Summarizing Definition from Wikipedia Articles

Zhicheng Zheng and Xiaoyan Zhu

State Key Laboratory of Intelligent Technology and Systems, Tsinghua National Laboratory for Information Science and Technology, Department of Computer Science and Technology, Tsinghua University, Beijing 100084, China
zhengzc04@gmail.com, zxy-dcs@tsinghua.edu.cn

Abstract. Definitional questions are quite important, since users often want to get a brief overview of a specific topic. It is a more challenging task to answer definitional questions than factoid questions. Since Wikipedia provides a wealth of structural or semi-structural information which covers a large number of topics, such sources will benefit the generation of definitions. In this paper, we propose a method to summarize definition from multiple related Wikipedia articles. First, we introduce the Wikipedia concepts model to represent the semantic elements in Wikipedia articles. Second, we further utilize multiple related articles, rather than a single article, to generate definition. The experiment results on TREC-QA demonstrate the effectiveness of our proposed method. The Wikipedia concept model outperforms the word model. Introducing multiple related articles helps find more essential nuggets.

Keywords: Definition, Summary, Wikipedia.

1 Introduction

In the daily life, people often want to know the definition of a specific topic. As shown in Voorhees [1], the definitional questions occur relatively frequently in logs of web search engines. Therefore, it is an important task to generate a definition, which consists of the most important and interesting aspects for the specific topic.

The task of definitional question answering, which aims to generate definitions for given topics, was included in the past QA tracks of TREC [2], and the evaluation results showed that the task is much more difficult than the factoid and list question answering [3]. Wikipedia, the largest online encyclopedia, contains a large number of topics. Each topic corresponds to a Wikipedia article describing it. For a given query, which represents exactly one topic, many search engines (e.g., Google, Yahoo!, Bing) often rank the corresponding articles in Wikipedia at top positions.

The first sentence or paragraph of a Wikipedia article may provide a brief and concise description of the corresponding concept. However, only one sentence or paragraph may be unable to cover enough important and interesting nuggets in which the users are interested, while the full article in Wikipedia may be too long

M. Zhou et al. (Eds.): NLPCC 2012, CCIS 333, pp. 260–271, 2012.

to read. Thus, it is necessary to summarize the Wikipedia articles to generate definitions.

Some researchers have done a variety of work on using Wikipedia to solve the definition problem. Summarizing definition from a single Wikipedia article can achieve a quite good result [4]. Besides the single Wikipedia article, we find that the related articles can help generating a better definition for the topic, due to the following reasons: (1) If the information about a specific topic is mentioned frequently in its related Wikipedia articles, generally the information is more important or interesting than other information to the topic. Hence, the definition should include this information. (2) The motivation of using related articles is similar with the explanation as Wan [5]: From human's perception, users would better understand a topic if they read more related articles. Hence, by adopting the enlarged knowledge within the related articles, the quality of definitions will be improved.

In this paper, we propose to summarize definition from multiple Wikipedia articles. We first introduce a model to use Wikipedia concept to represent semantic elements in the articles. By observing the advantage of other related articles, we further propose to use the related articles to summarize definition. Based on the related articles, we can compute better weights for semantic elements corresponding to the topic. The main contributions of the paper lie in two aspects: (1) We introduce the Wikipedia concept model, by which the semantic elements in the article can be represented more precisely. (2) We utilize related Wikipedia articles, rather than a single article, to generate definitions. We also measure the impact of the related article sets to different topics. Extensive experiments show that our method performs well for definitional questions.

The paper is organized as follows. Section 2 discusses some related work. The method of definition generation is described in Section 3. Experiments in Section 4 show the novelty and advantages of our work. Conclusions and future work are outlined in Section 5.

