International Journal of Computer Processing Of Languages
Vol. 22, Nos. 2 & 3 (2009) 205–218

DOI: 10.1142/S1793840609002044

Conveying Subjectivity of a Lexicon of One Language into Another Using a Bilingual Dictionary and a Link Analysis Algorithm

Jungi Kim*, Hun-Young Jung†, Yeha Lee‡, Jong-Hyeok Lee§
Division of Electrical and Computer Engineering,
Pohang University of Science and Technology,
San 31, Hyoja-Dong, Nam-Gu, Pohang, 790–784, Korea
*yangpa@postech.ac.kr
†blesshy@postech.ac.kr
‡sion@postech.ac.kr
§jhlee@postech.ac.kr

This paper proposes a method that automatically creates a sentiment lexicon in a new language using a sentiment lexicon in a resource–rich language with only a bilingual dictionary. We resolve some of the difficulties in selecting appropriate senses when translating lexicon, and present a framework that sequentially applies an iterative link analysis algorithm to enhance the quality of lexicons of both the source and target languages. The experimental results have empirically shown to improve the sentiment lexicon in the source language as well as create a good quality lexicon in the new language.

Keywords: Opinion Mining; Sentiment Lexicon; Lexicon Translation.

1. Introduction

With the recent pursuit of study in opinion mining such as subjectivity and polarity classification, opinion holder and opinion target extraction, and opinion summarization and trend prediction, much research effort has been spent on automating such tasks using various natural language processing approaches. Most previous researches, from constructing language resources [8, 21, 11, 19, 4, 1, 5, 6] and sentiment analysis [13, 24, 16, 17] to a variety of applications [9, 23, 25], have targeted English language only, and naturally many language resources for sentiment analysis have been created in English.

While a number of languages such as Japanese, Chinese, and German are directly employed in recent studies [20, 12, 10], other work [14, 17] have explored utilizing language resources in English to develop language resources and sentiment analysis tools in other languages.

Motivated by the latter approach, this paper presents a method for automatically creating a sentiment lexicon in a new language using a sentiment lexicon in a resource–rich language with the aids of 1) a bilingual dictionary of the two languages for translating the lexicon and 2) a link analysis algorithm for refining the relative rankings of the entries in the new lexicon, as well as the original lexicon. Translating sentiment lexicon using a bilingual dictionary faces several problems [17]: processing the inflected forms of entries without losing its subjective meaning, translating multi–word entries in the dictionary, and selecting the correct sense to translate among many candidates in an entry. Of the challenges mentioned, we address the problem of handling various senses in an entry, while leaving the rest as future work. Link analysis models have shown successful results in its recent applications to NLP tasks [16, 6]. Especially, [6] constructed a graph of WordNet synsets using glosses to create edges among the synsets, and learn opinion–related properties (ORPs) of synsets using PageRank, a popular random–walk model widely used in web searches, that ranks all the WordNet synsets in the graph according to evidence collected from its neighbors. The approach has shown to discern the ORPs of the synsets more accurately, especially when given an appropriate initial ORP value of the synsets. Adapting a similar framework, we have created a bipartite graph of lexicon entries, with entries of one language forming a cluster and the other language another, and applied a link analysis algorithm that is similar to both PageRank and HITS. The details of our link analysis model will be discussed in Section 3.2 of this paper.

Our work focuses on creating a sentiment lexicon in Korean utilizing sentiment lexicons in English; Korean is a relatively understudied language in sentiment analysis, and it is in urgent need of resources to jump–start its study. However, our work does not rely on any language–specific information but only requires a bilingual dictionary between the source and the target languages, making it easily applicable to other language pairs.

2. Related Work

Various sentiment lexicons have been used in many areas of opinion mining and sentiment analysis. Some lexicons are manually created [18, 22, 23] while others are the outcomes of the research efforts on automatically learning subjectivity from dictionary and thesaurus [11, 9, 13, 4, 17, 4, 5, 6] or from raw corpus [8, 21, 12].

There has also been efforts to utilize the language resources created in English for analyzing the sentiments in other languages; although in very limited fashion, [14] are the first to use English resources in German sentiment analysis, by translating a German e–mail into English, then applying English sentiment classifiers to the translated text. [17] was the first genuine multilingual work in sentiment analysis, in which languages resources developed for English are used for developing resources in Romanian by translating the sentiment lexicon using a freely available online dictionaries and creating a sentiment corpus through projection using a parallel corpus between English and Romanian and English subjectivity classifiers. Similar to the approach in [17], our work directly translates the sentiment lexicon in English into a target language. However, while they use a naive translation approach namely choosing the first sense of the translation candidates because dictionaries list the senses in order of the common usages hence the first sense being the most probable one, our work focuses on how to reduce the ambiguity errors while still maintaining a good number of translations.

