

ABOUT THE EFFECTS OF SENTIMENTS ON TOPIC DETECTION IN SOCIAL NETWORKS

KAREL GUTIERREZ-BATISTA, JESÚS R. CAMPAÑA, SANDRO MARTINEZ-FOLGOSO, M. AMPARO VILA &
MARIA J. MARTIN-BAUTISTA

Department of Computer Science and Artificial Intelligence, ETSIIT – University of Granada, 18071,
Granada, Spain

ABSTRACT

Topic detection from large textual data volumes extracted from Social Networks is an interesting research topic in the context of Big Data. The textual content present in Social Networks contains diverse information that can be exploited in order to obtain useful information. Topic detection and sentiment analysis in social networks are topics of widespread research. The study of both is sometimes intertwined as, usually, user messages revolve around a particular topic and express certain attitude of the user towards the topic discussed. However, this assumption is not valid for all messages as some of them express only general feelings or attitudes and do not refer to something in particular that covers up the topic discussed. In fact, these messages can influence the topic detection process. In this paper, we propose to obtain topics from massive quantities of text data extracted from social networks, without using previous information, and only with the use of unsupervised data mining techniques. We analyze the influence of sentiments in messages and how they affect the topic detection task. Terms related to sentiments provide useful information for a variety of applications, but not for topic detection where they represent a source of unnecessary noise. Experiments are conducted on data obtained from Twitter social network.

1 INTRODUCTION

In the last 10 years, the popularity and dramatic growth of the use of social networks has led to the creation of a huge amount of textual data by tens of millions of users on a daily basis [1]. A large amount of unstructured texts can be found in repositories from emails, work documents, surveys, opinions in websites and social networks, etc. The automatic processing of these textual data, which also lack structure, is very difficult. A solution to this problem is the automatic detection of the main topics discussed in the texts, which contain relevant information and can be presented to the users in a legible form.

Texts from social networks usually are harder to process as they contain informal language, and can focus on various topics or none at all. Social networks are known to contain a number of messages expressing the opinions of users. Sometimes, the structure and vocabulary used in these opinions can affect the process of topic detection.

This article proposes an analysis of the effect of sentiments on topic detection in social networks. We use a new approach to automatically detect the main topics in textual data from social networks by using data mining techniques and a multilingual knowledge base. This approach can group texts from a set of labels corresponding to a taxonomy, which replaces the original terms of the text after being syntactically processed and disambiguated. When those words have been replaced by the



This paper is part of the Proceedings of the International Conference on Big Data
(Big Data 2016)
www.witconferences.com

© 2016 WIT Press, www.witpress.com

ISSN: 1755-7437 (paper format), ISSN: 1755-7445 (online), <http://www.witpress.com/journals>

DOI: 10.2495/DNE-V11-N3-387-395

labels, the dimension of the problem is significantly reduced. The last step consists in labelling the obtained groups.

To achieve this, we propose a methodology based on the use of lexical resources, such as Multi-lingual Central Repository 3.0 (MCR 3.0) [2] and SentiWordNet [3], and on the tools Stanford Part-of-Speech (POS) Tagger [4] and Stanford Named Entity Recognition (NER) [5] combined with Hierarchical Clustering algorithms.

We use this methodology to generate clusters from text using two datasets generated from the same data, one with sentiment words and another with those words removed. This way we can see how the same procedure yields different results according to the data processing performed.

The rest of the article is structured as follows. A review of previous related work on this topic is presented in Section 2. Section 3 describes the methodology for the automatic topic detection in textual data from social networks. Then, Section 4 presents the experimental results and finally, Section 5 contains the conclusions and the future work.

2 RELATED WORK

In order to detect topics from text data we rely on hierarchical clustering on texts. Hierarchical clustering algorithms for textual data have been widely studied in the literature [6,7]. We can find a general overview of the different types of traditional agglomerative hierarchical clustering algorithms and how they can be applied to textual data in [8,9].

