TMSP: Topic Guided Manifold Ranking with Sink Points for Guided Summarization

Pan Du, Jin Zhang, Jiafeng Guo, Xueqi Cheng
Institute of Computing Technology
Chinese Academy of Sciences
Beijing 100190, China
{dupan,zhangjin,guojiafeng}@software.ict.ac.cn
cxq@ict.ac.cn

Abstract

Guided summarization is an extension of query-focused multi- document summarization. We proposed a novel ranking algorithm, *Topic Guided Manifold Ranking with Sink Points (TMSP)* for guided summarization tasks of TAC2010. TMSP is a topic extended version of *Manifold Ranking with Sink Points (MRSP)*, which handles the Update Summarization tasks of TAC2009 well. We adopt the TMSP and MRSP methods to guided summarization this year. The evaluation results show that our approaches are both competitive in practice.

1 Introduction

Guided summarization[1] task is to write a 100-word summary of a set of 10 newswire articles for a given topic, where the topic falls into a predefined category. Given a list of important aspects for each category, the summary must cover all these aspects if the information can be found in the documents. The summaries may also contain other information relevant to the topic. Besides, guided summarization also demands an update summary, similar to the update summarization[2] in TAC2009. Update summarization aims at generating summaries assuming the user has read some articles before. Specifically, given the topic, the task is to write two summaries, one for document set A and the other for document set B, that address the information need expressed in the corresponding topic statement. The summary for document set A is a query-focused multi-document summary. The update summary for document set B is also query-focused multi-document one but should be written under the assumption that the user of the summary has already read the documents in document set A. Each summary should be well-organized, in English, using complete sentences. Each summary can be no longer than 100 words.

As an effective and concise approach of helping users to catch the main points, document summarization has attracted much attention since the original work by Luhn [16] et al. A number of researchers have done good work in multi-document summarization (MDS). Recently, there emerged two novel demands for summarization. One is the aspect-specific requirement, the other is time-dependent requirement. A user expects the summary to contain information specific to the particular category of the event. Meanwhile, new information is created as the events develop. A user also wants the summary to contain mainly novel information, to save time. However, much of current work has focused on the specified static document collection without attempting to capture the changes over time or trying to give the aspect-based information. The classic problem of summarization is to take an information source, extract content from it, and present the most important content to the user in a condensed form

[1] http://www.nist.gov/tac/2010/Summarization/
[2] http://www.nist.gov/tac/2009/Summarization/

and in a manner sensitive to the user's or application's needs [17], which has been studied in many variations and has been addressed through a lot of summarization techniques [9, 2, 6, 29, 21, 25, 27, 3, 12, 23]. However, the demands of novel and aspect-specific information have not been fully recognized yet.

The goal of guided summarization task is to address these two new demands of summarization simultaneously. By providing concise, aspect-specific summaries of the periodical dynamic information devoted to a common topic, guided summary can save the users from browsing the web content with much redundancy. We can formulate the guided summarization task as aspect-based update summarization, which can be valuable for periodically monitoring the important changes of specific aspect from the documents varying over a given time period.

Guided summarization provides clearer requirements of automatic summary when faced with specific categories of documents. The difficulty lies in mining those specific aspects. Discovering the changes in the event is also a challenge. There are five categories in total, each category with a separate list of aspects. The categories and the corresponding aspects are listed as follows:

Accidents and Natural Disasters: what happened; date; location; reasons for accident/disaster; casualties; damages; rescue efforts/countermeasures

Attacks: what happened; date; location; casualties; damages; perpetrators; rescue efforts/countermeasures

Health and Safety: what is the issue; who is affected; how they are affected; why it happens; countermeasures

Endangered Resources: description of resource; importance of resource; threats to resource; countermeasures

Trials and Investigations: who is under investigation, who is investigating/suing; why (general); specific charges; sentence/consequences; how do they plead/react to charges

Summaries are supposed to find all the aspects corresponding to the category. Besides, an update summary is also required for each document collection B. Update summarization is essentially a temporal extension of topic-focused multi-document summarization. As defined in [1], the temporal summarization is to summarize from web documents over a given time interval. The temporal summarization focuses on the identification of changes between web documents. Both the requirements of novel information and aspect-specific information of guided summarization are seldom well addressed in current state of the art.