2 Related Work

Definitional Question Answering was firstly introduced into the TREC-QA Track in 2003. The definitional question answering is usually recognized as a difficult task. Some researchers attempt to retrieve definitional sentences using some hand-crafted patterns [6,7,8]. Knowledge intensive approaches can retrieve sentences of high quality; however, it requires experts to define all possible lexical or syntactic patterns. The method is not scalable as it is time and labor intensive. To overcome the deficiency, Cui et al. [9] propose to adopt soft pattern matching. Their method outperforms significantly those approaches with manually constructed patterns on the data set of TREC-QA 2004.

The pattern matching based methods are topic dependent, so even the soft pattern methods require that the topics in the training set are not biased. Another way to answer the definition question is to explore words or terms related to topics, then to use the words or terms to select the definitional sentences. In

TREC 2006, Kaisser et al. [10] collected signature words from snippets returned by search engine, and then selected the sentences containing most weight of the terms. In this way, their system outperformed all the other systems significantly. Since the centroid words or terms are quite important to such methods, Kor and Chua [11] proposed to build such centroid terms from Wikipedia, NewsLibrary, Google snippets and other resources like Wordnet.

Since each Wikipedia article is essentially an overview of a concept, Wikipedia is a good source of answering definitional questions. It is straightforward to extract definition for a topic from Wikipedia. The corresponding Wikipedia article is supposed to be highly focused on the topic, so each sentence in the article describes the topic no matter whether it contains the topic terms or not. Based on this assumption, Ye et al. [4] use EDCL to summarize the Wikipedia article as definition to the topic. EDCL is an extended document concept lattice model (DCL [12]) to combine Wikipedia concepts and non-textual features such as the outline and infobox.

However definitional question answering aims to find important and interesting nuggets for a topic, as mentioned by Kor and Chua [11], the important and interesting nuggets often come in the form of trivia, novel or rare facts about the topic that tend to strongly co-occur with direct mention of the topic keywords. Making use of related articles in Wikipedia will be helpful to generate the definition for the topic. On the other hand, it is challenging to summarize a single Wikipedia article. As analyzed in Ye et al. [4], the written style of Wikipedia articles is quite different from the free text used in traditional summarization tasks. In general, the guideline for composing Wikipedia articles is to avoid redundancy. Hence, there are low redundancies between the sentences within a Wikipedia article, compared to other types of documents. This is not compliant with the assumption that in traditional extractive summarization that the contents which are repeatedly emphasized should be included [13]. To overcome the problem, Ye et al. [4] try to group similar Wikipedia concepts and seek important contents by utilizing non-textual features such as outline and infobox. In addition, our proposal that incorporates other related articles is easier to identify key concepts with repeated mentions; hence, it will be more appropriate to use multiple document summarization methods.

The idea of generating a definition from related Wikipedia articles, as introduced above, to some extent is similar to single document summarization by expanding the single document to a small number of related documents [5,14,15]. Different from free-text documents, Wikipedia articles are organized structurally. Therefore, the categories of a Wikipedia article and the links between Wikipedia articles can be useful in finding related articles.

With related articles in hand, the definitional question answering can be approached as a multiple document summarization problem. Document summarization is a hot topic these years, and researchers have proposed many methods, such as Maximal Marginal Relevance (MMR [16]), document cluster centroid based [17] and some other graph-based methods (like LexRank [18] and TextRank [19]. Gillick and Favre [20] present an Integer Linear Program for exact inference under a maximum coverage model for summarization. In the summarization framework, it is important to investigate a good way to represent the content of text. Gabrilovich and Markovitch [21] explore Wikipedia articles to be semantic representation for text, which inspired researchers to explore encyclopedia knowledge to help summarization work. In this paper, we also propose a representation model making use of Wikipedia articles. Then, we use both the Maximal Marginal Relevance and Maximum Coverage methods to generate summary as definition, and analyze the effects of multiple related articles under different summarization algorithms.