[6] uses a graph representation of WordNet synsets and a random–walk model to simulate the dynamics of the WordNet synsets that have similar ORPs. In [6], a graph is constructed such that synsets of WordNet consist of nodes and edges connecting nodes with *similar*[a] meanings. While [6] employs glosses of WordNet entries to construct the edges between similar WordNet synsets, our work creates more reliable edges between vertices exploiting the bilingual dictionary such that a foreign word being the direct translation of a source word creates an edge between the two words.

3. Learning Sentiment Lexicon

To create a sentiment lexicon in Korean using an English sentiment lexicon, we adopt a three step approach; first, translate the English lexicon into Korean using a bilingual dictionary, refine the resulting lexicon using a link analysis model, then normalize the sentiment scores.

Sentiment lexicons vary in what information (*subjective/objective*, *positive/negative*) is tagged on which level of lexicon entries (word, POS–tagged word, sense) and how their strengths are measured (*weak/strong*, probability score $(0.0 \sim 1.0)$). We assume that our English sentiment lexicon contains English words with POS tags and semantic orientation with some measure for its strength (e.g. {abandon, verb, weak negative}, or {harm, verb, positive 0.0, negative 0.5, neutral 0.5}), and the Korean sentiment lexicon in similar format. However, our method could also be used to learn not only sentiment orientation but any ORPs whose strengths can be numerically transformed into scores to be used within our link analysis model.

3.1. Translating sentiment lexicon

Translating a sentiment lexicon into another language using a bilingual dictionary is a challenging task. Much of the subjective meaning of a lexicon can be lost when translating words that have different subjectivity in inflected forms, there are many multi–words that are not listed in the bilingual dictionary, and there are words that have various senses and different subjectivity associated with them [17].

[a]But not exploiting any relations defined in the WordNet such as synonymy, hypernymy, or hyponymy.

[17] relies on a heuristic method that translates only the first sense, since bilingual dictionaries usually order the translations such that more frequently used senses are listed before the less frequently used ones. Such a scheme would probably result in a lexicon with better quality in the sense of conveying subjectivity. However, it also reduces the size of the translated lexicon, limiting its application usages.

We present several naive heuristics that have different effects on the size and quality of the resulting lexicon, in a belief that more sophisticated heuristic would result in creating a lexicon with higher quality while maintaining a good number of entries. We assume that for each English word and its POS, our bilingual dictionary has multiple senses, with its rank in the reverse order of the usage frequency, and each sense also containing a number of translation candidates, whose rank is also ordered in reverse of its usage frequency.

First Word (FW) This approach assigns the sentiment scores of the English word to only the first word of the first sense. This translation scheme filters uncertain candidates, the size of the resulting lexicon being the smallest.

First Sense (FS) The approach taken in **FS** is similar to the one used in [17]. All the words in the first sense are assigned the sentiment scores of the English word, implying that different translation words with the same sense are equally likely to be translated.

All Senses (AS) **AS** assigns the sentiment scores of the English word to all the words in its translation candidates. This scheme produces the maximum number of Korean words, allowing unreliable words in the lexicon.

Sense Rank (SR) Korean words are assigned different scores by their sense ranks; words with higher sense ranks are assigned high sentiment scores, and vice versa. A simple formula of $\frac{NumSenses(w_e)-SenseRank(w_e)+1}{NumSenses(w_e)}$ is used.

Although these heuristics are very simple, they effectively control the size and reliability of the final translated lexicon, allowing us to observe the quality of the resulting lexicons in the evaluation process.

3.2. Refining sentiment lexicon with a link analysis algorithm

Similarly to [6], our approach uses a graph built from the words with ORPs as vertices, and the relations among the words as edges connecting the vertices. Unlike [6] that uses gloss of WordNet synsets to create semantic relations among the synsets, our approach utilizes a bilingual dictionary so that nodes connected by edges are direct translations of each other. These types of edges are more suited for building a much more semantically tight graph structure than the one using synset glosses.