In this paper, we introduce two models, MRSP and TMSP, to cope with update summarization and guided summarization correspondingly. MRSP is dedicated to update summarization, which aims to create summaries with high topic-relevance, importance, and information novelty and diversity simultaneously. TMSP is a topic-guided version of MRSP, which aims to give aspect-specific summaries by modeling aspects as subtopics. Both MRSP and TMSP are based on traditional manifold ranking approach. Experiments are conducted to on datasets of TAC2008, TAC2009, and TAC2010. The evaluation results provided by TAC2010 shows the effectiveness of proposed models.

In Section 2, we give an overview of the related works. The proposed model MRSP and TMSP are demonstrated in section 3 and 4. The experiments and evaluation followed in section 5. Finally, we conclude this paper with a summary and discussion of results in TAC 2010, and look ahead to future work.

2 Related Work

2.1 Update Summarization

Update summarization is a temporal extension of topic-focused multi-document summarization [24, 22, 27, 28, 30, 7, 10], by focusing on summarizing up-to-date information contained in the new document set given a past document set. A major approach for update summarization is extractive summarization [15, 9, 18]. In the extractive approach, update summarization is reduced to a sentence ranking problem, which composes a summary by extracting the most representative sentences from target document set. There are four goals a ranking algorithm for update summarization aims to achieve:

- *Topic Relevance*: The summary is based on a topic-related multi-document set, where a topic represents user's information need (either a short query or narrative). Therefore, the summary must stick to the topic users are interested in.

- *Importance*: Not all the sentences in the documents deliver information of equal importance about the topic. The summary has to neglect trivial content and include important information instead.

- *Diversity*: There should be less redundant information in the summary, so that the limited summary space can cover as much information as possible about the topic.

- *Novelty*: Given a specified topic and two chronologically ordered document sets, the summary needs to focus on the new information conveyed by the later dataset as compared with the earlier one under that topic.

Technically, *novelty* can be considered as a special kind of diversity since it focuses on the difference between sentences of newcoming documents and those of earlier documents, while *diversity* focuses on the difference between sentences selected already and those to be selected next.

Update summarization is most commonly used in a dynamic web environment. Allan et al. [1] generated temporal summaries over news stories on a certain event, which could be considered as an early form of update summarization. Recently, Boudin et al. [4] described a scalable sentence scoring method, SMMR derived from MMR [5], where candidate sentences were selected according to a combined criterion of query relevance and dissimilarity with previously read sentences. However, neither MMR nor SMMR took the influence of importance into consideration. Wan et al. [26] presented the TimedTextRank algorithm, a PageRank variation with a time factor, to select new and important sentences for update summarization. They achieved diversity through an additional penalty step based on cosine similarity measurement in a heuristic way. Li et al. [13] presented a positive and negative reinforcement ranking strategy PNR^2 to capture novelty for update summarization. They also penalized redundancy similarly as [26] to encourage diversity. It's hard to address the four goals of update summarization in a unified way.

2.2 Manifold Ranking

Manifold ranking[32, 31] is a ranking approach based on semi-supervised learning. The goal of it is to rank the data with respect to the intrinsic global manifold structure collectively revealed by a huge amount of data. Generally speaking, for many real world data types this would be superior to a local method, which rank data simply by pairwise Euclidean distances or inner products.

This approach was successfully adopted by Wan [27] et al. to deal with topic-focused multi-document summarization. The prior assumption of manifold ranking is: (1) nearby points are likely to have the similar ranking scores; (2) points on the same structure (typically referred to as a cluster or a manifold) are likely to have the same ranking scores.

In the summarization context, the data points are denoted by topic description and all sentences in the documents. The ranking function can is as follows [31]:

$$f^* = (1-\alpha)(I-\alpha S)^{-1}y. \tag{1}$$

The parameter of manifold-ranking weight α specifies the relative contributions to the ranking scores from neighbors and the initial ranking scores. Note that self-reinforcement is avoided since diagonal elements of the affinity matrix S are set to zeros.

