

What Is Needed for a Good Summary? – Two Different Types of Document Sets Yet Seemingly Indistinguishable to Human Users

Qunhua Zhao, Eugene Santos, Jr.,
Thayer School of Engineering
Dartmouth College
Qunhua.Zhao@dartmouth.edu
Eugene.Santos.Jr@dartmouth.edu

Hien Nguyen
Math. and Comp. Sci.
Univ. of Wisconsin
- Whitewater
nguyenh@uw.edu

Ahmed A. Mohamed
Dept. of Comp. Sci. and
Engineering
Univ. of Connecticut
amohamed@enr.uconn.edu

Abstract

Working with the DUC 2002 collection for multi-document summarization, we considered two types of document sets: sets consisting of closely correlated documents with highly overlapped content; and sets of diverse documents covering a wide scope of topics. Intuitively, this suggests that creating a quality summary would be more difficult for the latter case. The two types of document sets can be identified automatically by our document graph approach. However, human evaluators were shown to be fairly insensitive to this difference. This was identified when they were asked to rank the performance of automated summarizers. In this paper, we examine and analyze our experiment in order to better understand this phenomenon and how we might address it to improve summarization.

1. Introduction

The sheer amount of textual information available nowadays results in the well-known information overload problem [11]. The requirement for quickly obtaining relevant information in order to make timely decisions have further exacerbated this problem. Research into automatic text summarization has been proposed as one natural approach to dealing with this problem and is currently being carefully studied [1, 5]. Summaries are used to provide readers with the key information in the original texts, and to help readers to judge the relevancy of the text and decide if it is worth going through the full texts.

Text summarizations has been defined as “the process of distilling the most important information from a source (or sources) to produce an abridged version for a particular user (or users) and task (or tasks)” [5]. There are three elements that are included in this definition: 1) a summary is a condensed version that contains important information; 2) a

summary can be generated for a single document or a set of documents; and 3) a summary should be user- and task-oriented. Taking into account the possible uses that a summary may serve, a summary can be indicative by pointing out the topics which have been addressed in the text, or informative by covering as much as possible, important content and/or critical information that offers a critique of the source. Hence, the intention and coverage of summaries can be different according to the targeted tasks. Even for the same tasks, various people may have different opinions on what pieces of information are interesting and important. However, the prevalent form of summaries that we usually encounter is that of generic summaries that targets a wide range of readers. Only recently, user-sensitive (or user-centered) summarization has become an increasingly active research area, which studies how individual differences affect a user’s judgment on what should be selected to be included in a summary [1].

In the case of multi-document summarization, where a summary is generated from a set of related documents, it is typical that a variety of topics need to be covered. For example, assuming that there is a document set containing news reports on the tsunami disaster in south-east Asia, the topics included in this set may vary from what has happened in different locations, the loss of lives in different countries, the impact on the local economies, and the various international relief efforts. As has been pointed out, it is extremely difficult to determine what should be covered in a good summary because of the breadth of the multi-document set [7].

Within the standard testbed for the Document Understand Conference (DUC) 2002 data collection [10], there are basically two types of document sets for multi-document summarization. There are document sets which consist of closely related documents; and those of highly diverse texts.

Intuitively, it should be much more difficult to create a good summary for the document sets in the latter case. However, in conducting a user evaluation experiment, we found that human participants were fairly insensitive to this difference. In this paper, we present and analyze our experiment in order to better understand this phenomenon and how we might address it to improve summarization. We believe our result arose from the fact that there were no specific guidelines or tasks required during the summary evaluations. In this situation, we believe that human readers tend to accept summaries which simply contain general information as quality summaries. We describe our document graph (*DG*) approach, which captures semantic/structural information extracted from the text, as a possible approach to help us automatically identify when a summary has been created from a document-set that contains a diversity of information. This would then allow us to alert a user to not take the summary at face value. This can eventually lead to proper personalization of summaries with respect to user tasks and needs.