Naturally, edges of direct translations connect English words to Korean words only, and Korean words only to English words. This type of graph is called a bipartite graph, where vertices are partitioned into two disjoint sets with no edges connecting any two vertices in the same set.

HITS is a link analysis algorithm that rates vertices of a graph by determining their "hubness" (connectedness to vertices with high "authoritativeness") and "authoritativeness" (connectedness to vertices with high "hubness") values, iteratively and recursively computing the centrality of a vertex within the graph structure [15].

Considering the hubness of an English vertex as its sentiment score, and the authoritativeness of a Korean vertex as the vertex with connectedness to English vertices with high hubness, HITS algorithm applied to the bipartite graph of bilingual dictionary entries can effectively learn the refined sentiment scores of a Korean lexicon, given that English lexicon holds its hubness in the process of learning the authoritativeness of Korean lexicon. Since the sentiment (authoritativeness) scores of a Korean lexicon are not reliable in the initial iterations of the algorithm, it is necessary to lower the variability of the hubness scores of English lexicon while raising the variability of authoritativeness when learning the sentiment scores of a Korean lexicon. Damping factor in PageRank algorithm [2] has similar effects on variability of the graph structure. The prior knowledge from English sentiment lexicon and its translation to Korean provides good candidates for prior scores (referred to as *internal source* in [6], e_k and e_e in Equation (1)).

Combining the ideas results in Equation (1) where $TC(w)$ is the set of translation candidates of a word w, α and β are damping factors for Korean and English vertices.

$$\begin{aligned} AUTH(w_k) &= (1-\alpha) * e_k + \alpha * \sum_{w_e \in TC(w_k)} HUB(w_e), \\ HUB(w_e) &= (1-\beta) * e_e + \beta * \sum_{w_k \in TC(w_e)} AUTH(w_k) \end{aligned} \tag{1}$$

Larger α indicates higher variability of authoritativeness of Korean vertices, that hubness of English vertices are trustworthy and actively affect the authoritativeness of Korean vertices, and vice versa for β.

Once the sentiment scores of a Korean lexicon is refined, the sentiment scores of Korean and English lexicons can be re–learned using the same algorithm to maximize the quality of the English lexicon as well, using the Equation (2).

$$\begin{aligned} AUTH(W_e) &= (1-\alpha) * e_e + \alpha * \sum_{W_k \in TC(W_e)} HUB(W_k), \\ HUB(W_k) &= (1-\beta) * e_k + \beta * \sum_{W_e \in TC(W_k)} AUTH(W_e) \end{aligned} \tag{2}$$

In summary, refining the sentiment lexicons in English and Korean is carried out on our two phase link analysis framework: first, running HITS with Korean words such as "authorities" and English words such as "hubs" to learn the authoritativeness of Korean words, and secondly, running HITS again with English words such as "authorities" and Korean words such as "hubs" to re–learn the authoritativeness of English words.

The link analysis model in each phase should take different values for α and β to adjust the variability of vertices accordingly.

Our framework runs on positive, negative, and neutral sentiments independently, producing separate rankings of lexicons for positive, negative, neutral scores.

3.3. Normalizing sentiment scores

After refining a sentiment lexicon with a link analysis algorithm, Korean words are assigned new sentiment scores and English words with refined scores. The three scores for each word, positive, negative, and neutral, however, do not sum up to 1.0 as they did in the source language lexicon; such property is necessary for comparing words' sentiment scores objectively. For example, if word A has larger positive score than word B, then word A should carry a more positive meaning than word B does.

Due to the differences in the graphical representation for each word such as the number of out- and in- links and the way the link analysis algorithm works, sentiment scores of the resulting lexicon are no longer bounded by the the range of 0.0 to 1.0 by which the initial sentiment scores of English words were bounded. Words with many inbound and outbound links tend to have larger scores than words with small number of links. The numerical comparisons between sentiment scores between words become no longer meaningful. However, the positive, negative, and neutral scores of a word maybe compared to one another, because these scores are induced for the same word hence using the same graphical structure. We normalize the sentiment scores of lexicon such that the sum of positive, negative, and neutral scores of a word add up to 1.0, by dividing each score with the sum of all three scores.

