A Survey for Multi-Document Summarization

Satoshi Sekine

New York University
715 Broadway, 7th floor
New York, NY, 10003, USA
sekine@cs.nyu.edu

Chikashi Nobata

Communications Reserach Laboratory
2-2-2 Hikaridai, Seika-chou, Soraku-gun
Kyoto, 619-0289, Japan
nova@crl.go.jp

Abstract

Automatic Multi-Document summarization is still hard to realize. Under such circumstances, we believe, it is important to observe how humans are doing the same task, and look around for different strategies.

We prepared 100 document sets similar to the ones used in the DUC multi-document summarization task. For each document set, several people prepared the following data and we conducted a survey.

- A) Free style summarization
- B) Sentence Extraction type summarization
- C) Axis (type of main topic)
- D) Table style summary

In particular, we will describe the last two in detail, as these could lead to a new direction for multi-summarization research.

1 Introduction

Automatic Multi-Document summarization is still hard to realize. Like single document summarization for newspaper articles, where we don't have a notably better automatic summarization algorithm than a simple lead based method, automatic multi-document summarization faces very difficult challenges. Under such circumstances, we believe, it is important to observe how humans are doing on the same task, and look for possible different strategies.

Assume you are given several documents talking about the same topic, and are asked to summarize them, what might you do. The authors tried this by themselves. First we used a marker to mark the important phrases or sentences. Then we tried to connect them, in some cases by figuring out the main or common topics in the marked sentences, or in some cases, by making a list or a table to figure out the overview of the documents. When we looked at the result at this stage we noticed that these are very good summaries, even if they are not summaries in the conventional sense (a set of sentences to be read). The main topics are good to

understand the overall issues in the document set and the table is a good digest of the issues throughout the document set. If we can automatically create such data from document sets, we might be able to make a good summary. The questions arising here are what kinds of "main topic" we can make in general, and what percentage of document sets are suitable for table-style summarization.

The main topics we created in our hand summary experiment were like lists of keywords, but we found that there are more general types like "these documents are talking about a single person". As keyword extraction has been one of the techniques in summarization, we will focus on the types of the main topics in the following experiments.

We will describe the definition of our types and report on the experiment of manually creating table-style summarization, as well as analyses of free style summaries and sentence extraction type summaries. We prepared 100 document sets similar to the ones used in the DUC multi-document summarization task (DUC homepage). For each document set, annotators prepared the following data.

- A) Free style summarization
- B) Sentence Extraction type summarization
- C) Axis (type of main topic)
- D) Table style summary

In particular, we will describe the last two in detail, as these could lead to a new direction for multi-document summarization research.

2 Document Sets

First, we describe how we accumulated our 100 multi-document data. We found that the topics of DUC multi-document data are a bit biased as it is pre-filtered for evaluation purposes, i.e. DUC document sets are carefully chosen as described in the guidelines. The pre-filtering is useful for evaluation purposes, but it does not necessarily reflect the distribution of user needs or distribution of topics in the news. We would like to obtain relatively more balanced document sets. We

adopted the procedure described in the following, where the entire experiment was done using a Japanese newspaper corpus (Mainichi 1998 and 1999).

- Select an article randomly from the corpus (seed)
- Choose keywords from each article. Keywords are all nouns of frequency more than 1, except for some special types of nouns
- Use dice coefficient to retrieve articles similar to the seed article. Gather all documents that have coefficient more than 0.5.
- Select article sets that have more than 3 articles. About 300 such sets are obtained and among them, we selected 100 document sets, preferring more documents in a set and avoiding overlapping topics.

The average number of articles in a document set was 4.7 and the average number of sentences in a document was 12.9. Annotators read the articles in each set and detected if there were articles that are different from the topic throughout the document set. Such articles, which turned out to be very few in number, were excluded in the following experiments.

3 Task and annotator

We have four tasks and three annotators (indicated by a number). Annotator 1 and 2 did the same task, but annotator 3 did only a part of it. All of them have college degrees, in particular annotators 1 and 2 are Japanese native speakers and have majors in linguistics at US universities.

Some examples (free summaries for one document set, and axes and table data for three sets, all translated into English) are shown in the appendix.

Task	Annotator		
	1	2	3
Free style summary	100	40	
Sent. Extraction	100	20	100
Axis	100	100	100
Table summary	100	100	

Table 1. Task and annotator

4 Free style summarization

The first task is a free style summarization. The inter-annotator agreement based on the word vector metric adopted by TSC evaluation (TSC homepage) is calculated. This is a cosine metric of $tf*idf$ measure of the words in the summaries. Most of the pairs (37 sets out of 40 sets) had values of 0.5 or more, which is much

larger than that of automatic systems measured against the human made summaries in TSC-1 (ranging around 0.4 in 10% summary, 0.5 in 40% summary). We can reasonably believe the summaries are very reliable.