In the following text, we will first introduce the data collection we used in the experiment. Next, we describe the document graph approach that we used to identify the two types of document sets. Following that, we describe our experiments on evaluating the performance of various automatic summarization systems (summarizers) and provide a comparison between different summarization ranking approaches currently in use. Finally, we conclude with a discussion of our results.

2. DUC 2002 data collection for multi-document summarization

The National Institute of Standards and Technology launched a study on automatic text summarization and evaluation called the Document Understanding Conference (DUC). Since 2001, different types of summarization tasks have been studied, such as single document summarization, multi-document summarization, extract and abstract generation, and headline generation. A data collection is provided to serve as a testbed for state-of-the-art algorithms and systems.

We used the DUC 2002 data collection for multi-document summarization in our experiments: It has 59 document sets. Each set has 5-15 documents (with an average of 10), which are articles from the Wall Street Journal, AP newswire, San Jose Mercury News, Financial Times, LA Times, and FBIS records [10].

The document sets have been classified into four categories, as follows:

Category 1: Documents about a single natural disaster and created within at most a seven day window.

Category 2: Documents about a single event in any domain created within at most a seven day window.

Category 3: Documents about multiple distinct events of a single type (no limit on the time window).

Category 4: Documents that present biographical information mainly about a single individual.

We will refer to them as cat 1, cat 2, cat 3, and cat 4 sets.

There were ten participants that submitted summarization outputs generated by their automatic summarizers. For each document set, two model summaries were also created by human assessors. These summaries are extracts. Two different sizes of extracts, 200 words or 400 words, have been generated for each document set. For this paper, we worked with 200 word extracts in our experiments in order to reduce the time for user evaluations.

3. Using a document graph approach to identify different needs for summarization

Generating a summary is actually a process of extracting important relevant information and then presenting it to the user. Accordingly, we use an approach called document graph (*DG*) generation for information extraction and representation, which is described below, and use the resulting *DG* for the purpose of document similarity comparison. For more details on our *DG* generation research, see [13].

3.1. Document graph

A *DG* is a directed graph of concepts/entities and the relations between them. It contains two kinds of nodes, concept/entity nodes and relation nodes. Currently, only two kinds of relations, “*isa*” and “*related to*”, are captured for simplicity. The construction of a *DG* is an automated process, which contains following steps, 1) tokenizing a document (a summary is treated the same as a document) in the plain text format into sentences; 2) parsing each sentence by using Link Parser [14], 3) extracting noun phrases (NPs) from the parsing results; and, 4) generating relations between concepts/entities based on heuristic rules, and put them into the graph format.

The three heuristic rules for relation generation, in our current version, are:

- The NP-heuristic: it helps set up the hierarchical relations. For example, from a NP “folk hero stature”, we generate relations “folk hero stature *isa* stature”, “folk hero stature *related to* folk hero”, and “folk hero *isa* hero”.
- The NP-PP-heuristic: it attaches all prepositional phrases to adjacent NPs. For example, from “workers at a coal mine”, we generate a relation, “worker *related to* coal mine”.
- The sentence-heuristic: it relates all the concepts/entities contained within one sentence. The relations created by sentence-heuristic are then sensitive to verbs, since the interval between two noun phrases usually contains a verb. For example, from a sentence “workers at a coal mine went on strike”, we generate a relation “worker *related to* strike”. Another example, from “The usual cause of heart attacks is a blockage of the coronary arteries”, we generate “heart attack cause *related to* coronary artery blockage”. Figure 1 shows an example of a partial DG.

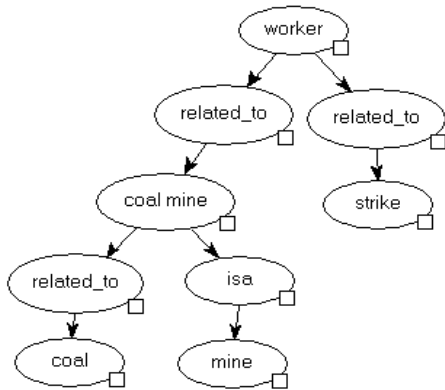


Figure 1. A partial DG.