4. Experiments

4.1. Setup

The English lexicons we use in our experiments are the sentiment lexicon used in OpinionFinder (**OF**) [24][b] and SentiWordNet 1.0.1 (**SentiWN**) [5].[c]

OF is a set of English words and sentiment annotations collected from a number of sources of which some are manually developed while others automatically gathered. Each word in **OF** has a POS tag and categories of *Positive*/*Negative* and *Weak*/*Strong*. For our experimental purposes, *Weak* sentiment words were assigned the score of 0.5, and *Strong* words with 1.0. Neutral scores of words are estimated as 0.0 if sentiment strength is *Strong*, 0.5 if *Weak*, and 1.0 if not listed in the lexicon.

SentiWN is a set of WordNet synsets with automatically assigned positive, negative, and neutral probability scores. In our experiments, each word in a synset

[b] http://www.cs.pitt.edu/mpqa/
[c] http://sentiwordnet.isti.cnr.it/

is treated separately with the sentiment scores of the synset as its own, ignoring the synonym information provided by WordNet synsets.

We use a online bilingual dictionary provided by a portal website.[d] For our experiments, a total of $63,001$ English entries were accessed, corresponding to $142,791$ translated words in Korean.

Using different translation schemes in Section 3.1, both English lexicons are translated into Korean. The link analysis algorithm in section 3.2 is then tested with various sets of initial scores: uniform weight **UW** ($\frac{1}{|NumberofVertices|}$), and every combinations of English lexicons (**OF** and **SentiWN**) with translation schemes (**FW**, **FS**, **AS**, and **SR**).

The parameters α and β in Equations (1) and (2) are optimized on a held–out data using values from 0.1 to 0.9 with a step of 0.1.

4.2. Evaluation method

We followed the evaluation scheme in [6], which uses a Micro–WNOp corpus [3] [e] as a gold standard and the *p–normalized Kendall* τ *distance* (τ_p) [7] as the evaluation measure.

Micro–WNOp is a subset of WordNet that are tagged with ORPs by the number of English majoring MSc students. Divided into three sections (*Common*, *Group1*, *Group2*), each section contains a number of synsets with its positive and negative scores. For our research, we use *Group1* as a held–out data and *Group2* as a test data. We extract one positive and one negative scores by averaging all scores of evaluators. For optimizing and evaluating Korean subjectivity lexicon, 496 synsets in *Group1* and 499 synsets in *Group2* of Micro–WNOp was translated into Korean by a knowledgeable evaluator, fluent both in English and Korean. Korean words not appearing in any of the lexicons in our experiments were removed, resulting in 87 words and their associated sentiment scores as the gold standard.

The *p–normalized Kendall* τ *distance* is a measure of how much two ranked lists of items agree with each other. Given a set of items $\{o_1...o_n\}$, all possible pairs of items are tested, such that the agreements of their partial orders are compared in each list, counting discordant and tied pairs for penalization, the distance is defined as

$$\tau_p = \frac{n_d + \frac{1}{2} \times n_u}{Z} \tag{3}$$

where n_d is the number of discordant pairs (pairs differently ordered in each list), n_u is the number of pairs ordered in the gold standard but tied in the prediction, and Z is the number of pairs ordered in the gold standard.

The measure for a predicted list whose items are ranked in the same order as the gold standard is 0, indicating that there is no discordant or undecided pair of

[d] http://endic.naver.com/
[e] http://www.unipv.it/wnop/

items. In the opposite case, if items in a list are in reverse order of the items in the gold standard, then τ_p equals 1. If a list does not order items but rather returns an unordered list, then the measure becomes 0.5.

5. Results

5.1. Kendall τ distance

The experimental results show our proposed translation heuristics worked as we had expected: heuristics that translate only reliable words tend to have low τ_p and smaller number of translated words, while heuristics that translate more words have higher τ_p and bigger number of translated words.

Direct evaluation of **OF** lexicon results in poor score (Table 1). It is due to the initialization where all *Strong* subjective words have the sentiment score of 1.0, and *Weak*, 0.5, arising many tied pairs that are penalized in our evaluation measure. Once translated, however, the quality of the lexicon is better than the ones translated from **SentiWN** because when translated, scores are averaged so that the words now have different values than 0.0, 0.5 or 1.0, and **OF** contains some manually–developed resources while **SentiWN** is created in completely automatic fashion.

After applying the refinement algorithm using link analysis, we see drastic decrease in the number of Korean words, especially for the translation heuristics that generates larger number of candidate words (Table 2). This is due to that many of these candidate words do not have the same sentiments as the original sense of the English words and many of these words did not have any inbound and outbound links hence removed from the graph. We observe that most of the translation

Table 1. *p–normalized Kendall τ distance (τ_p)* and lexicon size for English lexicons and Korean translations.