There are topic detection approaches specifically designed to deal with Big Data. A news search engine for interesting topics is described in [10]. The information is obtained from an information flow, and it is organized in a way that the user does not feel overwhelmed. Different topics are identified in a time window of Twitter messages for sports and politics, and are processed using a time dependent modification of tf-idf. In Gao *et. al.* [11] a new topic detection algorithm for internet news about major disasters is presented. The approach is based on the Group Average Hierarchical Clustering (GAHC). A new algorithm that combines characteristics from divisive and agglomerative algorithms is presented in [12]. The proposal is a new kind of agglomerative algorithm called ‘constrained agglomerative algorithm’, which is tested on different data sources. In all cases, data processing is conducted ignoring sentiment words that are treated as any other regular word.

Topic detection is widely used in social media to understand textual data and has many applications, such as content filtering [13,14], user interest modelling [15], interesting event monitoring [16,17] and sentiment analysis [18]. Sentiment analysis is a widespread application in social networks because users tend to express an opinion in their messages, that behaviour makes difficult topic detection as sentiments and topics are mixed in the texts. A work to empirically compare Twitter with a traditional medium, as the New York Times, was presented in [19]. The approach used topic modelling with a Twitter-LDA model. This work pointed out the great importance of opinionated text in Twitter and how it was different from traditional mediums.

These works deal with topic detection in social networks where the data belong to a specific domain. Text data is analyzed during a time interval and assigned to a specific topic. Sentiments and opinions are generally ignored in these approaches.

In this paper, we present a topic detection methodology for generic domains (such as in social networks), and analyze the effects of affective lexicon in this task.

3 METHODOLOGY

The methodology presents four stages: syntactic preprocessing, semantic preprocessing, hierarchical clustering for topic detection and cluster labelling. For a set of textual data, the methodology will

be able to detect the topic to which each text belongs. In each stage a different model is used to represent data in the form of a matrix.

These are the matrices used:

- *Term-Document Matrix* (M_1): Represents the terms in the documents before processing.
- *Processed Term-Document Matrix* (M_2): Represents the terms in the documents after syntactic preprocessing.
- *Label-Document Matrix* (M_3): Contains the weights of the labels obtained from the terms in the documents.
- *Label-Cluster Matrix* (M_4): Represents the labels associated to the clusters formed.

To test the influence of sentiments on topic detection, we apply the methodology by changing the semantic preprocessing stage where we generate two different sets of terms, one with all terms extracted from texts, and the other with the affective lexicon that denotes sentiments removed.

The next sections present a detailed description of the tasks performed in each stage.

3.1 Syntactic preprocessing

This first stage consists in executing a syntactic cleaning applying filters to the textual data to make their automatic processing easier. Terms are tagged with their POS, assigning a grammatical category (noun, adjective, verb, etc.), to each word from the text. Then, terms are annotated as entities (names or people, companies, and places) if appropriate. These processes are performed by means of the Stanford POS and Stanford NER tools, respectively. To finish standard syntactic preprocessing, tokenization is applied, which removes the punctuation marks that complicate the automatic processing of the texts.

After that, the necessary filters to delete any words that provide no useful information for the topic detection are applied. Those words are:

- stop words,
- words not identified as nouns by the POS tagger,
- words identified as proper names by the entity identifier,
- words that have associated a sentiment value according to the lexical resource SentiWordNet 3.0,
- words not included in the external knowledge base MCR 3.0.

It is important to distinguish between the first three types of words that are removed in the normal processing, and the last two that are used to remove sentiment words.

We create two different datasets, the first one where only the first three kinds of words are removed (not removing sentiment words), and the second one where all five types of words are removed including the ones with sentiment. We will apply the remaining stages of the methodology to both datasets, in order to get a comparative measure between data with, and without sentiment words. The data sources selected are social networks, where most of them can be expected to contain an opinion on a given topic. In this case, we are interested only in the topics and not in the sentiments expressed, so we discard all terms expressing any sentiment and work only with those that are relevant to topic detection. But, could the sentiment words be relevant to the topic detection task? That is what we study with the two datasets created. The results of this study can be seen in Section 4.2 where a comparison between results obtained for both datasets is presented.