We evaluate the similarity between the two documents based on the DGs that are generated from them. The simple similarity of two DGs, DG_1 to DG_2 , is given by the equation:

$$Similarity(DG_1, DG_2) = \frac{n}{2N} + \frac{m}{2M} \quad (1)$$

which is modified from Montes-y-Gómez et al. [9]. In the equation, N is the number of concept/entity nodes in DG_1 , and M stands for number of relations in DG_1 ; n is the number of matched concept/entity nodes in two DGs, and m is the number of matched relations. We say we find a matched relation in two different DGs, only when both of the two concept/entity nodes linked to that relation node are matched, and the relation node is also matched. Since we might compare two DGs that are significantly different in size (for example, DGs for an extract vs. its source document), we used the number of concept/entity nodes and relation nodes in the target DG as N and M , instead of the total number of nodes in both DGs. $Similarity(DG_1, DG_2)$ provides a percentage of DG_1 that is contained in DG_2 ; at the same time, $Similarity(DG_2, DG_1)$ represents how much of DG_2 has been covered in DG_1 . Next, an F -score can be calculated by Equation 2 (Rijsbergen, 1979).

$$F = \frac{2 \times P \times R}{P + R} \quad (2)$$

In the equation, P is precision and R is recall, where we define P to be $Similarity(DG_1, DG_2)$ and R as $Similarity(DG_2, DG_1)$. The F -score is then used as a direct measurement of the similarity between two documents and in ranking the performance of different summarizers.

Currently, we weight all the concepts/entities and relations equally. This can be fine tuned in the future [13].

3. 2. Two types of document sets

The DUC 2002 collection for multi-document summarization was used in our experiments. We generated a DG for each document in the collection, and then calculated the value of similarity (F -score) between each pair of DGs within the same document set, and finally obtained the average similarity score. Using our DG similarity, we can naturally and automatically divide the document sets into two simple types, say types A and B. In type A document sets, documents are more similar to each other; and in type B document sets, the documents are more diverse. The similarity data demonstrated that most of the document sets in cat 1 and cat 2 are of type A, and most of the document sets in cat 3 and 4 are of type B (as shown in Table 1).

As described in Section 2, document sets in cat 1 and cat 2 contain news articles about a single disaster or event within a period of seven days. Intuitively, these documents have more content overlap with each

other and are more similar to each other than those in the other two categories, where the document sets in cat 3 are about distinct events with no time window restriction, and biographical information about a person in cat 4 can be perceived from different points of view.

Table 1. The average F-scores for the document sets in different categories

Category	Cat 1	Cat 2	Cat 3	Cat 4
F-score	0.10	0.12	0.06	0.07

An example for the type A document set is D.79.E.200.A in the collection, which has been classified by DUC as cat 1 document set. It contains 9 articles, which are news reports about Hurricane Gilbert: when and where did it happen, and what kind of damage it caused. Obviously, the content of these news articles are highly overlapped. When creating the model extracts for this set, one human assessor picked 7 sentences and another picked 8; between them, 4 sentences are actually exactly the same. (As mentioned earlier in Section 2, the model summaries generated by human assessors are provided in the DUC 2002 data collection). A type B document set covers much broader topics, for example, the set D.106.E.200.G, which has been classified by DUC as cat 4 document set. It contains 10 articles consisting of stories such as: Northeastern University planning to award an honorary degree to Nelson Mandela; imprisoned African National Congress Leaders being allowed to visit Mandela; a family group visited Mandela; the Jewish-American group concerned about Mandela’s attitude towards Israel and Palestine Liberation Organization; and, news that he was going to be released.

The following are the two model extracts created for D.106.E.200.G set by human judges.