5 Sentence Extraction

Now we will look at the summarization by sentence extraction. Annotators 1 and 3 conducted the task for the entire data, so we will compare the results of those two. We asked the annotators to extract about 20% of the sentences as a summary of each document set, but the actual numbers of extracted sentences are slightly different between the two. Table 2 shows the number of sentences selected by the two annotators with inter-annotator agreement data. The number of sentences selected by both annotators (533) looks low, compared to the number of sentences selected by only one annotator (650 and 746). However, the chi-square test is 513.9, which means that these two results are strongly correlated (less than 0.0001% chance).

annotator 3	Annotator 1		Total
	selected	not selected	
selected	533	746	1279
not selected	650	4050	4700
Total	1183	4796	5979

Table 2. Number of selected sentences

6 Axis

Axis is based on the idea of (McKeown et al. 2001). They defined 4 categories of document sets based on the main topic of the document set for the purpose of using different summarization strategies (they actually used two sub-systems), shown in Table 3.

Category and Description
Single-Event (2) The documents center around one single event at one place and at roughly the same time, involving the same agents and actions
Person-centered (10) The documents deal with one event concerning one person
Multi-Event (7) Several events occurring at different places and times and usually with different protagonists, are reported together
Other (11) Document sets contain even more loosely related documents

Table 3. McKeown's categories

The number in brackets for each category indicates the number of document sets in the DUC 2001 training data. As can be seen, the number of “person centered” sets is quite high. We believe this is due to the pre-filtering in the DUC data. “Other” is also high, which means more categories may be needed.

We created new categories based on our study of document sets (other than the 100 sets reported here). We defined 13 categories, shown in Table 4, for what we will call the axis of the document set.

Single-person	Multi-Person
Single-location	Multi-location
Single-organization	Multi-organization
Single-facility	Multi-facility
Single-product	Multi-product
Single-event	Multi-event
Others	

Table 4. 13 Axes

The axis is a combination of two types of information; single or multi, and 6 kinds of named entities (person, location, organization, facility, product and event). “Single” means that all the articles are talking about a single event, person or other entity, whereas “Multi” articles are talking about multiple entities that might participate in similar types of events. We used 6 categories of entity types, which are the major categories defined in the MUC (Grishman and Sundheim 1996) or ACE project (ACE homepage). For example, if a document set is talking about Einstein’s biography, it should be tagged as “single-person”, and if a set is talking about earthquakes in California last year, it should be tagged as “multi-event”.

In order to demonstrate the validity of the categories, we tried to categorize the training data of DUC 2001’s multi-document sets into our categories. Two people assigned one or two categories to each set. We allow more than one axis to a document set, as some document sets should be inherently categorized into more than one axis. If we consider only the first choices, the inter annotator agreement ratio is 80% and if we include the second choices, the ratio is 93.3%. We believe the categorization is practical. Table 5 shows the distribution of axis categories tagged on our 100 data sets by three annotators. Note that annotators 1 and 2 assigned more than one axis to some data sets, so the totals exceed 100.

All the categories except multi-facility are used by at least two annotators. Because the axis “other” is used rarely, the set of axes are empirically found to have quite good coverage.

The inter-annotator agreements are 55, 61 and 67 % among the three annotators. Although the ratios are

lower than that on the DUC data (we believe this is because of the pre-filtering of document sets), the agreement is still high at 55-67% even though there are 13 kinds of axis. Note that chi-square test is not suitable to measure the data because of the data sparseness.

Axis	Annotator		
	1	2	3
s-event	36	32	37
s-facility	3	1	0
s-location	7	2	4
s-organization	6	12	11
s-person	12	13	12
s-product	14	14	8
m-event	15	30	8
m-facility	0	0	0
m-location	2	3	2
m-organization	9	12	13
m-person	2	7	1
m-product	2	2	0
Other	1	0	3

Table 5. Distribution of axis

There are 39 document sets that have the same axis assigned by the three annotators, and there are only 7 document sets that have three different axes by three annotators (no overlap at all). Even when different categories are tagged, sometimes all of them are understandable and we can say that these are all correct. So for some document sets, more than one category is instinctively correct. This result indicates that, for some large percentage of document sets, it is possible to assign axis(es). We believe for summarizing those document sets, knowing the axis before summarization could be quite helpful. We are seeking a method to automate the process of finding the axis(es).

7 Table

A table is a good way to summarize a document set talking about multiple events of the same type, a collection of similar events or chronological events. We asked annotators to make a table for each document set. Table 6 shows some statistics of the created tables. The average number of columns is 3.47 and 5.25 for annotator 1 and annotator 2, respectively. Regarding comparison between tables, the percentages of complete overlap (relationship of columns is 1 to 1 and the same information is collected in the column) are 58% and 38%. The percentages of overlap (relationship of columns is not 1 to 1, but the information of the columns is overlapping between the tables) are 94% and 70%. We can see that annotator 1 made fewer columns than annotator 2, and most columns made by annotator 1

overlap columns made by annotator 2. So the difference is probably due to the fact that annotator 2 made more detailed tables. (As this is the first such survey, it was not easy to create good instructions) In other words, it might be the case that most important information (which was turned into columns) is simultaneously found by the two annotators.

