 synsets, where a gloss g_i is represented by the set of all character trigrams contained in it. Similarity is computed via the Dice coefficient⁴

$$Dice(g_1, g_2) = \frac{2|g_1 \cap g_2|}{|g_1| + |g_2|} \quad (1)$$

In Equation 1 a higher $Dice(g_1, g_2)$ value means a stronger similarity. Given a Micro-WN(Op) synset s_i , its most similar WORDNET 3.0 gloss is determined, and the corresponding synset is chosen as the one matching s_i .

The Princeton research group had originally not used gloss similarity to produce the sense mappings used in Step 1 because, as reported in the documentation distributed with the mappings, ‘‘Glosses (...) are often significantly modified’’. We have found, by manually inspecting a sample of the results, that gloss similarity mapping was rather effective in our case.

⁴See also http://en.wikipedia.org/wiki/Dice_coefficient

The final result of this mapping process, that we call Micro-WN(Op)-3.0, is publicly available at <http://sentiwordnet.isti.cnr.it/>. It should be noted that the results of the automatic mapping process have not been *completely* checked for correctness, since checking if there is a better map for Micro-WN(Op) synset s than the current map requires in theory to search among all the WORDNET 3.0 synsets with the same POS. Therefore, the results of evaluations obtained on Micro-WN(Op)-3.0 are not directly comparable with those obtained on the original Micro-WN(Op).

4.2. Evaluation measure

In order to evaluate the quality of SENTIWORDNET we test how well it *ranks* by positivity (resp., negativity) the synsets in Micro-WN(Op)-3.0. As our gold standard we thus use a ranking by positivity (resp., negativity) of Micro-WN(Op)-3.0, obtained by sorting the Micro-WN(Op)-3.0 synsets according to their $Pos(s)$ (resp., $Neg(s)$) values). Similarly, we generate a ranking by positivity (resp., negativity) of the same synsets from the $Pos(s)$ and $Neg(s)$ values assigned by SENTIWORDNET 3.0, and compare them against the gold standard above.

We compare rankings by using the *p-normalized Kendall τ distance* (noted τ_p – see e.g., (Fagin et al., 2004)) between the gold standard rankings and the predicted rankings. The τ_p distance, a standard function for the evaluation of rankings that possibly admit ties, is defined as:

$$\tau_p = \frac{n_d + p \cdot n_u}{Z} \quad (2)$$

where n_d is the number of *discordant pairs*, i.e., pairs of objects ordered one way in the gold standard and the other way in the tested ranking; n_u is the number of pairs which are ordered (i.e., not tied) in the gold standard and are tied in the tested ranking; p is a penalization to be attributed to each such pair, set to $p = \frac{1}{2}$ (i.e., equal to the probability that a ranking algorithm correctly orders the pair by random guessing); and Z is a normalization factor (equal to the number of pairs that are ordered in the gold standard)

Table 2: τ_p values for the positivity and negativity rankings derived from SENTIWORDNET 1.0 and 3.0, as measured on Micro-WN(Op) and Micro-WN(Op)-3.0.

	Rankings	
	Positivity	Negativity
SENTIWORDNET 1.0	.349	.296
SENTIWORDNET 3.0	.281 (-19.48%)	.231 (-21.96%)

whose aim is to make the range of τ_p coincide with the $[0, 1]$ interval. Note that pairs tied in the gold standard are not considered in the evaluation. The lower the τ_p value the better; for a prediction which perfectly coincides with the gold standard, τ_p equals 0; for a prediction which is exactly the inverse of the gold standard, τ_p is equal to 1.

4.3. Results

Table 2 reports the τ_p values for the positivity and negativity rankings derived from SENTIWORDNET 1.0 and 3.0, as measured on Micro-WN(Op) and Micro-WN(Op)-3.0, respectively. The values for SENTIWORDNET 1.0 are extracted from (Esuli and Sebastiani, 2007b). As already pointed out in Section 4.1., we warn the reader that the comparison between the SENTIWORDNET 1.0 and 3.0 results is only partially reliable, since Micro-WN(Op)-3.0 (on which the SENTIWORDNET 3.0 results are based) may contain annotation errors introduced by the automatic mapping process.