Assessor A: ONE MAN'S STRUGGLE

From behind bars, Nelson Mandela has dominated the fight for black rights in South Africa. The following are some of the key events in his life. Mandela joins the African National Congress at age 26, later becoming president of the group's Youth League. When the National Party comes to power in 1948, the ANC begins planning a campaign of civil disobedience to fight the party's apartheid policy. April-June, 1964 Mandela and seven others are sentenced to life in prison. December, 1989 President Frederik W. de Klerk meets Mandela for the first time. World leaders welcomed South Africa's announcement Saturday of Nelson R. Mandela's

impending release from prison, and international rejoicing began to build for an event so long awaited by so many. Soweto and other black townships around Johannesburg have been hit by black factional fighting that has killed about 800 people since August. In the 15 months since he walked free from prison, Nelson Mandela has played the chief role in talks with President F.W. de Klerk and the white government. The radicals in the ANC may well use the trial and conviction of Mrs. Mandela to arouse members against Nelson Mandela's policy of peace and negotiation.

Assessor F: *Mandela the 70-year-old leader of the ANC jailed 27 years ago, was sentenced to life in prison for conspiring to overthrow the South African government. Mandela was allowed to meet last Friday with five other ANC members convicted with him in 1964 of trying to overthrow the government and still imprisoned in Cape Town, according to Dullah Omar, a family friend and attorney. President Pieter W. Botha offers to free Mandela if he renounces violence; Mandela refuses until the government takes steps to dismantle apartheid and grants full political rights to blacks. De Klerk lifts ANC ban and says Mandela will be released. Commonwealth Secretary General Shridath Ramphal said Mandela's safety must be guaranteed by the South African government and that until change there becomes irreversible, the Commonwealth will sustain its pressure for the end of apartheid. The arrival Wednesday of African National Congress leader Nelson Mandela prompted an outpouring of praise and encouragement from politicians and supporters. In the 15 months since he walked free from prison, Nelson Mandela has played the chief role in talks with President F.W. de Klerk and the white government. Mandela and other ANC leaders have been negotiating with President F.W. de Klerk to pave the way for full-scale talks on ending apartheid, the South African system of racial separation.*

Assessor A spent the first half of the summary listing important events in the life of Nelson Mandela, then continued with his role in talks with the white government in South Africa and the problems he was facing. While in Assessor F’s summary, the information was more about what was happening at the time of the news concerning Mandela’s release, after very briefly mentioning that he is an African National Congress (ANC) leader and has been jailed for 27 years. From the example, it is apparent that when facing such a complex task, people will try to pick up and organize information around their own

biases. Just as central is what the judges ignored in their write-ups.

It would be very difficult for an automated summarizer to decide what should be covered in the summary. For D.106.E.200.G, two machine generated extracts are as follows:

System 19: *1944 Mandela joins the African National Congress at age 26, later becoming president of the group's Youth League.*

When the National Party comes to power in 1948, the ANC begins planning a campaign of civil disobedience to fight the party's apartheid policy.

January, 1962

In August he is charged with incitement and leaving the country illegally; gets five-year jail term.

July, 1963

Mandela, seven others are charged with sabotage and conspiracy to overthrow the government .

April-June, 1964

Mandela and seven others are sentenced to life in prison .

February, 1990

De Klerk lifts ANC ban and says Mandela will be released.

On Feb. 11, the nation's best-known black leader becomes a free man at last .

World leaders welcomed South Africa's announcement Saturday of Nelson R.

Mandela's impending release from prison, and international rejoicing began to build for an event so long awaited by so many.

Mandela has embraced Yasser Arafat, chairman of the Palestine Liberation Organization, in public and has compared the struggle of Palestinians to that of black South Africans.

Mayor David Dinkins: ``Our honored guest ... is a George Washington, a Bolivar, a King, a Herzl.

System 21: *A 3-year campaign that has succeeded in getting Northeastern University to award an honorary degree to jailed South African nationalist leader Nelson Mandela has raised objections from some faculty and students.*

A representative of the African National Congress said Saturday the South African government may release black nationalist leader Nelson Mandela as early as Tuesday.

