Sentence-Level Contextual Opinion Retrieval

Sylvester Olubolu Orimaye
School of IT, Monash University, Malaysia
sylvester.orimaye@infotech.monash.edu.my
Supervised by: Saadat M. Alhashmi and Siew Eu-Gene
{saadat.m.alhashmi@infotech.monash.edu.my and siew.eugene@buseco.monash.edu.my}

ABSTRACT
Existing opinion retrieval techniques do not provide context-dependent relevant results. Most of the approaches used by state-of-the-art techniques are based on frequency of query terms, such that all documents containing query terms are retrieved, regardless of contextual relevance to the intent of the human seeking the opinion. However, in a particular opinionated document, words could occur in different contexts, yet meet the frequency attached to a certain opinion threshold, thus explicitly creating a bias in overall opinion retrieved. In this paper we propose a *sentence-level contextual model* for opinion retrieval using *grammatical tree derivations* and *approval voting* mechanism. Model evaluation performed between our contextual model, BM25, and language model shows that the model can be effective for contextual opinion retrieval such as faceted opinion retrieval.

Categories and Subject Descriptors
H.3.3 [Information Search and Retrieval]: *Information filtering, Query formulation, Retrieval models, Search process*

General Terms
Algorithms, Design, Experimentation, Performance

Keywords
Contextual opinion, sentence level, grammatical tree derivation, approval voting

1. INTRODUCTION
Understanding subjective contributions is a complex process, especially when it involves contributors with diversified styles and knowledge of making contributions. The arrival of web 2.0 interactive medias such as web logs (blogs) creates an avenue whereby people display their diversities in terms of style and knowledge when making subjective contributions. To make matter worse, when people contribute, they do so in their own ways; they write with self perception often conveyed by psychological inference; they present contributions with different language styles, and even play upon words.

The summary of the factors highlighted above is not far from *sentiment* or *opinion,* especially within a complex interactive space such as collection of blogs called blogosphere. In fact, there is reasonable amount of interest in mining different kind of knowledge from such huge human-oriented contribution space[29]. Organizations wish to know relevant opinions towards products or services rendered [26]; individuals wish to retrieve relevant information for personal use, and the need for understanding human behaviors through trends is becoming more apparent and necessary[1]. Therefore, we believe a technique that can take into account the need for contextual understanding of sentences would be appropriate for opinion retrieval tasks.

Many opinion retrieval systems avoid computational models that treats opinion as a process of language understanding [23], in fact they are mostly probabilistic or statistical [38]. Existing opinion retrieval models presume words that belong to specific targets [39], hence count the occurrence of such words in a given document to meet certain opinion threshold. We argue that in a particular document, words could occur in different contexts, yet meet the frequency attached to a certain opinion threshold, thus explicitly creates bias in overall opinion. For example, these two sentences, *the fight for academic success* and *I will fight you to finish*, have regular occurrence of the word "fight", which could imply *violence* opinion upon certain threshold, whereas, the word "fight" has appeared in two different contexts. More explanatory examples can be found in [10] on page 345.

Towards addressing this issue, we believe existing models and techniques in computational linguistics can be used, such as *context-sensitive* Combinatory Categorial Grammar (CCG)[18]. We believe opinion retrieval should benefit more from the impact and aggregation of textual understanding in context [32], rather than frequency of terms alone.

2. RELATED WORK
Existing opinion retrieval techniques do not provide context-dependent relevant results. Most of the approaches used by state-of-the-art techniques are based on frequency of query terms, such that all documents containing query terms are retrieved, regardless of contextual relevance to the intent of the human seeking the opinion[12,38]. Little have also been done on contextual understanding of sentences used in opinionated contributions [26, 32, 40]. We believe frequency of query terms can not imply subjectivity alone, without establishing the context at which a term must be frequent [23, 32]. In this work, we propose a context-dependent approach that can solve the *non-contextual relevance* problem common to existing opinion retrieval systems.

[1] http://www.nytimes.com/2009/08/24/technology/internet/24emotion.html

WWW 2011, March 28–April 1, 2011, Hyderabad, India.
ACM 978-1-4503-0637-9/11/03.