Once all the filters for word removal are applied, both datasets are processed the same way. In the next step, the texts are semantically preprocessed, as terms are processed for the topic detection. Terms not discarded at this stage are represented in matrix *Processed Term-Document* (M_2).

3.2 Semantic preprocessing

The objective of the semantic preprocessing is to homogenize the different syntactic representations of a concept in the text into a unique term. Each unique term will be replaced by its corresponding labels in the WordNet Domain taxonomy of the MCR 3.0 knowledge base.

MCR 3.0 is based on WordNet 3.0 and integrates WordNets in five languages, among them, English being interesting for this work. Besides, it integrates lexical resources such as WordNet Domains [20], a new version of Base Concepts, the AdimenSUMO ontology [21] and Top Ontology [22].

WordNet Domains (<http://wndomains.fbk.eu/>) is a lexical resource created in a semi-automatic way to provide WordNet with domain labels. Each synset of WordNet is noted with at least one domain label of WordNet Domains, which is hierarchically organized [20]. Mappings of text terms to WordNet Domains labels allow to create a semantic hierarchy of topics that represents the text. From the four levels present in WordNet Domains, the first two levels are too general domains (TOP, APPLIED_SCIENCE, FACTOTUM, FREE_TIME, HUMANITIES, PURE_SCIENCE and SOCIAL) and are excluded from the mappings. If all the terms in a text have no mapping to a valid domain, those texts will belong to the group defined as *Empty*.

In order to homogenize different syntactic representations and then map them to a domain, it is necessary to know the appropriate sense of the terms. The sense is obtained using the part of speech of the terms (obtained previously with the Stanford POS) and selecting a sense from the list of senses available in WordNet for the pair term/POS. Word sense disambiguation is a critical part of this stage, as it affects the correct identification of the labels of the term.

3.2.1 Disambiguation

Some electronic knowledge bases store values of the use of each sense of a word, making it possible to obtain the sense of a word through the use frequency. A trivial disambiguation algorithm consists in the selection of the most frequent sense of a term. Other algorithms execute the disambiguation based on dictionaries; their aim is to select the sense with a higher overlapping between the current context of the word and the different senses of such word in the dictionary.

Our approach to the disambiguation process is to use a heuristic one, which takes into account both considered approaches, using the MCR 3.0 knowledge base. First, we determine the sense selecting the one with a higher overlapping between its description in MCR 3.0 and the current context where the term has been used. Then we select the most frequent. If both algorithms select the same sense, this would be the sense selected for the terms; otherwise, the sense of such term would be a mixture of both. If, after the mapping, the term is assigned to a domain that is in the first two levels of MCR 3.0, the term is discarded as mentioned before.

After this stage, the data is ready to be used by the hierarchical clustering algorithms.

3.3 Automatic topic detection based on hierarchical clustering

Once the texts have been homogenized, the hierarchical clustering is executed using the assigned WordNet Domains labels. To do so, the weight matrix *Label-Document* (M_3) must be created. Next, the process for the creation of the weight matrix will be described. The following actions will be carried out for every term of a text:

1. The term is replaced by the label or labels with which the sense (synset) selected during the disambiguation process has been noted.
2. Each label e is assigned a value taking into account Equation 1, where C_e is the number of labels by which the corresponding term has been replaced and C_t , the number of terms of the analyzed text, after being syntactically preprocessed. If a term is discarded due to the fact that all its senses have been noted with labels corresponding to the first two levels of WordNet Domain, such term would not be taken into consideration to determine the value of the rest of the labels in the weight matrix.
3. If the label e appears more than once in one text, then, its final value will be the sum of all individual values.

$$v_e = \frac{1}{C_e * C_t} \tag{1}$$

In Fig. 1, a specific example of the algorithm used to create the weight matrix is shown. An original text is taken and syntactically analyzed, applying all filters to remove all those terms that do not provide any information for the detection of the topics. In this case, just the terms *ASSIGNMENT* and *EXAM* remain. Then, these terms are semantically preprocessed, establishing for each of them all possible senses located in the MCR 3.0 knowledge base, to later replace the term by the label or labels of the selected sense or synset during the disambiguation process.