Four imprisoned African National Congress leaders, convicted in 1964 with Nelson Mandela, were allowed to visit their ANC colleague at his isolated prison house, an activist said today.

The largest family group to visit black leader Nelson Mandela since he was jailed 27 years ago gathered in

Cape Town for a meeting Tuesday with the prisoner on his 71st birthday.

World leaders welcomed South Africa's announcement Saturday of Nelson R.

Mandela's impending release from prison, and international rejoicing began to build for an event so long awaited by so many.

Mayor David Dinkins warned Jewish groups against protesting the visit of anti-apartheid leader Nelson Mandela to New York City, saying the protests might insult the black community, a newspaper said today.

Unidentified assailants hurled a grenade and fired shots today at the home of relatives of Nelson Mandela, killing a baby girl and injuring her parents, police said.

For System 19, it tried to include the important events in Mandela's life and his fight against apartheid. While, for System 21, it is more likely that it picked one sentence from every article in the document set without any preference or focus.

In general, type A document sets consist of highly correlated documents forming a more coherent focus; and the type B document sets contains more diverse documents covering a wide scope of topics. This would naturally suggest that, it would be more difficult to generate quality summaries for type B document sets, since summaries would need to cover more topics. Our experiments in the next section demonstrate otherwise.

4. Comparing rankings obtained from different approaches for evaluation on the impact of document set types.

4.1. Hypothesis

As mentioned above, there are two general types of document sets in the DUC 2002 collection, which suggests that obtaining good summaries for type B document sets is more difficult than for those document sets belonging to type A. If we assume that some automated summarizers can generate quality summaries for both type A and type B document sets, while others performed worse for one document set type (more likely type B); then the ranking orders of summarizer performance by human assessors should be different, when they work with document sets that belongs to different types. There is a possibility that all summarizers performed similarly good or bad when working with certain document sets, however, from our examination of the automatically generated summaries, we believe this possibility to be low.

4.2. Experimental procedure

Three different approaches have been used to rank the performance of summarizers. We first ranked the 10 automated summarizers based on two automatic approaches. 1) Exact sentence matching, where system generated extracts were compared to model summaries created by human (which are also included in DUC 2002 data). 2) Document graph comparison, where automatically generated extracts are compared with the original documents based on the document graphs that generated, and average *F-scores* were calculated for comparison and ranking purposes [13].

We also established a ranking directly based on human judgments on the quality of the summaries, other than just comparing machine generated summaries with human created model summaries. Five people participated in the experiment. They are all graduate students pursuing their PhDs in computer science. One of them is a native English speaker, while the other four participants are from Asia and north Africa. Although English is a second language for four participants, there should have been little or no problem for them at understanding general news reports in English since they have been living in the U.S. and studying towards advanced science degrees for at least 5 five years.

Each participant was given 4 document sets, which were randomly picked from the 59 document sets in the 2002 DUC data collection; together with extracts generated by 10 automated summarizers for each document set. The summarizers have been numbered by DUC as systems 16, 19, 20, 21, 22, 24, 25, 28, 29, and 31 (Table 2). To avoid possible bias, the extracts have been renumbered. For example, in document set 61, summary number 1 is generated by System 22, but in document set 62, summary number 1 is generated by System 20. In total, 20 document sets have been evaluated in the experiment, within which 11 belong to type A and 9 belong to type B as identified.

The participants also received instruction on how to evaluate the extracts. They were asked to read carefully through the original document sets and identify the information they think that should be covered in a summary. After reviewing extracts, they should assign a score to each of the extracts using a five-point-scale scoring system, in which from 1 to 5, the quality of the summary would be from very poor to very good, respectively. They were instructed that for a good summary, its quality is based on how well it covers important information, while the order of the

sentences and the transition between two sentences in the extracts are **not** important factors. The time spent on each document set was 53.2 minutes on average (ranging from 31 minutes to 257 minutes, depending on the length the documents).