Some research works have treated opinion retrieval as *text classification* (TC) problem. The TC technique enables the introduction of both traditional supervised and unsupervised ML techniques with reasonable efficiencies recorded in most research works[29]. However, determining an effective labeling strategy for contextual opinion is still an open problem [27]. According to Pang and Lee[32] subjectivity can be interpreted differently in some cases as they come with different levels of challenges, such as “agreement” or “disagreement”[5] and “winner” or “looser”[4]. For example in [27], the use of *ironical* words or phrases and inverted polarity in opinionated documents led to lower precision for positive opinions with 77% accuracy. As a result of this, the choice of features for determining context in opinionated documents is still a problem factor [1, 9, 32].

Lexicon-based opinion generation has also been used for opinion retrieval, and many research works have combined both lexicon-based and TC for retrieving opinions [20]. However, this approach is domain specific [9, 31], and there have been instances of conflict between opinions and lexicon. For example, a positive comment about a product could have a negative meaning in the lexicon [22]. We believe that analysis of independent words to show meaning may not necessarily reflect opinion except when words are processed together in a sentence.

Probabilistic models have also been used to retrieve and rank opinions from documents[38], and some level of success have been recorded by the technique. However, the probabilistic approach is full of probabilistic assumptions based on frequency of query terms [16, 23, 32]. For example, in [38], a probabilistic model is used to determine proximity of words to the query terms, which may not necessarily reflect the context at which user information is required. Whereas, most opinionated documents have sentences that describe different opinions even within the same paragraph [5, 32], thus opinionated words occurring at proximity to query terms might have occurred in different context.

The *language model* approach has also been used for opinion retrieval [33, 35]. Most language models go as far as word level processing; sentence level processing; and paragraph level or bag-of-words. Some language models combine probabilistic techniques for efficient ranking of opinionated documents[14]. A common limitation of this approach is the estimation of model parameters [19], that is, an effective way to model a document to give a higher probability of relevance to user query is still very challenging. For example, [19] addressed the parameter tuning or optimization problem in higher order language models by incorporating different component models as features. A basic approach is to perform smoothing at varying levels for effective document retrieval. However, the methodology proposed in [19] leads to having higher number of model parameters, whereby all of such parameters would also require optimization at different levels.

3. CONTEXTUAL DERIVATION MODEL

In this work, we look at underlying meaning of each sentence and then compare with underlying meaning of the given query or opinion target. We believe sentences form the base of the overall opinion being expressed in a document, and opinionated information is better represented in sentences than individual query words. In our contextual model, we assume an ideal user would give query in the context at which the information is needed. We categorize the contextual model into three stages as we discourse in the following sections.

3.1 Contextual Query Formulation Model

We define opinion target T as a set of query terms Q:

$$T(Q) = Q(t) \leftarrow C_x, \tag{3.1}$$

where $Q(t)$ is a set of query terms composed of one or more words in its original sequence, C_x is the context to be derived from sequence $Q(t)$. Therefore, $Q(t) \leftarrow C_x$, shows the context of opinion target $T(Q)$ for each document D in the collection. We expand the original opinion target $T(Q)$ by deriving *related opinion targets* from $T(Q)$ through *contextual query formulation or substitution*[21] as shown below:

$$T(\partial) \xrightarrow[T(Q)]{yields} (T_1Q, T_2Q, \dots T_mQ). \tag{3.2}$$

Therefore, we can rewrite equation 3.1 to represent each of the opinion target contained in equation 3.2 as follows:

$$\forall T(Q) \in T(\partial) \mid T_m(Q) = Q(t) \leftarrow C_x . \tag{3.3}$$

Note that the composition of the above related opinion targets becomes a *reference profile*[2] to the original opinion target.