	EN							
	SentiWN				OF			
POS	0.365				0.490			
NEG	0.310				0.494			
Size	10,631				8,221			
	KR							
	SentiWN				OF			
	FW	FS	AS	SR	FW	FS	AS	SR
POS	0.301	0.278	0.312	0.312	0.179	0.142	0.122	0.122
NEG	0.300	0.304	0.261	0.261	0.214	0.167	0.192	0.192
Size	37,812	68,382	142,791	142,791	4,270	10,558	32,322	32,322

Table 2. Changes in *p–normalized Kendall τ distance (τ_p)* and lexicon size of Korean Lexicon, after the execution of the first phase of the proposed link analysis model framework, using Korean Words as authorities and English words as hubs.

KR as authority, $\alpha = 0.6, \beta = 0.9$								
POSITIVE								
	SentiWN				OF			
	FW	FS	AS	SR	FW	FS	AS	SR
Before	0.301	0.278	0.312	0.312	0.179	0.142	0.122	0.122
After	0.285	0.273	0.293	0.293	0.132	0.117	0.110	0.112
Diff	−5.32%	−1.80%	−6.09%	−6.09%	−26.3%	−17.6%	−9.84%	−8.20%
NEGATIVE								
	SentiWN				OF			
	FW	FS	AS	SR	FW	FS	AS	SR
Before	0.300	0.304	0.261	0.261	0.214	0.167	0.192	0.192
After	0.291	0.293	0.254	0.254	0.202	0.160	0.186	0.190
Diff	−3.00%	−3.62%	−2.68%	−2.68%	−5.61%	−4.19%	−3.13%	−1.04%
Size	9,199	39,228	39,335	39,335	39,184	39,184	39,191	39,191

Table 3. Changes in *p–normalized Kendall τ distance (τ_p)* and lexicon size of English Lexicon, after the execution of the second phase of the proposed link analysis model framework, using English Words as authorities and Korean words as hubs.

EN as authority, $\alpha = 0.1, \beta = 0.1$								
POSITIVE								
	SentiWN				OF			
	FW	FS	AS	SR	FW	FS	AS	SR
Before	0.365				0.490			
After	0.340	0.338	0.342	0.342	0.355	0.335	0.335	0.333
Diff	−6.85%	−7.40%	−6.30%	−6.30%	−27.6%	−31.6%	−31.6%	−32.0%
NEGATIVE								
	SentiWN				OF			
	FW	FS	AS	SR	FW	FS	AS	SR
Before	0.310				0.494			
After	0.309	0.305	0.313	0.314	0.290	0.298	0.306	0.304
Diff	−0.323%	−1.61%	+0.968%	+1.29%	−41.3%	−39.7%	−38.1%	−38.5%
Size	73,931	73,931	73,935	73,935	73,931	73,931	73,931	73,931

heuristics produce about the same number of Korean sentiment words, regardless of sizes of the given source lexicons. As for the quality of the lexicon, semi-automatically constructed lexicon **OF** is in all cases measured higher than fully-automatically constructed **SentiWN**.

English lexicons produced from the refinement algorithm all improved over the original lexicons, except the negative **SentiWN** lexicon with **AS** and **SR** translation heuristics. We also observe the increase in the size of the lexicons, due to the addition of translation candidates of Korean sentiment words. In most combinations of lexicon and translation heuristics, **OF** scores better than **SWN**, but the difference is not as much as in Korean lexicons.

In conclusion, We observe that the proposed framework with two link analysis models has a compensating effect in each phase that the lexicons mutually complement each other in turn. The quality of the lexicons in every approach has shown to range from slightly negative (+1.29%) to exceptional (−41.3%).

5.2. Score distribution and semantic orientation

In Figures 1 and 2, we show the score distribution and semantic orientation of English lexicon generated from **OF** using **FS**.

In Figure 1, we observe that sentiment scores of English words are clustered around certain regions, indicating that sentiment scores have not diverged from the original numbers. We also present the distribution of semantic orientations (tendency toward positive or negative sense) of English words. In Figure 2, we observe evenly distributed clusters of semantic orientations.