The rankings obtained based on the different approaches have been compared by using the Spearman rank correlation coefficient (r_s) [6].

4.3. Results

Three different ranking approaches have been applied and compared in our experiments. Sentence matching compares the machine generated summaries and the model summaries created by human; *DG* approach measures the information coverage of the extracts; and human evaluation is based on direct human judgment.

The ranking results are shown in Table 2, and the correlation data for different ranking approaches are shown in Table 3. Different ranking approaches ended up with different results, only the ranking from sentence matching and human judgment for type A document sets demonstrated that they are highly correlated.

For type A document sets, sentence matching and human judgment gave very similar ranking results ($r_s = 0.92$, $p < 0.01$), which were different from the *DG* approach results. This showed to us the difference among the three ranking approaches. Both sentence matching and human judgment involve human opinion on what is important, while the *DG* approach simply measures the information coverage. However, when applied to type B document sets, the correlation between sentence matching and human judgment was no longer statistically significant (Table 3), which may due to the reason that summarization for type B document sets is a more complex task and, hence, it is more difficult to reach agreement on which summarizers performed better.

We assumed that some summarizers could generate quality summaries for both types of document sets while others could have a fair performance for type A document sets but not for type B sets. Thus, the ranking orderings for the summarizers would be different when evaluated based on different types of document sets. The experiments showed that the *DG* approach indicated that there is a significant difference in summarizer performance when working on type A vs. type B document sets. The r_s of the ranking orders obtained when working with type A and type B document sets is only 0.45, and the correlation is not statistically significant (Table 4). In other words, at least some

summarizers performed differently with two different types of document sets. However, the sentence matching approach could not identify this difference as clearly as the *DG* approach, the correlation coefficient for the two rankings being 0.72, and the *p*-value indicated the correlation is statistically significant. Surprisingly, the human also failed to realize the difference ($r_s = 0.75$ and $p = 0.006$) (Table 4).

Table 2. Ranking order obtained based on three different approaches.

DG: document graph
S: Sentence matching
H: Human judge scoring

System	Type A			Type B		
	DG	S	H	DG	S	H
16	8	9	9	6	8	10
19	3	1	2	4	6	3.5
20	6	4	4	5	4	5
21	4	3	3	3	1	1.5
22	9	10	10	10	10	9
24	7	2	1	1	3	3.5
25	5	7	7.5	8	9	6
28	10	5	7.5	9	2	8
29	2	6	5	2	5	7
31	1	8	6	7	7	1.5

Although we can not completely rule out the possibility that all 10 automatic summarizers performed similarly well with type A document sets but not the type B sets, the low correlation between the two ranking orders for different types of document sets based on *DG* still supported our hypothesis.

Previous studies showed that human judgment differences could be one of the variations that affect the performance scores [3]. In our experiments, five participants generally agreed with each other on the performance of the automatic summarizers, where in pair-wised comparisons, 7 out of 10 pairs showed the correlation between two human judges was statistically significant (Table 5). Also, in this experiment, not only were the extracts provided by 10 summarizers included, but also the model summaries. In fact, the model summaries generated by humans were always ranked as the best ones, except in one case where it ranked as the third best.

Table 3. Correlation between the ranking obtained based on three different approaches.

Type A			
	DG	S	H
DG	-		
S	0.21	-	
H	0.45	0.92**	-
Type B			
	DG	S	H
DG	-		
S	0.54	-	
H	0.48	0.44	-

* $p < 0.05$

** $p < 0.01$

Table 4. Correlation between the ranking orders obtained by the same method on different types of document sets (type A and B).

DG	S	H
0.45	0.72**	0.75**

Table 5. Pair-wise correlation analyses on ranking orders obtained based on the scores assigned from different judges (The experiments were performed on all data, including both type A and type B document sets).