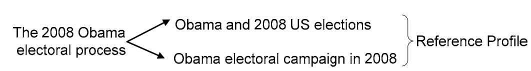

Figure 1. Similar opinion targets from a single original user query

3.2 Sentential Context Derivation Process

We identify initial document collection using BM25[37]. Using CCG [18], we understand the meaning of each sentence in order to derive context of a given sentence as $C_x(S)$. Upon sentential context derivation, we are particularly interested in sentences that have the same context with the opinion target. Also, we define *reference profile* $P = \{T(\partial)\}$ as a set of similar opinion target formed above. We shall show how P can be used in the contextual matching process using approval voting technique[7, 28].

We have two important properties for our model. The first property is the *tree derivation* ∂_t which shows the *global property* of phrase structure tree for S using a selected CCG rule r. The second property is the *semantic properties* ∂_{sm} derived from S which shows the grammatical realization of content words. We define the first property as follows:

$$\partial_t(S) = \partial_t\,(S \cap (r \in R)) \tag{3.4}$$

where r is a rule element in a set of CCG rules R, $S \cap (r \in R)$ is the appropriate rule that matches the complexity of S and automatically selected from set of CCG rules R, and ∂_t is the phrase tree derivation performed on S. The second property is defined as follows:

$$\partial_{sm}(S) = \partial_{sm}\,(w \in S), \tag{3.5}$$

[2] The reference profile defines a broader interpretation of the query

where w is a word element of sentence S, ∂_{sm} is the grammatical realization of each content word, and $\partial_{sm}(S)$ is semantic property of a set of grammatical realizations of each content word in *S*.

Therefore, the context flagged by *S* can be defined as follows:

$$C_x(S) = \partial_t(S) \cap \partial_{sm}(S) . \qquad (3.6)$$

In equation 3.6 above, a relationship is established between global property $\partial_t(S)$ and its corresponding semantic annotations; this relationship forms the *context* that defines the underlying meaning of the sentence. Now that we can derive sentential context, recall that our aim is to identify contextual sentences that match the opinion target; we therefore find the context C_x for query formulation $Q(t)$ by replacing *S* with $Q(t)$ as shown below:

$$C_x(Q(t)) = \partial_t(Q(t)) \cap \partial_{sm}(Q(t)), \qquad (3.7)$$

where $\partial_t(Q(t))$ and $\partial_{sm}(Q(t))$ are derived from equation 3.4 and 3.5 respectively by substituting $Q(t)$ for *S*. Therefore, a contribution to user information need is established when each $C_x(S)$ is identical to $C_x(Q(t)) \in T(\partial)$, and $C_x(Q(t))$ can be seen as a representation of original opinion target or its *reference profile P* defined in 3.2 above. This process is shown below:

$$T_m(Q) \leftarrow C_x(Q(t)), \qquad (3.8)$$

$$C_x(S).P \equiv C_x(Q(t)). \qquad (3.9)$$

3.3 Contextual Matching by Voting

In this process we use approval voting [7, 24, 28] technique in order to differentiate between the size of relevant opinions and non-relevant opinions present in the document. Each contextual sentence that can describe the opinion target is an *implicit vote* for the opinion target[28]. The size of the opinion relevant to the opinion target can be defined as the respective accumulation of contextual sentences $C_x(S).P \equiv C_x(Q(t))$ shown below:

$$||O_{pc}|| = \sum C_x(S).P \equiv C_x(Q(t)). \qquad (3.10)$$

However, our research interest is not just to know sentences that support the opinion target only, as opinionated documents can have diverse opinions[26, 32], we are also interested in sentences that describe other context too, which is why the voting technique comes in. A very important question in every voting technique is *"who the candidates are?"* Therefore, using the *sentential context derivation process* described in section 3.2, we are able to recursively define candidates from each opinionated document, to include the opinion target, and other unique contexts different from the opinion target.

$$C_x(S) \neq C_x(Q(t)), \qquad (3.11)$$

$$C_x(S) \leftarrow C_x'. \qquad (3.12)$$

Equation 3.12 above describes a situation where a particular sentence does not support the opinion target; therefore a contextual sentence C_x' can be seen as a potential candidate y' in a set of contesting candidates C_d. Recall that each candidate has a profile *P* described as a broader representation of the context flagged by the particular candidate. This process is shown below:

$$C_x' \leftarrow y', \qquad (3.13)$$

$$C_d = (y_0 P, y'_1 P, y'_2 P \dots y'_n P) \qquad (3.14)$$

In equation 3.14 above, C_d is a set of candidates, y_0 is the context of the opinion target $T(Q)$ as the first candidate, subsequent contexts are then added to the set. For the approval voting process, we assume a *leader rule*[3] strategy with element of *sincerity* and *admissibility* [24], and then define a voter's choice, that is, when a sentence can vote for one or more candidates.