3 Method

In the paper, we propose to summarize a definition from multiple related Wikipedia articles to a given topic. We firstly identify the corresponding Wikipedia article for a given topic. Then, we expand the Wikipedia article to a related Wikipedia article set, and we finally use multi-document summarization techniques to extract the definition for the topic. In this section, we will present the problem formulation, and then we will describe the article expansion methods and the multi-document summarizing methods in detail.

3.1 Representation Model

We first give some symbols of the representation model. Denote that d represents a Wikipedia article. For a given topic t, we aim to give a brief definition def_t to t. As mentioned, there is a Wikipedia article d_t, which focuses on describing the topic t. And with some article expansion methods, we get a Wikipedia article set D_t which contains the Wikipedia articles related to d_t. Here, Dt at least contains d_t. Each article d_i consists of n_i sentences, denoted as s^i_j ($j = 1$ to n). Each sentence is represented as a concept vector. Assuming that there are totally K concepts, denoted as c_k ($k = 1$ to K). For each s^i_j and c_k, we calculate a weight $w^i_j(k)$ according to the importance of c_k in s^i_j. Similarly, we calculate $W^t(k)$ representing the importance of c_k to specific topic t.

Next, we will give two different methods to construct the concept vector space and calculate the weights for each concept in a topic: (1) A simple way to represent concepts by words; and (2) A more precise method to represent the concept by Wikipedia concepts.

Words Model. It is an intuitive way to represent concept by each word. We can calculate $w^i_j(k)$ as TF-IDF weight:

$$w^i_j(k) = tf^i_j(k) * idf(k)$$

Here, $tf^i_j(k)$ is the term frequency of c_k in s^i_j, and $idf(k)$ is the inverse document frequency of c_k, which is calculated on the whole Wikipedia corpus.

We calculate $W^t(k)$ by summing up the weight of all the sentences in all related articles:

$$W^t(k) = \sum_{d_i \in D_t} \sum_{j=1}^{n_i} w_j^i(k)$$

Wikipedia Concept Model. The word model could not precisely represent the exact concept in a text. Consider a person's name: '*Jordan Hill*' with the other two different names: '*Jordan Farmer*' and '*Grant Hill*'. By word model, *Jordan Hill* overlaps with the other two names in the concept vector space. However, they should be distinguished as three different concepts. Inspired by Gabrilovich and Markovitch's work [21], we investigate to using Wikipedia articles to represent a text instead of word model. We define the Wikipedia concepts as the concepts that can be mapped to Wikipedia articles, and we construct the concept vector space by using Wikipedia concepts, in which way we expect that it is able to better represent the content of sentence.

There are many inner links in Wikipedia. Each link consists of an anchor text and a target Wikipedia concept that it links to. We collect the anchor text and target Wikipedia concepts of all the inner links. All the anchor text form a phrase set A, and according to the inner links each phrase a in the set is mapped to a set of Wikipedia concepts, which is denoted as $C(a) = \{c_k\}$. And by counting link pairs of a and c_k, which is denoted as $l_k(a)$, we could assign a prior probability $p_a(k)$ to each c_k in $C(a)$ as:

$$p_a(k) = \frac{l_k(a)}{\sum_{c_i \in C(a)} l_k^a}$$

For each sentence s_j^i in d_i, with the dictionary consisting of phrases in A, we detect a set of anchor phrases $A_j^i = \{a_z\}$ by applying forward maximum matching in s_j^i. And then we assign a Wikipedia concept to each a_z according to both the context of the Wikipedia article and the prior distribution of Wikipedia concept. We operate as follows:

1. Collect the inner links in d_i, and count the number of links with different target Wikipedia concepts (Denote the number of links with target Wikipedia concept c_k as l_k^i).
2. For each a_z in A_j^i, the probability of that c_k is assigned to a_z is calculated as:
$$p(c_k|a_z, d_i) = p_{a_z}(k) \cdot \frac{l_k^i + \alpha}{|C(a_z)| \cdot \alpha + \sum_{j=1}^{|C(a_z)|} l_j^i}$$
where α is a smooth factor to avoid zero probability when $l_k = 0$ (In experiments, we set $\alpha = 0.1$), and $|C(a_z)|$ is the size of the set of Wikipedia concepts that a_z is mapped to.
3. For each a_z, we assign the Wikipedia concept c_k with highest $p(c_k|a_z, d_i)$.