2.3 Probabilistic Latent Semantic Analysis

Probabilistic Latent Semantic Analysis (pLSA) [11] is a statistical model which has been called aspect model. The aspect model is a latent variable model for co-occurrence data which associates an unobserved class variable with each observation. pLSA is based on a mixture decomposition derived from a latent class model, this coincides with the aspect-specific demanding of guided summarization.

3 MRSP and TMSP

In this paper, we propose a novel approach MRSP [8] to address diversity as well as relevance and importance in ranking in a unified way. Specifically, MRSP assumes all the data and query objects are points sampled from a low-dimensional manifold and leverages a manifold ranking process [31, 32]. Such a ranking process tends to give the objects that are close to the query on the manifold and that have strong centrality higher rank. Therefore, it can naturally find the most relevant and important objects.

Meanwhile, to address the diversity in ranking, we first introduce the concept of sink points into the data manifold. The sink points are data objects whose ranking scores are fixed at the minimum score (zero in our case) during the ranking process. Hence, the sink points will never spread any ranking score to their neighbors. Intuitively, we can imagine the sink points as the "black holes" on the manifold, where ranking scores spreading to them will be *absorbed* and no ranking scores would *escape* from them. This way, the ranking scores of other points close to the sink points (i.e. objects sharing similar information with the sink points) will be penalized during ranking.

Our **overall algorithm** follows an iterative structure. At each iteration, we use manifold ranking to find one or more most relevant points. Then, we turn the ranked points into sink points, update scores, and repeat. By turning ranked objects into sink points on data manifold, we can effectively prevent redundant objects from receiving a high rank. Note here that the key idea of MRSP is similar to absorbing random walk [33]. However, absorbing random walk uses two different measures, *stationary distribution* and *expected number of visits*, to select the top ranked object and the remaining objects. It is largely different from MRSP where all the objects are ranked using a consistent strategy (i.e., using their ranking scores).

We now describe our MRSP algorithm in detail. Let $\chi = \chi_q \cup \chi_s \cup \chi_r \subset \mathbb{R}^m$ denote a set of data points over the manifold, where $\chi_q = \{x_1, \ldots, x_q\}$ denotes a set of query points, $\chi_s = \{x_1, \ldots, x_s\}$ denotes a set of sink points, and $\chi_r = \{x_1, \ldots, x_r\}$ denotes the set of points to be ranked, called *free points*. Let $f : \chi \rightarrow R$ denote a ranking function which assigns a ranking score f_i to each point x_i. We can view f as a vector $f = [f_1, \ldots, f_N]^T$, where $N = q + s + r$. We also define a vector $y = [y_1, \ldots, y_N]^T$, in which $y_i = 1$ if x_i is a query, and $y_i = 0$ otherwise. The MRSP algorithm works as follows:

1. Form the affinity matrix W for the data manifold, where $W_{ij} = sim(x_i, x_j)$ if there is an edge linking x_i and x_j. Note that $sim(x_i, x_j)$ is the similarity between objects x_i and x_j.

2. Symmetrically normalize W as $S = D^{-1/2}WD^{-1/2}$ in which D is a diagonal matrix with its (i, i)-element equal to the sum of the i-th row of W.

3. Repeat until χ_r is empty:

   (a) Iterate $f(t+1) = \alpha S I_f f(t) + (1-\alpha)y$ until convergence, where $0 \leq \alpha < 1$, and I_f is an indicator matrix which is a diagonal matrix with its (i, i)-element equal to 0 if $x_i \in \chi_s$ and 1 otherwise.

   (b) Let f_i^* denote the limit of the sequence $\{f_i(t)\}$. Rank points $x_i \in \chi_r$ according to their ranking scores f_i^* (largest ranked first).

   (c) Pick the top ranked point x_m. Turn x_m into a new sink point by moving it from χ_r to χ_s.

As we can see, the major difference between MRSP and the traditional manifold ranking algorithm is the introduction of sink points, which in turn affect the ranking process as shown in step $3(a)\sim(c)$. In the core iteration of the ranking process step $3(a)$, an indicator matrix I_f is used to fix the ranking scores of sink points at zero. As a result, the sink points will not spread any ranking score to their neighbors during the ranking process.