In the example of Fig. 1, the term *ASSIGNMENT* presents six senses, having one of them the label {*LAW*}, another one, the labels {*LAW, ECONOMY*}, and the rest, {*FACTOTUM*}. The last are not taken into account during the disambiguation process, as has been previously commented. Once the term has been disambiguated, it is replaced by the labels {*LAW, ECONOMY*}, with which the selected sense was noted after the disambiguation of the term *ASSIGNMENT*. On the other hand, the term *EXAM* has just one sense, which has the label {*SCHOOL*}. Lastly, the weight matrix is created to then apply the hierarchical clustering algorithms. For the previous

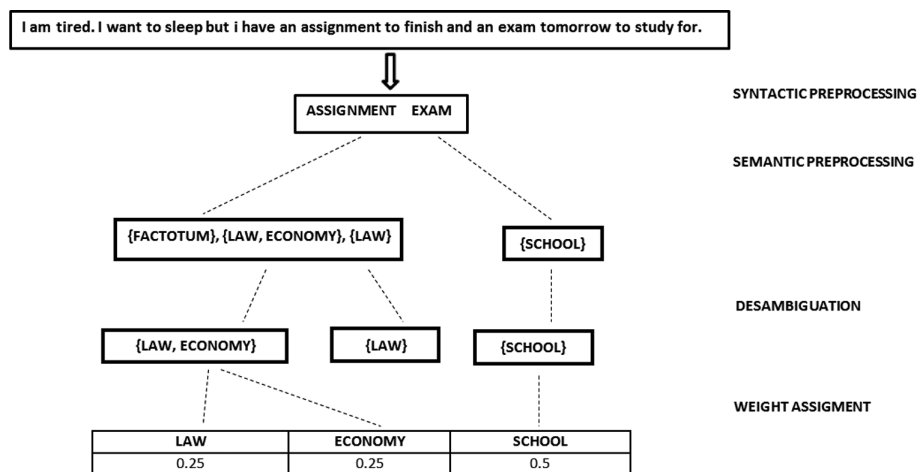


Figure 1: Creation of the labels-documents weight matrix.

specific case, the term *ASSIGNMENT* is replaced by two labels, being $C_e = 2$ and $C_t = 2$, which means that $v_e = \frac{1}{2*2} = 0.25$, whereas the term *EXAM* is replaced by one label, due to which $C_e = 1$ and $C_t = 2$, being $v_e = \frac{1}{1*2} = 0.5$.

After the creation of the weight matrix, the hierarchical clustering algorithm is executed to create the groups of texts from the semantics provided by WordNet Domains. This paper will analyze the behaviour of three classic hierarchical clustering algorithms (Complete Link, Average and Ward's Method), using for all cases the cosine similarity measure.

3.4 Cluster labelling

Cluster labelling tries to select descriptive labels for the groups obtained after applying a clustering algorithm. The labels are obtained after analyzing the documents belonging to a group or cluster. Labels should summarize the content of a cluster of documents, and differentiate it from the rest of the clusters. Labels selection is a highly relevant task, particularly in applications related to data analysis, where the user needs to know the content of a given cluster [23]. The correspondence matrix *Label-Cluster* (M_d) is created.