After assigning Wikipedia concepts to all anchor phrases in A_j^i, we calculate the concept frequency $cf_j^i(k)$ of c_k by counting the occurrence number of c_k in s_j^i.

On the other hand, for each c_k, we also count the number of links whose target Wikipedia concept is c_k. Denote the number as l_k, it can be obtained by $\sum_a l_k(a)$. And the original importance of c_k, denoted as $coi(k)$, is calculated as:

$$coi(k) = \log l_k$$

The weight of c_k in s_j^i can be calculated by multiplying $cf_j^i(k)$ and $coi(k)$:

$$w_j^i(k) = cf_j^i(k) \cdot coi(k)$$

We can get $W^t(k)$ by the same way as the words model.

Additionally, we add each attributes of infobox of d_i as pseudo concepts. We rank the sentences in d_i according to the similarity of the sentences and the attribute, and add the pseudo concept to the top 2 sentences. The original importance of the concept is manually assigned, and then we treat them as common Wikipedia concepts.

3.2 Wikipedia Article Expansion

The motivation to summarize definition from multiple related Wikipedia articles is as follows: (1) The principle of typical extractive summarization approaches is that the contents which are repeatedly emphasized should be included. (2) As Wikipedia articles are human-written overview pages in which redundancy has been avoided, it is difficult to weight the importance of different concepts just using a single Wikipedia article. Therefore, it is more appropriate to summarize definition with multiple related articles. (3) A concept may be important and interesting when the concept mentioned in highly related Wikipedia articles.

In the paper, to make better use of Wikipedia's structural information, we retrieve the related Wikipedia articles to a specified Wikipedia article d_t using its inner links. However, even d_t contains an inner link links to another Wikipedia article $d_{t'}$, it does not always imply that d_t is related with $d_{t'}$. In fact many phrases in a Wikipedia article link to other articles just because there are entries for the corresponding Wikipedia concepts. To verify the relatedness, $d_{t'}$ is added into D_t if and only if $d_{t'}$ and d_t link with each other.

3.3 Multi Wikipedia Articles Summarization

Related articles also bring in more noises when generating definition by the extractive summarization approaches. To avoid the noises, we make a constraint that when generating definition for topic t, we only extract the sentences from d_t. As the purpose of the Wikipedia editors, the corresponding Wikipedia article d_t should always focus on the topic t. With the limitation, we miss some nuggets that exist in other related Wikipedia articles, but we ensure the relatedness between the extracted sentences and the topic t.

We utilize two multiple documents summarization algorithms: Maximal Marginal Relevance (MMR) and Maximum Coverage (MC). As we model the articles and make a constraint, we will describe the summarizing methods next.

Maximal Marginal Relevance. A good summary should meet the following two conditions: (1) The summary focus on the topic of the article set; (2) The summary need to avoid redundancy. The MMR algorithm considers both factors, and repeatedly select the sentence that can be more representative for the article set and has less redundancy with the sentences already selected.

As we assume that the content of the corresponding article d_t should focus on t, so the representativeness of a sentence s_j^t for article set D_t, denoted as $RP(s_j^t|D_t)$, is calculated as:

$$RP(s_j^t|D_t) = \sum_{c_k \in s_j^t} W_k^t$$

And the redundancy between two sentences s_j^t and $s_{j'}^t$, denoted as $RD(s_j^t|s_{j'}^t)$, is calculated as:

$$RD(s_j^t|s_{j'}^t) = \sum_{c_k \in s_j^t \cap s_{j'}^t} W_k^t$$

So under the condition of existing summary sentence set S_t, the maximal marginal relevance sentence s_{mmr} is calculated as:

$$s_{mmr} = \arg \max_{s \in d_i \setminus S_t} \{RP(s, D_t) + \max_{s' \in S_t} \{RD(s, s')\}\}$$