The main difference between MRSP and TMSP is the initial assignment of prior vector y in above algorithm. In TMSP, we first fit the pLSA model with an EM algorithm [11] to get the topic distributions of the documents. Then we make use of this distribution to assign the prior needed by above MRSP algorithm. In this way, we hope the aspect-specific requirement can be captured.

4 Experiments

4.1 Data Set

Summarization has been evaluated in TAC 2008 and TAC 2009, each task having a gold standard data set consisting of document clusters and reference summaries. The test data set of TAC 2010 is composed of 46 topics. Each topic has been assigned to a category and has 20 relevant documents which have been divided into 2 sets: Document Set A and Document Set B. Each document set has 10 documents, and all the documents in Set A chronologically precede the documents in Set B.

In our experiments, data set from TAC 2009 including the two chronologically ordered document sets A and B, together with their corresponding human model summaries are used for training and parameter tuning, data set from TAC 2010 without gold standard models are used for testing, and the results are evaluated and published by the organizers of TAC 2010. As a preprocessing step, the stop words in each sentence were removed and the remaining words are stemmed.

4.2 Evaluation Metric

ROUGE[14], Recall Oriented Understudy for Gisting Evaluation, is a metric adopted by TAC for automatic summarization evaluation. There are several variants that can be used in practice with provided toolkits. ROUGE-N measures summary quality by counting overlapping units of n-gram between the candidate summary (peer) and the reference summaries (model). The evaluation metrics we adopted in our training process are ROUGE-1, ROUGE-2 and ROUGE-SU4 respectively. ROUGE-N is computed as follows:

$$ROUGE-N = \frac{\sum_{S \in ReferenceSummries} \sum_{gram_n \in S} Count_{match}(gram_n)}{\sum_{S \in ReferenceSummries} \sum_{gram_n \in S} Count(gram_n)} \quad (2)$$

Where n stands for the length of n-gram, $gram_n$, and $Count_{match}(gram_n)$ is the maximum number of n-grams co-occurring in a candidate summary and a set of reference summaries. ROUGE-SU4 is a skip-bigram co-occurrence measure with addition of unigrams as counting unit.

The ROUGE toolkit reports scores for 1-, 2-, 3-, and 4-gram. We show three of the ROUGE metrics in the experimental results, at a confidence level of 95%: ROUGE-1, ROUGE-2, ROUGE-SU4.

Pyramid[19, 20] is a manual metric used for summary evaluation in TAC 2009. Its kernel concept is Summary Content Units, referred as SCUs, which are semantically motivated, sub-sentential units that are variable in length but no bigger than a sentential clause. SCUs emerge from annotation of a collection of human summaries for the same input. They are identified by noting information that is repeated across summaries, whether the repetition is as small as a modifier of a noun phrase or as large as a clause. The weight an SCU obtains is directly proportional to the number of reference summaries that support that piece of information. The evaluation method that is based on overlapping SCUs in human and automatic summaries is described in the Pyramid method.

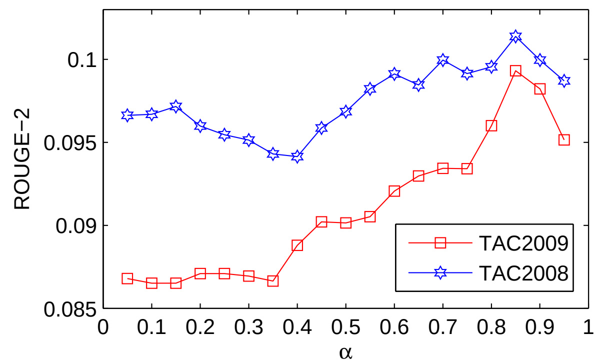

Figure 1: ROUGE-2 vs. Parameter α on MRSP

4.3 Experimental Results

4.3.1 Parameter Tuning

Figure 1 demonstrates the influence of parameter α in our MRSP algorithm. Three ROUGE-2 scores are obtained in our training process conducted on data set of TAC 2008 and 2009, update summarization task.

Scores arrive at their climax near the point of $\alpha = 0.85$. The high value of α means that the scores transposed to the point's neighbors from the prior score of its own at each iteration step are very low $(1-\alpha)y$, while much of the score transposed to their neighbors is the score that it has cumulated from the iteration process.