$$V_{choice} = \forall y \in C_d \,|\, P(y_{i=0}) \cap P(y_n) \qquad (3.15)$$

$$|| O_{pc}' || = \sum C_x'(s) \qquad (3.16)$$

Therefore, equation 3.10 is the size or weight of the opinion $||O_{pc}||$ relevant to the opinion target from a given document and equation 3.16 is the size or weight of non-opinion $|| O_{pc}' ||$ derived from the same given document. A document's relevance R_{doc} to user's intent is determined when the size of relevant opinion is **greater than** the size non-relevant opinion clouds derived from the document. Opinion weights for relevant documents can then be used to rank documents hierarchically. Relevance of a document can be defined below:

$$R_{doc} = \forall \; || O_{pc}' || \quad (||O_{pc}|| > || O_{pc}' ||). \qquad (3.17)$$

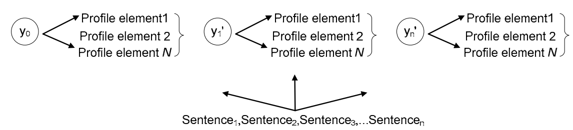

Figure 2. Contextual voting process using candidates' profiles

4. Evaluation using AIC Model Selection

We used Akaike Information Criterion (AIC)[4] to perform model selection between BM25[5] [37], language models [19], and our contextual model. By this, we are able to understand the approximating capability of our contextual model by simply computing and comparing the AIC values for each of the candidate models, the model with the *least* AIC value is believed to be the best fitted model for the opinion retrieval task[3, 6, 8]. Moreover, AIC has strong theoretical foundations built over Kullback-Leibler distance[30], for selecting appropriate approximating model for statistical inference from many types of empirical data [2-3, 6, 8].

We show AIC's criteria for estimating the expected relative distance between models as follows:

$$\textbf{AIC} = -2\log(L(\hat{\theta}\,|y)) + 2K. \qquad (4.1)$$

4.1 Baselines for AIC Model Selection

In equation 4.1, the numerical value of the log likelihood at its maximum point is represented by $\log(L(\hat{\theta}|y))$ where K is the number of estimable parameters in each participating model[8]. We set the *optimized parameters* and *values* for BM25 and Bigram Language Model (LM) in [19] as our baselines to be compared with our contextual model using AIC values. Arguably, we assume

[3] A behavioral rule with rational foundation [24].

[4] Selects best approximating model from finite samples by minimizing the K-L information through parameters estimation [2-3, 6, 8].

[5] BM25 has been used in [11, 17, 34, 38].

Normalized Discounted Cumulative Gain (NDGC)[6] values@1 reported in [19] as our *estimated* likelihood values for the two models above.

4.2 Likelihood Ratio and Parameter Estimation for the Contextual Model

For our contextual model we compute the opinion score based on the Likelihood Ratio (LR) method presented in [13]. The major difference between this method and our approach is that, rather than "frequency of terms" we used *frequency of contextual opinionated sentence* as described in section 3.2, that is, the frequency of sentences that support the context of the opinion target. Therefore, we represent set of documents with opinions as Doc_{op} and relevant documents with R_{doc} and then compute the frequency of *relevant sentence* in a document as follows:

$$p(C_x|Doc_{op}) = \frac{\sum d \in Doc_{op}\, N(C_x, d)}{\sum d \in Doc_{op}\, |S|}. \quad (4.2)$$