In this work, the Arithmetic Mean has been used to establish the most relevant labels of each cluster of texts. The Arithmetic Mean, is an internal labelling method that calculates the mean of all vectors of a cluster. This technique was selected due to the fact that the clusters of texts obtained from applying the hierarchical clustering methods can share one label, i.e. a topic can appear in more than one cluster, presumably with different relevance values.

4 EXPERIMENTAL RESULTS AND DISCUSSION

After applying the whole methodology to the two sets of data defined in Syntactic Preprocessing (removing sentiment words, and without removal), this section evaluates the results obtained. Both datasets were built from social network data and the only difference is the apparition of lack of sentiment words.

Evaluation is a critical task to determine which of the datasets created produces better results. There are a number of several different approaches to topic detection using hierarchical clustering algorithms. Most of these approaches deal with evaluation using datasets tagged manually by experts, and use evaluation measures such as *purity* and *entropy*, widely used in supervised problems [12, 24]. These measures are not applicable to our approach as texts are not previously categorized. The only valid measures that can be used to evaluate our approach are non-supervised measures.

We use the Silhouette Coefficient as evaluation measure [25], because it can determine the quality of separation and cohesion in clusters.

4.1 Dataset

The original data used has been extracted from Twitter, and downloaded from Sentiment140 (<http://www.sentiment140.com/>). This set contained 5000 tweets and 3396 distinct terms. It is important to note that, although these data are oriented towards the Sentiment Analysis, they are a very important source of data for experimentation, as they are useful to test the proposed approach, where the terms related to feelings are discarded to just work with those terms that provide useful information

for the detection of topics. This Twitter dataset is also the only one that can be downloaded at no cost for research purposes.

We have tested three classical hierarchical clustering algorithms (Complete Link, Average and Ward’s Method) together with our approach to topic detection, to have a benchmark to establish a comparison with the results obtained using our methodology, applying the filter for those terms expressing sentiments and without applying it. The Silhouette Coefficient has been computed by cuttings in 17, 25, 40, 60, 80, 100 and 120 clusters, making a total number of 42 experiments.

4.2 Evaluation method for hierarchical clustering algorithms

Table 1 shows the results obtained using the Silhouette Coefficient. It can be seen that when the proposed ontology-based methodology with the sentiment words removed is applied (**No Sentiment (NS)**), the results improve considerably with respect to when the sentiment words are maintained (**With Sentiment (WS)**). The Silhouette Coefficient values range from 0.06 to 0.36, which is a very good result for real texts and the most significant values are obtained by cuttings from 80 to 120 clusters.

As can be seen, removal of sentiment words helps to reduce the semantic noise introduced by those terms, which interferes with the topic detection task.

5 CONCLUSIONS AND FUTURE WORK

In this article, we have measured the effect of sentiments on topic detection tasks performed on Social Networks data. In order to measure the effect, an ontology-based methodology for automatic topic detection in social networks has been applied to the same data, divided into two datasets with the only difference that one has sentiment words removed and the other does not.

The proposed methodology can be applied to different social networks regardless of the topics mentioned in such social networks thanks to the knowledge base (MCR 3.0).

The methodology organizes the texts into clusters, where each cluster is represented as a set of topics. Three different hierarchical clustering algorithms are tested (Complete Link, Average and Ward’s Method). The output of the methodology is measured for both datasets using the Silhouette Coefficient, a well-known unsupervised measure that takes into account separation and cohesion between clusters. Experimental results obtained using Twitter dataset show that when sentiment words are removed the values obtained are better.

As future work, new semantic resources as ontologies and other electronic dictionaries, or the combination of both, are going to be tested in order to improve current results.

Table 1: Silhouette coefficient for Twitter data set (5,000 documents).