Maximum Coverage. Gillick and Favre proposed a summarization method based on maximum coverage. The method selects sentences with a globally optimal solution that also address redundancy globally. They choose to represent information at a finer granularity than sentences, with concepts, and assume that the value of a summary is the sum of the values of the unique concepts it contains. They also present the Integer Linear Program for exact inference under the model. We formulate our ILP problem as follows:

$$\begin{aligned}
\text{Maximize:} \quad & \sum_k W^t(k) \cdot SumC_k \\
\text{Subject to:} \quad & \sum_j o_j^t \leq L \\
& o_j^t \cdot Occ_j^t(k) \leq SumC_k, \forall k, j \\
& \sum_j o_j^t \cdot Occ_j^t(k) \geq SumC_k, \forall k
\end{aligned}$$

$$SumC_k \in \{0,1\}, \forall k$$
$$o_j^t \in \{0,1\}, \forall j \tag{1}$$

Here, $SumC_k$ represents whether c_k occurs in the summary or not, o_j^t means whether s_j^t is selected in the summary or not, and $Occ_j^t(k)$ equals 1 if and only if $cf_j^t(k) \geq 1$. Since the concept weight is always positive in the model we defined, (1) can be equally transformed to:

$$\begin{aligned}
\text{Maximize:} \quad & \sum_k W^t(k) \cdot SumC_k \\
\text{Subject to:} \quad & \sum_j o_j^t \leq L \\
& \sum_j o_j^t \cdot Occ_j^t(k) \geq SumC_k, \forall k \\
& SumC_k \in \{0,1\}, \forall k \\
& o_j^t \in \{0,1\}, \forall j
\end{aligned} \tag{2}$$

(2) removes the second constraint in (1), and the complexity of (2) is reduced.

4 Experiment

4.1 Experiment Setting

We evaluate our method on the corpus of TREC-QA in 2004-2006 (TREC 13-15). For each topic, we retrieve the corresponding Wikipedia article. Because the focus of the paper is on summarization evaluation, we simply ignore the topics in TREC-QA where the corresponding articles do not exist in Wikipedia.

We evaluate the summarization performance by pourpre [22]. Like prior studies [9,4], we also treat the answers of factoid/list questions as essential nuggets, and add them to the gold standard list of definition nuggets. To avoid the influence of those nuggets that do not exist in Wikipedia corpus, we only consider the nuggets that could be found in Wikipedia. So we first explore the available answers in Wikipedia Corpus, and the result is shown in Table 1. Among 215 topics in TREC 13-15, we could obtain 190 Wikipedia articles corresponding to the exact topics. As our corpus was downloaded in 2009, the available topic number and available nuggets number in the single article are both larger than those in 2007, which indicates that the Wikipedia not only covers more and more new topics, but also covers more old topics. So the idea of using Wikipedia as a resource to answer definitional questions is feasible. As mentioned in Section 3.3, we limit our algorithm to select sentences only from d_t. We observe that this constraint caused a lost of 25% essential nuggets could be found in related Wikipedia article set. It seems to be a huge loss in recall, but we can benefit the precision. We evaluate the precision of the two special ways of summarization, using the entire corresponding article as summary (S_{alls}) and using all the related articles in D as summary (S_{allm}). The precision of S_{alls} is 0.169 while the precision of S_{allm} is 0.029. Without the constraint on sentence selection, the number of the nuggets that can be retrieved will increase, but it introduces too many noises. With comprehensive consideration about recall and precision, the sentence selection constraint is reasonable.