In equation 4.2 above, $N(C_x, d)$ is the number of *contextual relevant sentence* C_x in document d, $|S|$ is the total number of sentences in document *d*, and $p(C_x|R_{doc})$ can be computed as the ratio of number of sentences with at least one of the terms (see section 3.2) of the opinion target to the total number of sentences. Therefore, the LR of *contextual relevant sentence* can be seen as the ratio of $p(C_x|R_{doc})$ to $p(C_x|Doc_{op})$ as shown below:

$$Opinion_{Score}C_x = \frac{p(C_x|Doc_{op})}{p(C_x|R_{doc})}. \quad (4.3)$$

Before we compute the contextual model's *estimated* value of the opinion score for AIC, we wish to quickly state that the contextual model can be interpreted using the Hidden Markov Models (HMM) [36], whereby we derived parameter γ as the number of appropriate *hidden* states that generates the best *probability distribution* for our contextual model. For example, the appropriate number of *related opinion targets* to be generated such as described in section 3.1.

Since the likelihood values for our baselines models were derived for one page (i.e. NDGC@1) [19], for an estimated $Opinion_{Score}C_x$, we manually observed the *first* blog result[7] shown by *Google blogs*[8] using "The 2008 Obama electoral process" as query, and we set our parameter $\gamma = 1$. For example, we assume section 3.1 has generated *one* related opinion target such as "Obama and 2008 US elections", and then observe the total number of sentences that can support or vote for the above two targets (see section 3.3). From our observation $N(C_x, d)$ is equal to 3 and $|S|$ is equal to 25 for the query given above. In this particular blog, we could however observe other targets such as "Support for US elections" attract more sentences or *votes* other than the initial query. With $N(C_x, d)$ equal to 8 for "Support for US elections", obviously our contextual model would not have retrieved this blog for the query "The 2008 Obama electoral process", as it will consider other context that outweighs the query, such as "Support for US elections".

We assume the latter query (i.e."Support for US elections") as the likely target for the blog above, therefore we used $N(C_x, d)$ = 8, $|S|$ = 25, and computed $p(C_x|R_{doc})$ using number of sentences with at least one unique *query term or synonyms* equals 13, thus $p(C_x|R_{doc})$ would be 13/25 = 0.52. Using equation 4.2, $p(C_x|Doc_{op})$ is computed as 8/25 = 0.32, and using equation 4.3, $Opinion_{Score}C_x$ is computed as 0.32/0.52 = 0.62.

Table 1. Parameters and likelihood values for AIC candidate models[19].

Model	**Parameters**	**No of parameters**	**Likelihood**
BM25	$k1 = 1.0$, $b = 0.5$	2	0.2556
Bigram LM	$\lambda 1 = 0.24$, $\lambda 2 = 0.29$, $\lambda 3 = 0.94$, $\mu 1 = 1800$, $\mu 2 = 400$, $\mu 3 = 792$, $\mu 4 = 900$	7	0.2630
Contextual Model	$\gamma = 1$	1	0.62

4.3 Results

Table 2. AIC values for the candidate models derived based on [6, 8] with the appropriate formulas shown in footnotes on this page. The least AIC or AIC_c and the highest w_i values respectively denote the best approximating model[6, 8].

Model	K	AIC[9]	AIC_c[10]	Δ_i[11]	w_i[12]
Contextual Model	1	**2.956072**	**0.956072**	0	**0.867519**
BM25	2	6.728283	4.728283	3.772211	0.131569
Bigram LM	7	16.6712	14.6712	13.71513	0.000912

Using the above table we can compute the evidence ratio ($\frac{w_i}{w_j} \equiv \frac{w_{min}}{w_j}$) as a rule of thumb [3, 6, 8], which shows the best approximating model, given the data and the candidate models. The strength of evidence for our *sentence-level contextual model* against BM25 and then Bigram LM can be computed as 0.867619/0.131569 ≈ 6.6, and 0.867619/0.000912 ≈ 951.5 respectively. Obviously, the evidence for *sentence-level contextual model* is very strong against Bigram LM while BM25 seems reasonable. In addition, candidate model with w_i less than 10% of the highest cannot be considered as good approximating model[3, 6, 8]. In this case, 10% of the highest w_i (i.e. 0.867519 * 0.10**)** is 0.086752; therefore since 0.867519 is greater than 0.086752, we can

[6] Used for performance measurement of models in IR[19].