Taking into account the above warning, we can observe that SENTIWORDNET 3.0 is substantially more accurate than SENTIWORDNET 1.0, with a 19.48% relative improvement for the ranking by positivity and a 21.96% improvement for the ranking by negativity.

We have also measured (see Table 3) the difference in accuracy between the rankings produced by SENTIWORDNET 3.0-semi and SENTIWORDNET 3.0, where by ‘‘SENTIWORDNET 3.0-semi’’ we refer to the outcome of the semi-supervised learning step (described in Section 3.1.) that led to the generation of SENTIWORDNET 3.0. The reason we have measured this difference is to check whether the random-walk step of Section 3.2. is indeed beneficial. The relative improvement of SENTIWORDNET 3.0 with respect to SENTIWORDNET 3.0-semi is 17.11% for the ranking by positivity, and 19.23% for the ranking by negativity; this unequivocally shows that the random-walk process is indeed beneficial.

It would certainly have been interesting to also measure the impact that the manually disambiguated glosses available from the Princeton WordNet Gloss Corpus have had in generating SENTIWORDNET, by comparing the performance obtained by using them (either in the semi-supervised learning step, or in the random-walk step, or in both) with the performance obtained by using the automatically disambiguated ones from EXTENDEDWORDNET. Unfortunately, this is not possible, since the former glosses are available for WORDNET-3.0 only, while EXTENDEDWORDNET is available for WORDNET-2.0 only.

Table 3: τ_p values for the positivity and negativity rankings derived from (a) the results of the semi-supervised learning step of SENTIWORDNET 3.0, and (b) SENTIWORDNET 3.0, as measured on Micro-WN(Op)-3.0.

	Rankings	
	Positivity	Negativity
SENTIWORDNET 3.0-semi	.339	.286
SENTIWORDNET 3.0	.281 (-17.11%)	.231 (-19.23%)

5. References

- S. Cerini, V. Compagnoni, A. Demontis, M. Formentelli, and G. Gandini. 2007. Micro-WNOp: A gold standard for the evaluation of automatically compiled lexical resources for opinion mining. In Andrea Sansò, editor, *Language resources and linguistic theory: Typology, second language acquisition, English linguistics*, pages 200–210. Franco Angeli Editore, Milano, IT.
- Andrea Esuli and Fabrizio Sebastiani. 2006. SENTIWORDNET: A publicly available lexical resource for opinion mining. In *Proceedings of the 5th Conference on Language Resources and Evaluation (LREC’06)*, pages 417–422, Genova, IT.
- Andrea Esuli and Fabrizio Sebastiani. 2007a. Random-walk models of term semantics: An application to opinion-related properties. In *Proceedings of the 3rd Language Technology Conference (LTC’07)*, pages 221–225, Poznań, PL.
- Andrea Esuli and Fabrizio Sebastiani. 2007b. SENTIWORDNET: A high-coverage lexical resource for opinion mining. Technical Report 2007-TR-02, Istituto di Scienza e Tecnologie dell’Informazione, Consiglio Nazionale delle Ricerche, Pisa, IT.
- Andrea Esuli. 2008. *Automatic Generation of Lexical Resources for Opinion Mining: Models, Algorithms, and Applications*. Ph.D. thesis, Scuola di Dottorato in Ingegneria ‘‘Leonardo da Vinci’’, University of Pisa, Pisa, IT.
- Ronald Fagin, Ravi Kumar, Mohammad Mahdiany, D. Sivakumar, and Erik Veez. 2004. Comparing and aggregating rankings with ties. In *Proceedings of ACM International Conference on Principles of Database Systems (PODS’04)*, pages 47–58, Paris, FR.
- Sanda M. Harabagiu, George A. Miller, and Dan I. Moldovan. 1999. WordNet 2: A morphologically and semantically enhanced resource. In *Proceedings of the ACL Workshop on Standardizing Lexical Resources (SIGLEX’99)*, pages 1–8, College Park, US.
- Bo Pang and Lillian Lee. 2008. Opinion mining and sentiment analysis. *Foundations and Trends in Information Retrieval*, 2(1/2):1–135.
- Peter D. Turney and Michael L. Littman. 2003. Measuring praise and criticism: Inference of semantic orientation from association. *ACM Transactions on Information Systems*, 21(4):315–346.