Table 1. Availability Analysis

	Wikipedia 07	Wikipedia 09
Available Topics	180/215	190/215
Available Nuggets (single article)	47%	55%
Available Nuggets (multiple articles)		72%

We examine the quality of definition summary by nugget recall (NR, only consider the nuggets can be found in d_t) and an approximation to nugget precision (NP) on answer length. NR and NP are then combined using F1 and F3 measures. The evaluation is automatically conducted by Pourpre v1.1.

4.2 Performance Evaluation

To measure the performance of different models, we evaluate the quality of definition produced by MMR and MC algorithms combined with different representation models and article sets: Word model with a single article (Word), Wikipedia concept model with a single article (Concept), and Wikipedia concept model with related article set (C + M). The maximum number of sentences in a summary was set to 10. As the result shown in Table 2.

Table 2. Evaluated Result (sentences num = 10)

	NP	NR	F3	F1
Word (MMR)	0.576	0.574	0.561	0.545
Concept (MMR)	0.634	0.580	0.572	0.573
C+M (MMR)	0.649	0.609	0.601	0.599
Word (MC)	0.593	0.591	0.579	0.564
Concept (MC)	0.646	0.604	0.596	0.593
C+M (MC)	0.667	0.637	0.629	0.623

In both the summarization algorithms, the Wikipedia concept model outperforms the word model and the related article set helps improving the performance. We get following observations from the results:

1. Both the two algorithms benefit from the Wikipedia concept model. On all the evaluation metrics, the Wikipedia concept model outperforms the word model. For example, the Wikipedia concept model outperforms word model by about 5.1% on F1 score with both MMR algorithm and MC algorithm.

2. The related article set can help improving the performance in both the two algorithms. When using the related articles, all the evaluation metrics achieve improvements. For example, on F3 score, the MMR algorithm improves by 5.1% and the MC algorithm improves by 5.5%.
3. The Wikipedia concept model contributes more to precision than to recall. The Wikipedia concept weight can better represent the basic information elements in a sentence, and it helps to avoid selecting redundant sentences in a summary. Since we don't cluster the Wikipedia concepts, the redundancy in a Wikipedia article is too low to give guidance to get the essential nuggets. So when using Wikipedia concepts, the improvement on recall metric derives from that when more different nuggets are selected, more essential nuggets may be covered. The analysis can also explain why the recall of the Wikipedia concept model is even lower when the summary is short.
4. The related article set leads to more improvement in terms of nugget recall than the Wikipedia concept model. The more important the concepts are, the more frequent they occur in other related articles. The related article set will enhance the weight of these concepts, then the summarization algorithms will tend to choose the sentences containing these concepts. Hence, article expansion helps to extract more essential nuggets. For example, Table 3 lists most important concepts for topic *Manchester United Football Club* while using a single article and the related article set. Since in the main article about *Manchester United Football Club*, *win, season, player* are mentioned many times, they obtained a quite high weight in the single article although their original importance in Wikipedia is low. Related articles repeat the concepts of the relevant matches and main opponents, which are more important to the topic. Hence, while using related article set, the summarization algorithms will prefer to extracting the sentences containing these important concepts, which will improve the performance.

Table 3. Important Concepts for Manchester United Football Club

Rank	Single Article	related Articles
1	Manchester United F.C.	Manchester United F.C.
2	Association football	FA Cup
3	English language	Premier League
4	Win (baseball)	Liverpool F.C.
5	Season (sports)	UEFA Champions League
6	Player (game)	Arsenal F.C.

5 Conclusion

In the paper, we present a framework of summarizing definition from multiple Wikipedia articles. Experiments with different summarization algorithms demonstrate that the explicit semantic representation via Wikipedia concepts benefits the extraction of definition. The experiment results also show that the related articles can weight concepts more effectively than a single article, particularly for those general and popular topics. The framework proposed in the paper achieves excellent results on TREC-QA data, which demonstrates the feasibility of our methods.

However, using the extractive summary as definition still faces some problems, such as the discourse consistency between the extracted textual segments. In the future, we are considering using some generative summarization techniques (such as compression, reordering) to improve the consistency quality.