4.3.2 System Comparison

Our proposed approaches for guided summarization, which is an extension of update summarization, has shown competitive performance in Text Analysis Conference of 2010. The guided summarization task in TAC 2010 requires the generation of 100-word summaries for 46 topics. Each topic has a topic category and 20 relevant documents which have been divided into 2 sets: Document Set A and Document Set B. Each document set has 10 documents, and all the documents in Set A chronologically precede the documents in Set B. The generated summaries are evaluated by the National Institute of Standards and Technology (NIST[3]). All summaries were truncated to 100 words before being evaluated by manual and automatic metrics. The evaluation results of MRSP system of Run 8 are demonstrated in Table 1:

Table 1: Evaluation Results of MRSP in TAC10.

Metric	Score	Rank
Pyramid - A	0.344	21
Pyramid - B	0.276	3
BE - A	0.04475	22
BE - B	0.04350	3
ROUGE-2 - A	0.07700	20
ROUGE-2 - B	0.07251	4
ROUGE-SU4 - A	0.11104	20
ROUGE-SU4 - B	0.11039	5

Except for the automatic metrics of ROUGE family, NIST also conducted a manual evaluation of summary content based on the Pyramid Method[4].

The evaluation results of TMSP based approach - Run 30 in TAC 2010 - are demonstrated in Table 2:

Comparing to other participants, the performance of our system 8 is rather stable on ranks of sets B, which means that it captured the update nature of set B successfully. At the same time, its performance on sets A is not good enough, which might be caused by the unique setting of parameter α. Update summarization and topic-focused multi-document summarization should be treated with respect to parameter settings. Our run 30 does not perform well enough. This might be caused by the naive use of aspect model in TMSP. We simply extract the coarse distribution of subtopics and use the aspect information as a prior in MRSP. Perhaps this method does not capture the real aspects required by guided summarization, which leads to the poor performance of our second run, numbered 30.

[3]http://www.nist.gov/

[4]http://www1.cs.columbia.edu/~becky/DUC2006/2006-pyramid-guidelines.html

Table 2: Evaluation Results of TMSP in TAC10.

Metric	Score	Rank
Pyramid - A	0.351	18
Pyramid - B	0.273	8
BE - A	0.04529	21
BE - B	0.03962	7
ROUGE-2 - A	0.07623	21
ROUGE-2 - B	0.06957	10
ROUGE-SU4 - A	0.11042	21
ROUGE-SU4 - B	0.10703	16

Given this analysis, we believe that by keeping the well-performing parts of systems 8 and a finely developed aspect model, a promising performance can be expected in our future work.

5 Conclusions and Future Work

In this paper, we proposed two approaches for guided summarization task of TAC 2010. The first one is a MRSP-based summarization approach, in which the macro-structure of the information can be preserved, reflecting better topic-relevance, hence sentences with high topic-relevance, importance, novelty and diversity are extracted as summary candidates. Information redundancy is eliminated by the function of sink points on the sentence manifold. The second one is TMRSP-based approach, in which the aspects extracted by pLSA are adopted as the prior of the MRSP, trying to capture the aspect-specific feature of guided summarization.

MRSP has achieved promising results in TAC 2010 of sets B under of the evaluation metrics of ROUGE, BE and Pyramid. The MRSP approach is better at capturing the update essence of the information according to the evaluation results provided by NIST. moreover, it achieved a stable performance on every evaluation metric. However, the TMSP approach does not perform well enough. A possible reason for the poor performance might be that the aspect model adopted in TMSP fails to capture the aspect requirement of the guided summarization.

We will consider a more refined usage of the aspect model to improve the performance of TMSP for guided summarization. By combining the capability of capturing novelty and diversity of MRSP and of the ability of obtaining aspect by pLSA, we believe a more competitive system can be acquired in our future work.

6 Acknowledgments

This research work was supported by the State Key Program of National Natural Science Foundation of China (Grant No. 60933005), the Program of National Natural Science Foundation of China (Grant No.60903139 and Grant No.61003166).