[7] http://www.american-election.com/2010/10/05/michelle-obama-i-need-your-help/

[8] http://blogsearch.google.com/

[9] Formula is shown in equation 4.1

[10] AIC requires bias-adjustment if n/K < 40 [8], therefore, $\text{AIC}_c = \textbf{AIC} + 2 * \text{K} + (2 * \text{K} * (\text{K} + 1)) \;/\; (\text{n} - \text{K} - 1)$, where n = number of observations and K = number of parameters, in our case we have 1 observation and K is 1, therefore n/K < 40.

[11] Differences between AIC, that is, $\Delta_i = \text{AIC}_i - \text{AIC}_{min}$ [8].

[12] Akaike weight is the normalized relative likelihood, that is, $w_i = exp(-0.5\,\Delta_i) / \sum_{r=1}^{R} exp(-0.5\,\Delta_r)$, where the denominator is the sum of all the relative likelihoods of the candidate models[8].

conclude that *sentence-level contextual model* is most likely to be a good approximating model for opinion retrieval task.

5. CONCLUSION

In this paper, we introduced and successfully evaluated a contextual sentence-level contextual model for effective opinion retrieval task. The technique considered in this model is based on NLP technique (i.e. CCG) which gives a clear representation of meaning and context for each sentence. Our argument is that, rather than word counts or probabilistic processes, contextual meaning of opinionated sentences can be derived. Using approval voting technique, the aggregates of votes show the size of opinion for each contextual candidate. Contextual candidate with the largest opinion size best describes the overall opinion flagged by each document. Relevant opinionated document can be ranked hierarchically using their respective relevant opinion sizes. For future work, we plan to implement our model on TREC blog 2009 data set for detecting faceted opinions and also revisit our model to include multi-sentence contextual dependencies.

6. REFERENCES

[1] Abbasi, A., Chen, H. and Salem, A. (2008). Sentiment analysis in multiple languages: Feature selection for opinion classification. In ACM Trans. Inf. Syst.,1-34.

[2] Akaike, H. (1987). Factor analysis and AIC. In Psychometrika 317-332.

[3] Akaike, H. (1981). Likelihood of a model and information criteria. In Econometrics, 3-14.

[4] Tumasjan, A., Sprenger, T.O., Sandner, P.G., and Welpe, I.M. (2010). Predicting Elections with Twitter:What 140 Characters Reveal about Political Sentiment. In WSM '10.

[5] Bermingham, A. and Smeaton, A.F. (2009). A study of inter-annotator agreement for opinion retrieval. In SIGIR, 784-785.

[6] Bozdogan, H. (1987). Model Selection and Akaike's Information Criterion (AIC): The general theory and its analytical extentions. In Psychometrika, 345-370.

[7] Fishburn, P.C. and Brams, S.J. (2010). Approval Voting, Condorcet's Principle, and Runoff Elections. In Public Choice '10, 89-114.

[8] Burnham, K.P., Anderson, D.R. (2002). Model Selection and Multimodel Inference. Springer-Verlag New York, Inc.,

[9] Choi, Y., Kim, Y. and Myaeng, S-H. (2009). Domain-specific sentiment analysis using contextual feature generation. In CIKM '09, 37-44.

[10] Chung, C. and Pennebaker, J. (2009). The Psychological Functions of Function Words. In Social Communication, 343 - 359.

[11] Macdonald, C, Ounis, I., and Soboroff, I. (2009). Overview of the TREC2009 Blog Track. In TREC '09.

[12] Hannah, D., Macdonald, C., Peng, J., He, B., and Ounis, I. (2007). University of Glasgow at TREC 2007: Experiments in Blog and Enterprise Tracks with Terrier. In TREC '07.

[13] Gerani, S., Carman, M.J., and Crestani, F. (2009). Investigating Learning Approaches for Blog Post Opinion Retrieval. In ECIR '09. 313-324.

[14] Hiemstra, D., (2000) Using language models for information retrieval. PhD Thesis, Centre for Telematics and Information Technology, The Netherlands.