Acknowledgement. This work was carried out with the aid of a grant from the International Development Research Center, Ottawa, Canada, number:104519-006. This work was supported by the Chinese Natural Science Foundation under grant No.60973104.

References

1. Voorhees, E.M.: Overview of the trec 2001 question answering track. In: Text REtreival Conference (2001)
2. Voorhees, E.M.: Overview of the trec 2004 question answering track. In: Text REtreival Conference (2004)
3. Voorhees, E.M., Dang, H.T.: Overview of the trec 2005 question answering track. In: Text REtreival Conference (2005)
4. Ye, S., Chua, T.S., Lu, J.: Summarizing definition fromwikipedia. In: ACL 2009 (2009)
5. Wan, X., Yang, J.: Single document summarization with document expansion. In: AAAI 2007 (2007)
6. Blair-Goldensohn, S., McKeown, K.R., Schlaikjer, A.H.: A hybrid approach for qa track definitional questions. In: Text REtreival Conference (2003)
7. Xu, J., Licuanan, A., Weischedel, R.: Trec 2003 qa at bbn: Answering definitional questions. In: Text REtreival Conference (2003)
8. Harabagiu, S., Moldovan, D., Clark, C., Bowden, M., Hickl, A., Wang, P.: Employing two question answering systems at trec 2005. In: Text REtreival Conference (2005)
9. Cui, H., Kan, M.Y., Chua, T.S.: Generic soft pattern models for definitional question answering. In: SIGIR 2005 (2005)
10. Kaisser, M., Scheible, S.: BonnieWebber: Experiments at the university of edinburgh for the trec 2006 qa track. In: Text REtreival Conference (2006)
11. Kor, K.W., Chua, T.S.: Interesting nuggets and their impact on definitional question answering. In: SIGIR 2007 (2007)
12. Ye, S., Chua, T.S., Kan, M.Y., Qiu, L.: Document concept lattice for text understanding and summarization. Information Processing and Management 43(6), 1643–1662 (2007)
13. Silber, H.G., McCoy, K.F.: Efficiently computed lexical chains as an intermediate representation for automatic text summarization. Computational Linguistics 28(4), 487–496 (2002)

14. Wan, X., Yang, J., Xiao, J.: Incorporating Cross-Document Relationships Between Sentences for Single Document Summarizations. In: Gonzalo, J., Thanos, C., Verdejo, M.F., Carrasco, R.C. (eds.) ECDL 2006. LNCS, vol. 4172, pp. 403–414. Springer, Heidelberg (2006)
15. Wan, X., Yang, J.: Collabsum: exploiting multiple document clustering for collaborative single document summarizations. In: SIGIR 2007 (2007)
16. Carbonell, J., Goldstein, J.: The use of mmr, diversity-based reranking for reordering documents and producing summaries. In: SIGIR 1998 (1998)
17. Radev, D.R., Jing, H., Stys, M., Tam, D.: Centroid based summarization of multiple documents. Information Processing and Management 40, 919–938 (2004)
18. Erkan, G., Radev, D.R.: Lexpagerank: Prestige in multi-document text summarization. In: EMNLP 2004 (2004)
19. Mihalcea, R., Tarau, P.: Textrank - bring order into texts. In: EMNLP 2004 (2004)
20. Gillick, D., Favre, B.: A scalable global model for summarization. In: ACL 2009 (2009)
21. Gabrilovich, E., Markovitch, S.: Computing semantic relatedness using wikipedia-based explicit semantic analysis. In: Proceedings of the 20th International Joint Conference on Artificial Intelligence, vol. 6, p. 12. Morgan Kaufmann Publishers Inc. (2007)
22. Lin, J.J., Demner-Fushman, D.: Methods for automatically evaluating answers to complex questions. Information Retrieval 9(5), 565–587 (2006)