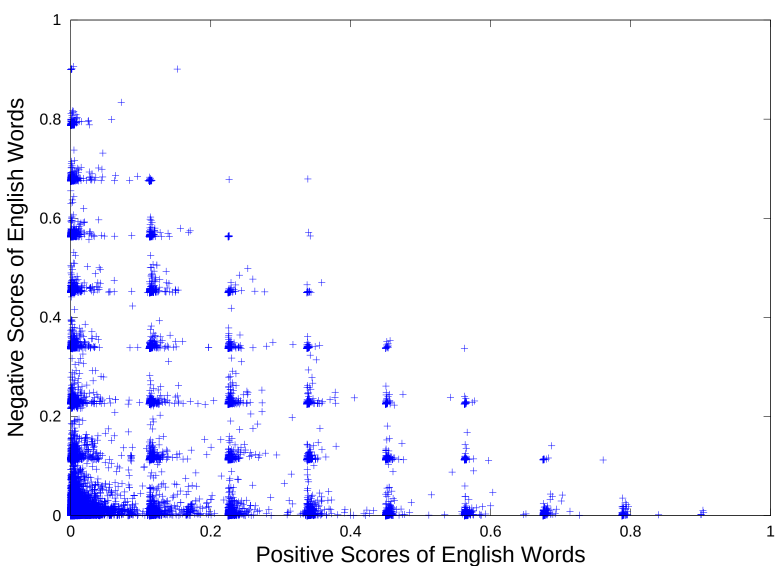

Figure 1. Positive and Negative Scores Distribution of English Lexicon Constructed from **OF** with **FS**.

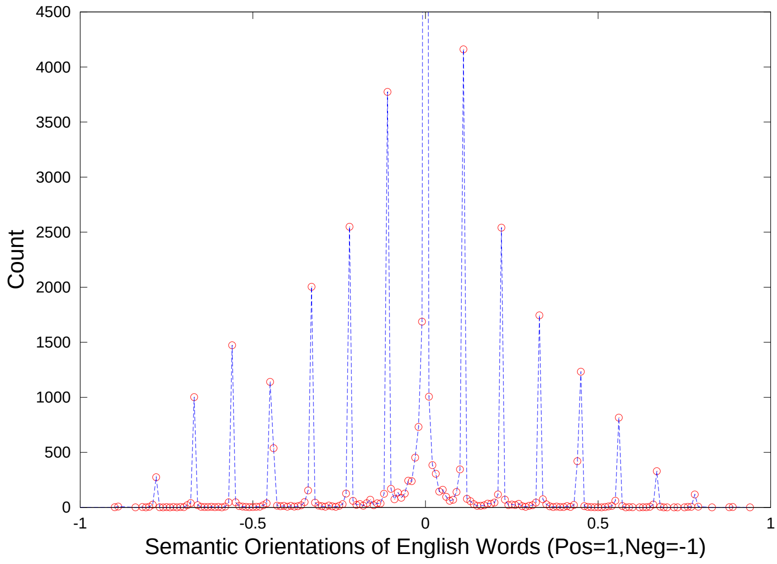

Figure 2. Semantic Orientations (*Positive Score – Negative Score*) of English Lexicon Constructed from **OF** with **FS**.

Figures 3 and 4 are the score distribution and semantic orientation graphs of a Korean lexicon constructed from **OF** with **FS**. Positive and negative scores of Korean lexicon are more evenly distributed all over the lower triangle of the first quadrant. In Figure 4, we also observe evenly distributed semantic orientations of Korean words, except for three notable clusters around −1.0, −0.5, 0.0, 0.5, and 1.0.

6. Conclusion

This paper investigated the feasibility of exploiting a sentiment lexicon in one language to developing a sentiment lexicon in another language with a bilingual dictionary as the only available language resource. Our proposed method of first translating the lexicon using the bilingual dictionary with several translation heuristics, then applying a framework that sequentially applies an iterative link analysis algorithm and score normalization technique to enhance the quality of lexicons of both the source and the target languages has been empirically shown to create good quality lexicons.

Unlike previous work, we have explored the possibility of regarding a language translation process as a subjectivity projection operation. We have also attempted to draw compensation interactions using a graph structure as a medium between the language pair.