Judge	1	2	3	4	5
1	-				
2	0.37	-			
3	0.80*	0.39	-		
4	0.60*	0.62*	0.64*	-	
5	0.92**	0.45	0.67*	0.57*	-

In our user evaluation experiments, in trying to eliminate possible bias, we told the participants that a quality summary should cover important content, but deliberately avoided directing what kind of content should be considered as important (is it high level analysis? or is it detailed information? etc.). Unfortunately, this also resulted in the fact that no specific task was assigned for the summarization process. In this situation, human judges might assign

the same scores to the summaries that covered some general topics and detailed information, although their coverage and focuses could be quite different. In evaluations using model summaries, it is a common practice to use summaries created by several judges, such as in DUC 2002 data collection. Here the model summaries have been created by 9 judges with each document set having 2 model summaries. Then, the individual differences among the judges might be masked by the compensation from different people.

A further analysis revealed an interesting result. We calculated pair-wise correlations among five human judges separately on type A and type B document sets. When working with type B sets, there was only 1 pair in a total of 10 that correlated with each other at a level of significant, statistically, with average r_s 0.19; for type A sets, there were 4 pairs rankings showing statistically significant correlation, and the average r_s increased to 0.47. It suggested that, as individuals, the human judges had more disagreements with each other when working with diverse document sets, type B sets; while when they worked as a group, this difference is not obvious (Table 4).

5. Related Work and Conclusion

The process of automatic summarization can be decomposed into three steps: analyzing the input text, extracting important information, and synthesizing an appropriate output. In [15], Spark-Jones pointed out that for automatic summarization “our methodology should be on *context factors*”, where *purpose factors* are the most important ones among the *context factors*. The *purpose factors*, which are categorized under *situation*, *audience*, and *use*, are actually directly linked to the user and the user’s information needs. For user- and task-oriented summarization, a step of modeling the user and analyzing the user’s summarization needs should be added, and which is also the greatest challenge we are facing today.

Various technologies have been tried in generating summaries, such as term frequency [4], predefined templates [8], and latent semantic indexing [2]. However, less has been done in identifying how individual differences affect the perception of a good/quality summarization. There has been a closely related effort, which used the utility of query biased summaries to help users identify relevant documents [12], in which Local Context Analysis (LCA) has been used to expand topics contained in the baseline with additional words and phrases. LCA is a technique for automatic query expansion using pseudo feedback. It examines the context surrounding

the topic terms in the top ranked documents for query expansion. It has been found that users could judge the relevance of documents based on their summaries, almost as accurately as if they accessed the full texts. Sakai and Masuyama [11] proposed an interactive approach for multi-document summarization realizing a user’s summarization need. Their system extracts keywords from a document set and shows k best keywords with scores to a user on the screen. Then the user has the opportunity to select those that reflect his/her information needs. The approach helped improve the system performance.

Our results showed that the *DG* approach can automatically identify if a summary is created from a broad and diverse document set (as opposed to a highly focused set). Human evaluators, as a group, were not very sensitive to this difference; however, as individuals, they had more disagreement with each other when they working with document sets covering diverse content. Probably more important, the difference between these two types of document sets actually reflects the requirements for a good summary, which should have a common core of the most general topic and various pieces of information from different aspects or details (branches) that are relevant to the core. Thus, for user-centered summarization, the task is to identify what kind of “branching” information needs to be retrieved and kept in the summary to satisfy the specific user interests according to the user’s knowledge and the specific tasks he/she has.

McKeown et al [7] found that for the multi-event input document sets, the difficulty in generating a quality summary comes from the breadth and diversity of the documents in the sets. Sometimes, even humans seem to have a hard time determining how to produce a good summary. In this situation, their summaries were often quite different from each other. The method proposed to handle this situation is to apply different strategies/techniques on different categories of documents in the input sets. For DUC 2002 data, they used four different strategies, one for single events, one for multiple related events, one for biographies and one for discussion of an issue with related events [7]. The problem associated with this approach is that, the categories are classified based on the content of the documents. There should be more categories than the four being used in DUC 2002 collection, and a method is needed to automatically classify documents (McKeown et al. suggested create document sets and define a set of criteria by automatically filtering and clustering large online data [7]), while it is a manual process in DUC 2002.