References

[1] J. Allan, R. Gupta, and V. Khandelwal. Temporal summaries of new topics. In *SIGIR '01: Proceedings of the 24th annual international ACM SIGIR conference on Research and development in information retrieval*, pages 10–18, New York, NY, USA, 2001. ACM.

[2] R. Barzilay and M. Elhadad. Using lexical chains for text summarization. In *In Proceedings of the ACL Workshop on Intelligent Scalable Text Summarization*, pages 10–17, 1997.

[3] S. Berkovsky, T. Baldwin, and I. Zukerman. Aspect-based personalized text summarization. In *AH '08: Proceedings of the 5th international conference on Adaptive Hypermedia and Adaptive Web-Based Systems*, pages 267–270, Berlin, Heidelberg, 2008. Springer-Verlag.

[4] F. Boudin, M. El-Bèze, and J.-M. Torres-Moreno. A scalable MMR approach to sentence scoring for multi-document update summarization. In *Coling 2008: Companion volume: Posters*, pages 23–26, Manchester, UK, August 2008. Coling 2008 Organizing Committee.

[5] J. Carbonell and J. Goldstein. The use of mmr, diversity-based reranking for reordering documents and producing summaries. In *SIGIR '98: Proceedings of the 21st annual international ACM SIGIR conference on Research and development in information retrieval*, pages 335–336, New York, NY, USA, 1998. ACM.

[6] J. M. Conroy and D. P. O'leary. Text summarization via hidden markov models. In *SIGIR '01: Proceedings of the 24th annual international ACM SIGIR conference on Research and development in information retrieval*, pages 406–407, New York, NY, USA, 2001. ACM.

[7] J. M. Conroy and J. D. Schlesinger. Classy query-based multidocument summarization. In *In Proceedings of DUC'2005*, 2005.

[8] P. Du, J. Guo, J. Zhang, and X. Cheng. Manifold ranking with sink points for update summarization. In *CIKM '10: Proceeding of the 19th ACM conference on Information and knowledge management*, Toronto, Canada, 2010. ACM.

[9] G. Erkan and D. R. Radev. Lexrank: graph-based lexical centrality as salience in text summarization. *J. Artif. Int. Res.*, 22(1):457–479, 2004.

[10] E. Hovy, C. yew Lin, L. Zhou, and J. Fukumoto. Automated summarization evaluation with basic elements. In *In Proceedings of the Fifth Conference on Language Resources and Evaluation (LREC)*, 2006.

[11] K. B. Laskey and H. Prade, editors. *Probabilistic Latent Semantic Analysis*. Morgan Kaufmann, 1999.

[12] L. Li, K. Zhou, G.-R. Xue, H. Zha, and Y. Yu. Enhancing diversity, coverage and balance for summarization through structure learning. In *WWW '09: Proceedings of the 18th international conference on World wide web*, pages 71–80, New York, NY, USA, 2009. ACM.

[13] W. Li, F. Wei, Q. Lu, and Y. He. PNR2: Ranking sentences with positive and negative reinforcement for query-oriented update summarization. In *Proceedings of the 22nd International Conference on Computational Linguistics (Coling 2008)*, pages 489–496, Manchester, UK, August 2008. Coling 2008 Organizing Committee.

[14] C.-Y. Lin. Rouge: A package for automatic evaluation of summaries. In S. S. Marie-Francine Moens, editor, *Text Summarization Branches Out: Proceedings of the ACL-04 Workshop*, pages 74–81, Barcelona, Spain, July 2004. Association for Computational Linguistics.

[15] C.-Y. Lin and E. Hovy. Manual and automatic evaluation of summaries. In *Proceedings of the ACL-02 Workshop on Automatic Summarization*, pages 45–51, Morristown, NJ, USA, 2002. Association for Computational Linguistics.

[16] H. P. Luhn. The automatic creation of literature abstracts. *IBM Journal of Research Development*, 2(2):159–165, 1958.

[17] I. Mani. Recent developments in text summarization. In *CIKM '01: Proceedings of the tenth international conference on Information and knowledge management*, pages 529–531, New York, NY, USA, 2001. ACM.

[18] R. Mihalcea and P. Tarau. Textrank: Bringing order into texts. In D. Lin and D. Wu, editors, *Proceedings of EMNLP 2004*, pages 404–411, Barcelona, Spain, July 2004. Association for Computational Linguistics.