Cluster/Methods		17	25	40	60	80	100	120
Complete Link	WS	0.03	0.08	0.18	0.23	0.25	0.24	0.23
	NS	0.06	0.12	0.28	0.33	0.35	0.35	0.3532
Average	WS	0.11	0.15	0.2	0.25	0.26	0.25	0.24
	NS	0.17	0.22	0.28	0.34	0.35	0.36	0.36
Ward’s Method	WS	0.15	0.18	0.21	0.21	0.18	0.17	0.18
	NS	0.23	0.25	0.29	0.33	0.35	0.35	0.36

ACKNOWLEDGEMENT

The research reported in this paper was partially supported by the Andalusian Government (Junta de Andalucía) under projects P11-TIC-7460 and P10-TIC-6109 and by the Energy IN TIME project, funded from the European Union in the Seventh Framework Programme FP7-NMP under grant agreement 608981.

REFERENCES

- [1] Guille, A., Hacid, H., Favre, C. & Zighed, D.A., Information diffusion in online social networks: A survey. *SIGMOD Record*, **42**(2), 2013.
<http://dx.doi.org/10.1145/2503792.2503797>
- [2] Gonzalez-Agirre, A., Laparra, E. & Rigau, G., Multilingual central repository version 3.0. *Proceedings of the Eight International Conference on Language Resources and Evaluation (LREC'12)*, European Language Resources Association (ELRA): Istanbul, Turkey, 2012.
- [3] Esuli, A. & Sebastiani, F., Sentiwordnet: A publicly available lexical resource for opinion mining. *In Proceedings of the 5th Conference on Language Resources and Evaluation (LREC06)*, pp. 417–422, 2006.
- [4] Toutanova, K., Klein, D., Manning, C.D. & Singer, Y., Feature-rich part-of-speech tagging with a cyclic dependency network. *Proceedings of the 2003 Conference of the North American Chapter of the Association for Computational Linguistics on Human Language Technology – Volume 1*, Association for Computational Linguistics, NAACL '03: Stroudsburg, PA, USA, pp. 173–180, 2003.
<http://dx.doi.org/10.3115/1073445.1073478>
- [5] Finkel, J.R., Grenager, T. & Manning, C., Incorporating non-local information into information extraction systems by gibbs sampling. *Proceedings of the 43rd Annual Meeting on Association for Computational Linguistics*, Association for Computational Linguistics, ACL '05: Stroudsburg, PA, USA, pp. 363–370, 2005.
<http://dx.doi.org/10.3115/1219840.1219885>
- [6] RaghavaRao, N., Sravankumar, K. & Madhu, P., A survey on document clustering with hierarchical methods and similarity measures. *International Journal of Engineering Research & Technology (IJERT)*, **1**(7), 2012.
- [7] Deshmukh, D., Kamble, S. & Dandekar, P., Survey on hierarchical document clustering techniques fhc & f2 ihc. *International Journal of Advanced Research in Computer Science and Software Engineering*, **3**(7), pp. 157–161, 2013.
- [8] Voorhees, E.M., Implementing agglomerative hierarchic clustering algorithms for use in document retrieval. *Information Processing and Management*, **22**(6), pp. 465–476, 1986.
[http://dx.doi.org/10.1016/0306-4573\(86\)90097-X](http://dx.doi.org/10.1016/0306-4573(86)90097-X)
- [9] Willett, P., Recent trends in hierarchical document clustering: A critical review. *Information Processing and Management*, **24**(5), pp. 577–597, 1988.
[http://dx.doi.org/10.1016/0306-4573\(88\)90027-1](http://dx.doi.org/10.1016/0306-4573(88)90027-1)
- [10] Martin, C., Corney, D. & Goker, A., Mining newsworthy topics from social media. *BCS SGAI Workshop on Social Media Analysis*, Cambridge, UK, pp. 35–46, 2013.
- [11] Gao, N., Gao, L., He, Y., Wang, H. & Sun, Q., Topic detection based on group average hierarchical clustering. *International Conference on Advanced Cloud and Big Data (CBD, 2013)*, IEEE, pp. 88–92, 2013.
<http://dx.doi.org/10.1109/cbd.2013.38>