Our *DG* approach has the potential to overcome this problem. *DG* generation is a process of information extraction and representation. As a result, the important concepts, entities and relations in the text are captured. We can then generate a summary from each *DG* of each document. First, a core (also in *DG* format) that contains the most general information is needed. For document sets, the core can be constructed by majority vote; while for a single document, the core can be created based on the weights of the relations. We then expand the core *DG* by inserting relevant relations based on the underlying graph structure. The most relevant relations would be decided with the help from a user model that captures a user's knowledge and foci/interests. Finally, we generate a summary based on this *DG*, which should be biased towards a user's individual interests, and better meet his/her information needs. We are currently pursuing this effort and are focused on formally defining the appropriate graph theoretic measures for expanding *DGs* from multiple documents.

Acknowledgement. This work was supported in part by the Advanced Research and Development Activity (ARDA) U.S. Government. Any opinions, findings and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the U. S. Government.

6. References

- [1] Elhadad, N. "User-Sensitive Text Summarization", AAAI Doctoral Consortium, San Jose, CA. pp. 987-988, 2004
- [2] Gong, Y. and Liu, X. "Generic Text Summarization Using Relevance Measure and Latent Semantic Analysis". In *Proceedings of the 24th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval*. pp 19-25. New Orleans, Louisiana. 2001.
- [3] Harman, D., and Over, P. "The effects of human variation in duc summarization evaluation". In *Text summarization branches out workshop at ACL'2004*. 2004.
- [4] Luhn, H.P. "The automatic creation of literature abstracts". *IBM Journal*, pp 159-165, 1958
- [5] Mani, I. and Maybury, M.T. "Advances in Automatic Text Summarization". The MIT Press. 1999
- [6] Myers, J. L. and Well, A. D. "Research Design and Statistical Analysis". Lawrence Erlbaum Associates, publishers, New Jersey pp. 488-490, 1995
- [7] McKeown, K., Barzilay, R. and Blair-Goldensohn, S. "The Columbia Multi-Document Summarizer for DUC 2002". 2002
- [8] McKeown, K and Radev, D.R. "Generating Summaries from Multiple News Articles". In *Proceedings of the 18th Annual ACM SIGIR Conference*, pp. 74-82. Seattle, WA. 1995
- [9] Montes-y-Gómez, M., Gelbukh, A., and López-López, A. "Comparison of Conceptual Graphs". In *Proceeding of MICAI-2000*, In *1st Mexican International Conference on Artificial Intelligence*. Acapulco, Mexico, 2000
- [10] Over, P. and Liggett, W. "Introduction to DUC-2002: an Intrinsic Evaluation of Generic News Text Summarization Systems". Document Understanding Conference website (<http://duc.nist.gov/>), 2002
- [11] Sakai, H. and Masuyama, S. "A Multiple-Document Summarization System Introducing User Interaction for reflecting User's Summarization Need". Working Notes of NTCIR-4, Tokyo, 2-4 June 2004.
- [12] Sanderson, M. "Accurate User Directed Summarization from Existing Tools". In *Proceedings of the 7th International Conference on Information and Knowledge Management*. Bethesda, Maryland. pp 45-51, 1998
- [13] Santos, E. Jr., Mohamed, A.A and Zhao, Q. "Automatic Evaluation of Summaries Using Document Graphs". In *Proceedings of ACL 2004, Workshop on Text Summarization Branches out*, Barcelona, Spain, pp 66-73, 2004
- [14] Sleator, D. D. and Temperley, D. "Parsing English with a link grammar". In *Proceedings of the 3rd International Workshop on Parsing Technologies*. pp 277-292, 1993
- [15] Sparck-Jones, K. "Automatic summarizing: factors and directions. In *Advances in automatic text summarization*. Mani and Maybury, 1998