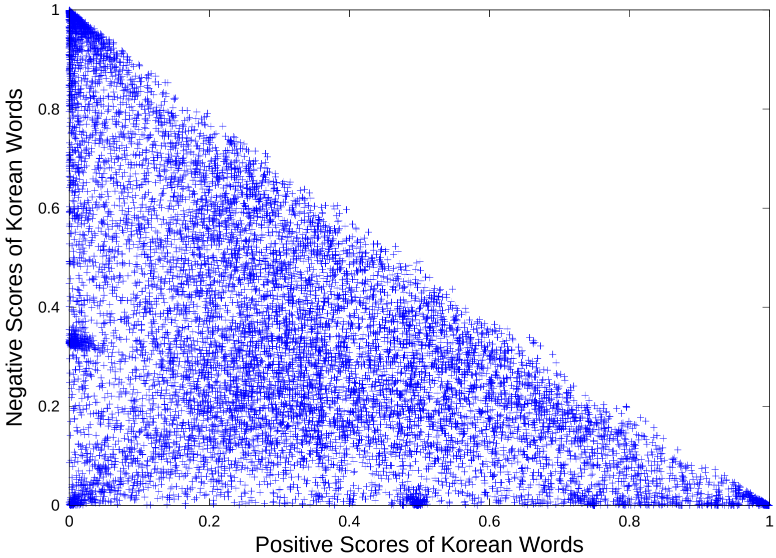

Figure 3. Positive and Negative Scores Distribution of Korean Lexicon Constructed from **OF** with **FS**.

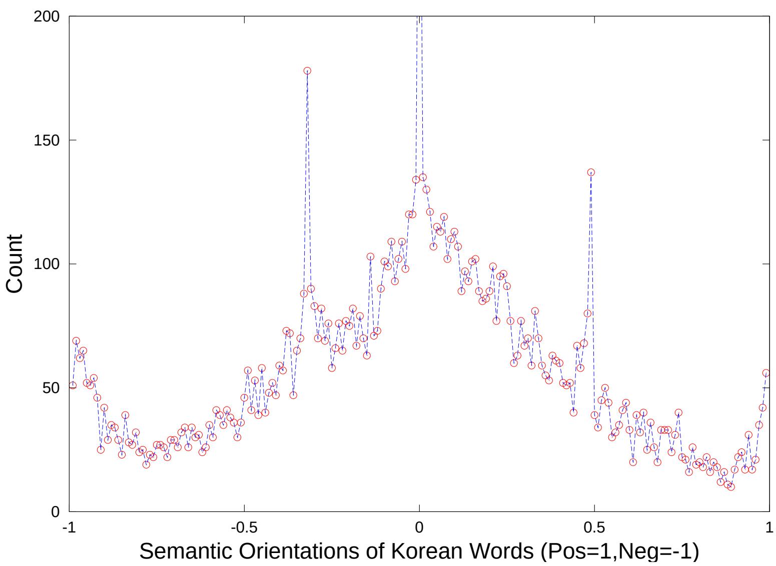

Figure 4. Semantic Orientations ($PositiveScore - NegativeScore$) of Korean Lexicon Constructed from **OF** with **FS**.

Our future work includes incorporating the word sense of the target language into the translation process and creating a sense-level lexicon, and extending to different language pairs.

Acknowledgments

This work was supported by the BK 21 Project in 2009.

References

[1] A. Andreevskaia and S. Bergler, Mining wordnet for fuzzy sentiment: Sentiment tag extraction from wordnet glosses, in *Proc. 11th Conf. European Chapter of the Association for Computational Linguistics* (*EACL'06*), 2007, pp. 209–216.

[2] S. Brin and L. Page, The anatomy of a large-scale hypertextual web search engine, in *Computer Networks and ISDN Systems*, 1998, pp. 107–117.

[3] S. Cerini, V. Compagnoni, A. Demontis, M. Formentelli, and C. Gandini, Micro-wnop: A gold standard for the evaluation of automatically compiled lexical resources for opinion mining, in *Language Resources and Linguistic Theory: Typology, Second Language Acquisition, English Linguistics*, 2007.

[4] A. Esuli and F. Sebastiani, Determining the semantic orientation of terms through gloss analysis, in *Proc. 14th ACM Int. Conf. on Information and Knowledge Management* (*CIKM'05*), 2005, pp. 617–624.

[5] A. Esuli and F. Sebastiani, Sentiwordnet: A publicly available lexical resource for opinion mining, in *Proc. 5th Conf. on Language Resources and Evaluation* (*LREC'06*), 2006, pp. 417–422.

[6] A. Esuli and F. Sebastiani, Pageranking wordnet synsets: An application to opinion mining, in *Proc. 45th Annual Meeting of the Association for Computational Linguistics* (*ACL'07*), 2007, pp. 424–431.