- [12] Zhao, Y. & Karypis, G., Evaluation of hierarchical clustering algorithms for document datasets. *Proceedings of the Eleventh International Conference on Information and Knowledge Management (CIKM '02)*, ACM: New York, NY, USA, pp. 515–524, 2002.
<http://dx.doi.org/10.1145/584792.584877>
- [13] Duan, J. & Zeng, J., Web objectionable text content detection using topic modeling technique. *Expert Systems with Applications*, **40**, pp. 6094–6104, 2013.
<http://dx.doi.org/10.1016/j.eswa.2013.05.032>
- [14] Martinez-Romo, J. & Araujo, L., Detecting malicious tweets in trending topics using a statistical analysis of language. *Expert Systems with Applications*, **40**, pp. 2992–3000, 2013.
<http://dx.doi.org/10.1016/j.eswa.2012.12.015>
- [15] Pennacchiotti, M. & Gurumurthy, S., Investigating topic models for social media user recommendation. *20th International Conference Companion on World Wide Web*, ACM: New York, NY, USA, pp. 101–102, 2011.
<http://dx.doi.org/10.1145/1963192.1963244>
- [16] Chung-Hong, L., Unsupervised and supervised learning to evaluate event relatedness based on content mining from social-media streams. *Expert Systems with Applications*, **39**(18), pp. 13338–13356, 2012.
<http://dx.doi.org/10.1016/j.eswa.2012.05.068>
- [17] Wu, J., Gao, W., Zhang, B., Liu, J. & Li, C., Cluster based detection and analysis of internet topics. *4th International Symposium on Computational Intelligence and Design, ISCID 2011*, Vol. 2, pp. 371–374, 2011.
<http://dx.doi.org/10.1109/iscid.2011.195>
- [18] Lin, C. & He, Y., Joint sentiment/topic model for sentiment analysis. *18th ACM Conference on Information and Knowledge Management 8CIKM09*, ACM: New York, NY, USA, pp. 375–384, 2009.
- [19] Zhao, W.X., Jiang, J., Weng, J., He, J., Lim, E.P., Yan, H. & Li, X., Comparing twitter and traditional media using topic models. *33rd European Conference on Advances in Information Retrieval (ECIR11)*, Springer-Verlag: Berlin, Heidelberg, pp. 338–349, 2011.
http://dx.doi.org/10.1007/978-3-642-20161-5_34
- [20] Magnini, B. & Cavaglia, G., Integrating subject field codes into wordnet. *LREC*, European Language Resources Association, 2000.
- [21] Pease, A., Niles, I. & Li, J., The suggested upper merged ontology: A large ontology for the semantic web and its applications. *Working Notes of the AAAI-2002 Workshop on Ontologies and the Semantic Web*, p. 2002, 2002.
- [22] Álvez, J., Atserias, J., Carrera, J., Climent, S., Laparra, E., Oliver, A. & Rigau, G., Complete and consistent annotation of wordnet using the top concept ontology. *Proceedings of the Sixth International Conference on Language Resources and Evaluation (LREC'08)*, European Language Resources Association (ELRA): Marrakech, Morocco, 2008. <http://www.lrec-conf.org/proceedings/lrec2008/>.
- [23] Manning, C.D., Raghavan, P. & Schütze, H., *Introduction to Information Retrieval*. Cambridge University Press: New York, NY, USA, 2008.
- [24] Zhao, Y. & Karypis, G., Empirical and theoretical comparisons of selected criterion functions for document clustering. *Machine Learning*, **55**(3), pp. 311–331, 2004.
<http://dx.doi.org/10.1023/B:MACH.0000027785.44527.d6>
- [25] Rousseeuw, P., Silhouettes: A graphical aid to the interpretation and validation of cluster analysis. *Journal of Computational and Applied Mathematics*, **20**(1), pp. 53–65, 1987.
[http://dx.doi.org/10.1016/0377-0427\(87\)90125-7](http://dx.doi.org/10.1016/0377-0427(87)90125-7)