

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

Jungi Kim, Hun-Young Jung, Sang-Hyob Nam, Yeha Lee, and Jong-Hyeok Lee

Division of Electrical and Computer Engineering
Pohang University of Science and Technology
San 31, Hyoja-Dong, Nam-Gu, Pohang, 790-784, Republic of Korea
{yangpa, blessing, namsang, sion, jhlee}@postech.ac.kr

Abstract. This paper proposes a method that automatically creates a subjectivity lexicon in a new language using a subjectivity 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 subjectivity lexicon in the source language as well as create a good quality lexicon in a new language.

Keywords: Opinion Analysis, Sentiment Lexicon, Lexicon Translation.

1 Introduction

With the recent pursuit of study in subjectivity analysis tasks such as sentiment classification, opinion holder extraction, and opinion summarization, much research effort has been spent on automating such tasks using various natural language processing approaches. Most previous researches, from constructing language resources [1,2,3,4,5,6,7,8] and sentiment analysis [9,10,11,12] to a variety of applications [13,14,15], have targeted English language only, and naturally many language resources in subjectivity analysis have been created in English.

While a number of languages such as Japanese, Chinese, and German are directly employed in recent studies [16,17,18], some works [19,12] have explored utilizing language resources in English to develop language resources and subjectivity analysis tools in other languages.

Motivated by the latter approach, this paper presents a method for automatically creating a subjectivity lexicon in a new language using a subjectivity 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 subjectivity lexicon using a bilingual dictionary

faces several problems [12]: 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 [11,8]. Especially, [8] 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 the evidences 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 subjectivity lexicon in Korean utilizing subjectivity lexicons in English; Korean is a relatively understudied language in subjectivity 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 subjectivity lexicons have been used in many subjectivity analysis tasks. Some lexicons are manually created [20,21,14] while others are the outcomes of the research efforts on automatically learning subjectivity from dictionary and thesaurus [3,13,9,5,12,5,7,8] or raw corpus [1,2,17].

There also has been efforts to utilize the language resources created in English for analyzing the subjectivity in other languages; although in very limited fashion, [19] are 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. [12] was the first genuine multilingual work in subjectivity analysis, in which subjectivity analysis resources developed for English are used for developing resources in Romanian, by translating the subjectivity lexicon and creating a subjectivity corpus through projection using a parallel corpus between English and Romanian and a subjectivity classifier in English. Similar to the approach in [12], our work directly translates the subjectivity lexicon in English to 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.

[8] uses a graph representation of WordNet synsets and a random-walk model to simulate the dynamics of the vertices that have similar ORPs. While [8] obtains the clues for the edges from glosses of WordNet entries, our work creates more secure and 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 Subjectivity Lexicon

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

Subjectivity 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 ~ 1.0)). We assume that our English subjectivity lexicon contains English words with POS tags and sentiment orientation with some measure of its strength (e.g. {abandon, verb, weak negative}, or {harm, verb, positive 0.0, negative 0.5, neutral 0.5}), and the Korean subjectivity 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 Subjectivity Lexicon

Translating a subjectivity lexicon into another language using a bilingual dictionary is a very 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 [12].

[12] 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 a 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 [12]. 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 the Lexicon with a Link Analysis Algorithm

Similarly to [8], our approach uses a graph built from the words with ORPs as vertices, and the relations among the words as edges connecting the vertices. While [8] used gloss of WordNet synsets to create some semantic relations among the synsets, with the hypothesis that gloss of a synset will usually contain terms belonging to synsets with similar ORPs, our approach utilizes the 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 [22].

Considering the hubness of an English vertex as its sentiment score, and a 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 [23] 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 [8], 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.

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

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.

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

In summary, refining the subjectivity lexicons in English and Korean is carried out on our two phase link analysis framework: first, running HITS with Korean words as “authorities” and English words as “hubs” to learn the authoritativeness of Korean words, and secondly, running HITS again with English words as “authorities” and Korean words 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 and negative sentiment independently, producing separate rankings of lexicons for positive and negative sentiment scores.