[15] Hiemstra, D., Robertson, S., and Zaragoza, H. (2004). Parsimonious language models for information retrieval. In SIGIR '04, 178-185.

[16] Huang, X. and Croft, W.B. (2009). A unified relevance model for opinion retrieval. In CIKM '09, 947-956.

[17] Ounis, I., Macdonald C., and Soboroff, I. (2008). Overview of the TREC2008 Blog Track. In TREC '08.

[18] Hockenmaier, J. and Steedman, M. (2007). CCGbank: A Corpus of CCG Derivations and Dependency Structures Extracted from the Penn Treebank. Computational Linguistics 33(3), 355-396.

[19] Javanmardi, S., Gao, J., and Wang, K. (2010). Optimizing two stage bigram language models for IR. In WWW '10, 1125-1126.

[20] Bollen, J., Pepe, A., and Mao, H. (2010). Modeling public mood and emotion:Twitter sentiment and socioeconomic phenomena. In WWW ' 10.

[21] Jones, R., Rey, B., Madani, O., Greiner, W. (2006). Generating query substitutions. In WWW '06, 387-396.

[22] Kanayama, H. and Nasukawa, T. (2006). Fully automatic lexicon expansion for domain-oriented sentiment analysis. In EMNLP '06, 355-363.

[23] Krahmer, E. (2010). What Computational Linguists Can Learn from Psychologists (and Vice Versa). In ACL '10, 285-294.

[24] Laslier, J.-F. and Maniquet, F. (2010). Classical Electoral Competition Under Approval Voting. In Handbook on Approval Voting, 415-429.

[25] Kullback, S. and Leibler, R.A. (1951). On Information and Sufficiency. In The Annals of Mathematical Statistics, 79-86.

[26] Liu, B. (2010). Sentiment Analysis and Subjectivity. In HNLP.

[27] Sarmento, L., Carvalho, P., Silva, M.J., and Eugénio de Oliveira. (2009). Automatic creation of a reference corpus for political opinion mining in user-generated content. In CIKM '09, 29-36.

[28] Macdonald, C. (2009). The Voting Model for People Search. PhD Thesis, University of Glasgow.

[29] Macdonald, C. (2010). Blog track research at TREC. In SIGIR '10, 58-75.

[30] Nederhof, M.-J. and Satta, G. (2004). Kullback-Leibler distance between probabilistic context-free grammars and probabilistic finite automata. In ACL '04, 71.

[31] Pan, S.J., (2010). Cross-domain sentiment classification via spectral feature alignment. In WWW '10,751-760.

[32] Pang, B. and Lee, L. (2008). Opinion Mining and Sentiment Analysis. In Found. Trends Inf. Retr,1-135.

[33] Jiang, P., Yang, Q., Zhang, C, and Niu, Z. (2009). BIT at TREC 2009 Faceted Blog Distillation Task. In TREC '09.

[34] Pickens, J. and MacFarlane, A. (2006). Term context models for information retrieval. In CIKM '06, 559-566.

[35] Ponte, J.M. and Croft, W.B. (1998). A language modeling approach to information retrieval. In SIGIR, 275-281.

[36] Rabiner, L.R. (1989). A Tutorial on Hidden Markov Models and Selected Applications. In Speech Recognition, IEEE, 257-286.

[37] Robertson, S., and Zaragoza, H. (2009). The Probabilistic Relevance Framework: BM25 and Beyond. In Found., and Trends in Information Retrieval, 333–389.

[38] Gerani, S., Carman, M.J., and Crestani, F. (2010). Proximity-Based Opinion Retrieval. In SIGIR '10, 978.

[39] Tausczik, Y.R., and Pennebaker, J.W. (2010). The Psychological Meaning of Words: LIWC and Computerized Text Analysis Methods. In Journal of Language and Social Psychology, 24-54.

[40] Wilson, T., Wiebe, J., and Hoffmann, P. (2009). Recognizing Contextual Polarity: an exploration of features for phrase-level sentiment analysis. Computational Linguistics 35(3).