[19] A. Nenkova, R. Passonneau, and K. McKeown. The pyramid method: Incorporating human content selection variation in summarization evaluation. *ACM Trans. Speech Lang. Process.*, 4(2):4, 2007.

[20] A. Nenkova and R. J. Passonneau. Evaluating content selection in summarization: The pyramid method. In *HLT-NAACL*, pages 145–152, 2004.

[21] A. Nenkova, L. Vanderwende, and K. McKeown. A compositional context sensitive multi-document summarizer: exploring the factors that influence summarization. In *SIGIR '06: Proceedings of the 29th annual international ACM SIGIR conference on Research and development in information retrieval*, pages 573–580, New York, NY, USA, 2006. ACM.

[22] J. Otterbacher, G. Erkan, and D. Radev. Using random walks for question-focused sentence retrieval. In *Proceedings of Human Language Technology Conference and Conference on Empirical Methods in Natural Language Processing*, pages 915–922, Vancouver, British Columbia, Canada, October 2005. Association for Computational Linguistics.

[23] Y. Ouyang, W. Li, S. Li, and Q. Lu. Applying regression models to query-focused multi-document summarization. *Information Processing & Management*, In Press, Corrected Proof:–, 2010.

[24] H. Saggion, K. Bontcheva, and H. Cunningham. Robust generic and query-based summarisation. In *EACL '03: Proceedings of the tenth conference on European chapter of the Association for Computational Linguistics*, pages 235–238, Morristown, NJ, USA, 2003. Association for Computational Linguistics.

[25] D. Shen, J.-T. Sun, H. Li, Q. Yang, and Z. Chen. Document summarization using conditional random fields. In *IJCAI*, pages 2862–2867, 2007.

[26] X. Wan. Timedtextrank: adding the temporal dimension to multi-document summarization. In *SIGIR '07: Proceedings of the 30th annual international ACM SIGIR conference on Research and development in information retrieval*, pages 867–868, New York, NY, USA, 2007. ACM.

[27] X. Wan, J. Yang, and J. Xiao. Manifold-ranking based topic-focused multi-document summarization. In *IJCAI 2007, Proceedings of the 20th International Joint Conference on Artificial Intelligence*, pages 2903–2908, Hyderabad, India, January 6-12 2007.

[28] F. Wei, W. Li, Q. Lu, and Y. He. Query-sensitive mutual reinforcement chain and its application in query-oriented multi-document summarization. In *SIGIR '08: Proceedings of the 31st annual international ACM SIGIR conference on Research and development in information retrieval*, pages 283–290, New York, NY, USA, 2008. ACM.

[29] H. Zha. Generic summarization and keyphrase extraction using mutual reinforcement principle and sentence clustering. In *SIGIR '02: Proceedings of the 25th annual international ACM SIGIR conference on Research and development in information retrieval*, pages 113–120, New York, NY, USA, 2002. ACM.

[30] J. Zhang, X. Cheng, G. Wu, and H. Xu. Adasum: an adaptive model for summarization. In *CIKM '08: Proceeding of the 17th ACM conference on Information and knowledge management*, pages 901–910, New York, NY, USA, 2008. ACM.

[31] D. Zhou, O. Bousquet, T. N. Lal, J. Weston, and B. Schölkopf. Learning with local and global consistency. In S. Thrun, L. Saul, and B. Schölkopf, editors, *Advances in Neural Information Processing Systems 16*. MIT Press, Cambridge, MA, 2004.

[32] D. Zhou, J. Weston, A. Gretton, O. Bousquet, and B. Schölkopf. Ranking on data manifolds. In S. Thrun, L. Saul, and B. Schölkopf, editors, *Advances in Neural Information Processing Systems 16*. MIT Press, Cambridge, MA, 2004.

[33] X. Zhu, A. Goldberg, J. Van Gael, and D. Andrzejewski. Improving diversity in ranking using absorbing random walks. In *Human Language Technologies 2007: The Conference of the North American Chapter of the Association for Computational Linguistics; Proceedings of the Main Conference*, pages 97–104, Rochester, New York, April 2007. Association for Computational Linguistics.