4 Experiments

4.1 Setup

The English lexicons we use in our experiments are the subjectivity lexicon used in OpinionFinder (**OF**) [10]¹ and SentiWordNet 1.0.1 (**SentiWN**) [7]².

¹ <http://www.cs.pitt.edu/mpqa/>

² <http://sentiwordnet.isti.cnr.it/>

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*. *Weak* subjectivity words were assigned the score of 0.5, and *Strong* words with 1.0.

SentiWN is a set of WordNet synsets with automatically assigned positive, negative, and neutral probability scores. In our experiments, each word in a synset is treated separately with the sentiment scores of the synset as its own, ignoring the synonym information provided by WordNet synsets.

We use an online bilingual dictionary provided by a portal website³. 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** ($\frac{1}{|\text{Number of Vertices}|}$), and every combination of English lexicons (**OF** and **SentiWN**) 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 [8], which uses a Micro-WNOp corpus [24]⁴ as a gold standard and the *p-normalized Kendall τ distance* (τ_p) [25] 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 $\{o_1 \dots o_n\}$, 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

$$\tau_p = \frac{n_d + \frac{1}{2} \times n_u}{Z} \quad (3)$$

³ <http://endic.naver.com/>

⁴ <http://www.unipv.it/wnop/>

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 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

The experimental results show our proposed translation heuristics worked as we had expected: heuristics that translate only reliable words tend to have smaller τ_p and a lower number of translated words, while heuristics that translate more words have a bigger translated τ_p . Direct evaluation of **OF** lexicon results in poor scores (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.

The proposed framework with two link analysis models has a compensating effect in each phase that the lexicons mutually complement each other in turn (Table 2 and Table 3). The quality of the lexicons in every approach has shown to range from slightly negative (+1.29%) to exceptional (-41.3%).

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, 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, 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

6 Conclusions

This paper investigated the feasibility of exploiting a subjectivity lexicon in one language to developing a subjectivity 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 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.

Our future work includes incorporating word sense into the translation process and extending to different language pairs.

Acknowledgments

This work was supported in part by MKE & IITA through IT Leading R&D Support Project and also in part by the BK 21 Project in 2008.

References

1. Hatzivassiloglou, V., Mckeown, K.R.: Predicting the semantic orientation of adjectives. In: Proceedings of the 35th Annual Meeting of the Association for Computational Linguistics (ACL 1997), madrid, ES, pp. 174–181 (1997)
2. Turney, P.D., Littman, M.L.: Measuring praise and criticism: Inference of semantic orientation from association. *ACM Transactions on Information Systems* 21(4), 315–346 (2003)
3. Kamps, J., Marx, M., Mokken, R.J., Rijke, M.D.: Using wordnet to measure semantic orientation of adjectives. In: Proceedings of the 4th International Conference on Language Resources and Evaluation (LREC 2004), Lisbon, PT, pp. 1115–1118 (2004)
4. Takamura, H., Inui, T., Okumura, M.: Extracting semantic orientations of words using spin model. In: Proceedings of the 43rd Annual Meeting of the Association for Computational Linguistics (ACL 2005), Ann Arbor, USA, pp. 133–140 (2005)
5. Esuli, A., Sebastiani, F.: Determining the semantic orientation of terms through gloss analysis. In: Proceedings of the 14th ACM International Conference on Information and Knowledge Management (CIKM 2005), Bremen, DE, pp. 617–624 (2005)
6. Andreevskaia, A., Bergler, S.: Mining wordnet for fuzzy sentiment: Sentiment tag extraction from wordnet glosses. In: Proceedings of the 11th Conference of the European Chapter of the Association for Computational Linguistics (EACL 2006), Trento, IT, pp. 209–216 (2007)
7. Esuli, A., Sebastiani, F.: Sentiwordnet: A publicly available lexical resource for opinion mining. In: Proceedings of the 5th Conference on Language Resources and Evaluation (LREC 2006), Geneva, IT, pp. 417–422 (2006)
8. Esuli, A., Sebastiani, F.: Pageranking wordnet synsets: An application to opinion mining. In: Proceedings of the 45th Annual Meeting of the Association for Computational Linguistics (ACL 2007), Prague, CZ, pp. 424–431 (2007)
9. Kim, S.M., Hovy, E.: Determining the sentiment of opinions. In: Proceedings of 20th International Conference on Computational Linguistics (COLING 2004), Geneva, CH, pp. 1367–1373 (2004)
10. Wilson, T., Wiebe, J., Hoffmann, P.: Recognizing contextual polarity in phrase-level sentiment analysis. In: Proceedings of the Conference on Human Language Technology and Empirical Methods in Natural Language Processing (HLT-EMNLP 2005), Vancouver, CA, pp. 347–354 (2005)