[7] R. Fagin, R. Kumar, M. Mahdian, D. Sivakumar, and E. Vee, Comparing and aggregating rankings with ties, in *Proc. ACM Int. Conf. on Principles of Database Systems* (*PODS'04*), 2004, pp. 47–58.

[8] V. Hatzivassiloglou and K. R. Mckeown, Predicting the semantic orientation of adjectives, in *Proc. 35th Annual Meeting of the Association for Computational Linguistics* (*ACL'97*), 1997, pp. 174–181.

[9] M. Hu and B. Liu, Mining and summarizing customer reviews, in *Proc. 10th ACM SIGKDD Int. Conf. on Knowledge Discovery and Data Mining* (*KDD'04*), 2004, pp. 168–177.

[10] Y. Hu, J. Duan, X. Chen, B. Pei, and R. Lu, A new method for sentiment classification in text retrieval. in *Proc. 2nd Int. Joint Conf. on Natural Language Processing* (*IJCNLP'05*), 2005, pp. 168–177.

[11] J. Kamps, M. Marx, R. J. Mokken, and M. D. Rijke, Using wordnet to measure semantic orientation of adjectives, in *Proc. 4th Int. Conf. on Language Resources and Evaluation* (*LREC'04*), 2004, pp. 1115–1118.

[12] H. Kanayama and T. Nasukawa, Fully automatic lexicon expansion for domain-oriented sentiment analysis, in *Proc. 2006 Conf. on Empirical Methods in Natural Language Processing* (*EMNLP'06*), 2006, pp. 355–363.

[13] S.-M. Kim and E. Hovy, Determining the sentiment of opinions, in *Proc. 20th Int. Conf. on Computational Linguistics* (*COLING'04*), 2004, pp. 1367–1373.

[14] S.-M. Kim and E. Hovy, Identifying and analyzing judgment opinions, in *Proc. Human Language Technology Conf. on the NAACL* (*HLT/NAACL'06*), 2006, pp. 200–207.

[15] J. M. Kleinberg, Authoritative sources in a hyperlinked environment, *J. ACM* **46**(5) (1999) 604–632.

[16] R. Mihalcea, Random walks on text structures, in *Proc. 7th Int. Conf. on Computational Linguistics and Intelligent Text Processing* (*CICLing'06*), 2006, pp. 249–262.

[17] R. Mihalcea, C. Banea, and J. Wiebe, Learning multilingual subjective language via cross-lingual projections, in *Proc. 45th Annual Meeting of the Association of Computational Linguistics* (*ACL'07*), 2007, pp. 976–983.

[18] P. J. Stone, D. C. Dunphy, M. S. Smith, and D. M. Ogilvie, *The General Inquirer: A Computer Approach to Content Analysis* (MIT Press, 1966).

[19] H. Takamura, T. Inui, and M. Okumura, Extracting semantic orientations of words using spin model, in *Proc. of the 43rd Annual Meeting of the Association for Computational Linguistics* (*ACL'05*), 2005, pp. 133–140.

[20] H. Takamura, T. Inui, and M. Okumura, Latent variable models for semantic orientations of phrases, in *Proc. 11th Conf. on the European Chapter of the Association for Computational Linguistics* (*EACL'06*), 2006, pp. 201–208.

[21] P. D. Turney and M. L. Littman, Measuring praise and criticism: Inference of semantic orientation from association, *ACM Trans. on Information Systems* **21**(4) (2003) 315–346.

[22] C. Whitelaw, N. Garg, and S. Argamon, Using appraisal groups for sentiment analysis, in *Proc. 14th ACM Int. Conf. on Information and Knowledge Management* (*CIKM'05*), 2005, pp. 625–631.

[23] J. Wiebe and R. Mihalcea, Word sense and subjectivity, in *Proc. 44th Annual Meeting of the Association for Computational Linguistics* (*ACL'06*), 2006, pp. 1065–1072.

[24] T. Wilson, J. Wiebe, and P. Hoffmann, Recognizing contextual polarity in phrase-level sentiment analysis, in *Proc. Conf. on Human Language Technology and Empirical Methods in Natural Language Processing* (*HLT-EMNLP'05*), 2005, pp. 347–354.

[25] H. Yu and V. Hatzivassiloglou, Towards answering opinion questions: Separating facts from opinions and identifying the polarity of opinion sentences, in *Proc. Conf. on the Empirical Methods in Natural Language Processing* (*EMNLP'03*), 2003, pp. 129–136.