11. Mihalcea, R.: Random walks on text structures. In: Gelbukh, A. (ed.) *CICLing 2006*. LNCS, vol. 3878, pp. 249–262. Springer, Heidelberg (2006)
12. Mihalcea, R., Banea, C., Wiebe, J.: Learning multilingual subjective language via cross-lingual projections. In: *Proceedings of the 45th Annual Meeting of the Association of Computational Linguistics (ACL 2007)*, Prague, CZ, pp. 976–983 (2007)
13. Hu, M., Liu, B.: Mining and summarizing customer reviews. In: *Proceedings of the 10th ACM SIGKDD international conference on Knowledge discovery and data mining (KDD 2004)*, New York, USA, pp. 168–177 (2004)
14. Wiebe, J., Mihalcea, R.: Word sense and subjectivity. In: *Proceedings of the 44th Annual Meeting of the Association for Computational Linguistics (ACL 2006)*, Sydney, AU, pp. 1065–1072 (2006)
15. Yu, H., Hatzivassiloglou, V.: Towards answering opinion questions: Separating facts from opinions and identifying the polarity of opinion sentences. In: *Proceedings of 2003 Conference on the Empirical Methods in Natural Language Processing (EMNLP 2003)*, Sapporo, JP, pp. 129–136 (2003)
16. Takamura, H., Inui, T., Okumura, M.: Latent variable models for semantic orientations of phrases. In: *Proceedings of the 11th Conference of the European Chapter of the Association for Computational Linguistics (EACL 2006)*, Trento, IT, pp. 201–208 (2006)
17. Kanayama, H., Nasukawa, T.: Fully automatic lexicon expansion for domain-oriented sentiment analysis. In: *Proceedings of the 2006 Conference on Empirical Methods in Natural Language Processing (EMNLP 2006)*, Sydney, AU, pp. 355–363 (2006)
18. Hu, Y., Duan, J., Chen, X., Pei, B., Lu, R.: A new method for sentiment classification in text retrieval. In: Dale, R., Wong, K.-F., Su, J., Kwong, O.Y. (eds.) *IJCNLP 2005*. LNCS, vol. 3651, pp. 1–9. Springer, Heidelberg (2005)
19. Kim, S.M., Hovy, E.: Identifying and analyzing judgment opinions. In: *Proceedings of the Human Language Technology Conference of the NAACL (HLT/NAACL 2006)*, New York, USA, pp. 200–207 (2006)
20. Stone, P.J., Dunphy, D.C., Smith, M.S., Ogilvie, D.M.: *The General Inquirer: A Computer Approach to Content Analysis*. MIT Press, Cambridge (1966)
21. Whitelaw, C., Garg, N., Argamon, S.: Using appraisal groups for sentiment analysis. In: *Proceedings of the 14th ACM international conference on Information and knowledge management (CIKM 2005)*, Bremen, DE, pp. 625–631 (2005)
22. Kleinberg, J.M.: Authoritative sources in a hyperlinked environment. *J. ACM* 46(5), 604–632 (1999)
23. Brin, S., Page, L.: The anatomy of a large-scale hypertextual web search engine. In: *Computer Networks and ISDN Systems*, pp. 107–117 (1998)
24. Cerini, S., Compagnoni, V., Demontis, A., Formentelli, M., Gandini, C.: Micrownop: 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*, Milano, IT (2007)
25. Fagin, R., Kumar, R., Mahdian, M., Sivakumar, D., Vee, E.: Comparing and aggregating rankings with ties. In: *Proceedings of the ACM International Conference on Principles of Database Systems (PODS 2004)*, Paris, FR, pp. 47–58 (2004)