	Annotator 1	Annotator 2
Ave. num. of column	3.47	5.25
Complete overlap	58%	38%
Overlap	94%	70%

Table 6. Statistics of created tables

When we compared the tables created by the two annotators one by one, we categorized the results into 5 categories.

- A) Two tables are completely the same
- B) The information in the tables is the same, but the way of segmenting information into columns is different. For example, one of the tables has a column “visiting activity (of a diplomat)” including information about visiting place, person and purpose, whereas the other table has columns “visiting place”, “the person to meet” and “purpose of the visit”.
- C) Missing one or two columns from either or both of tables (in total). This means one of the tables has one or two fewer columns and the information in the columns is not mentioned in the other table. As we can guess from Table 6, most of the missing columns were found in tables of annotator 1.
- D) Missing more than two columns from tables.
- E) The two tables are completely different in structure, because of the table creator’s different point of view.

Table 7 shows the result of this survey.

Description	Num. of sets
A) Same table	8
B) Only segmentation	15
C) Missing one or two column	34
D) Missing more than two column	17
E) Completely different table	26
Total	100

Table 7. Comparison of tables

There are only a small number of document sets (8) from which the annotators made completely the same table. However, for more than half the document sets, the tables created by the two annotators are quite similar (including “same table”, “only segmentation” and “missing one or two columns”). This is complementary to the result shown in Table 6; for many document sets, the tables by annotator 2 have additional information compared to the tables by annotator 1.

We also asked the annotators to judge if each document set is suitable to summarize into a table. We made three categories for the survey.

- A) Table is natural for summarizing the document set
- B) Information can be summarized in table format
- C) Table is not suitable to summarize the document set

The result for the two annotators is shown in Table 8. Annotators 1 and 2 judged 40 and 45 sets to be suitable for a table, 36 and 38 are OK and 24 and 17 are not suitable. This is an interesting result - that for so many document sets (40-45%) a table is judged to be natural for summarizing. Compared to that, only a smaller fraction (17-24%) are judged unsuitable. The relationships between the two annotators’ judgments are also shown in Table 8. The Chi-test is 17.94 and the probability is 0.13%; that means that the two judges are highly correlated.

Annotator 2	Annotator 1			total
	A	B	C	
A	28	12	5	45
B	9	16	13	38
C	3	8	6	17
total	40	36	24	100

Table 8. Suitability of table

8 Discussion

We reported a survey for multi-document summarization. We believe the results are encouraging for the pursuit of some novel strategies of multi-document summarization.

One of them is the notion of axis. As we observed that for some percentage of the document sets, the axis can be tagged with some certainty, we might be able to make an automatic system to find it. Once the axis is correctly found, it might be useful for multi document summarization. For example, if a set is “single-person” then the summary for the set should be centered on the person. This may suggest, for example, generating a

summary of type ‘biography’ (Mani 2001). If a document set is found to be “multi-event”, then the summary should focus on the differences of the events.

The other result found in the experiment is that a quite large percentage of document sets can be summarized in table format. As this is a preliminary experiment, there is incompleteness in the instruction and we believe further study on this topic is necessary. In addition to setting guidelines for the degree of detail, the style of cell contents shall be more uniform. Currently, cells contain words, phrases and sentences. We believe that by making more careful instructions for annotation, the comparison between different tables can be more systematized. In other words, a systematic evaluation may be possible.

9 Future Work

Obviously, the future work suggested by these results includes automatic methods to find what the human found in this experiment.

We have started finding the axis automatically by observing the distribution of named entities, words and phrases.

Once we are able to find the axis and suitability of table summary automatically for a given document set, the next stage of research will involve using this information to select an appropriate way to summarize the set, i.e. table summary, sentence extraction summary or summary including rewriting. This will be extended if the document set is created dynamically from a user’s query. The type of query could be a helpful clue in selecting the way to summarize the retrieved document set.

The technology to summarize a document set in table format is studied in Information Extraction. However, it has a hard limitation that the topic of the document set has to be known in advance and the knowledge to build table has to be created by hand, which usually takes a long time. There have been efforts to automate the knowledge creation (Riloff 1996) (Yangarber 2000) (Sudo 2001); we hope to make a bridge between such automatic IE knowledge discovery and automatic summarization efforts.

10 Acknowledgements

This research is supported by the Defense Advanced Research Projects Agency as part of the Translingual Information Detection, Extraction and Summarization

(TIDES) program, under Grant N66001-001-1-8917 from the Space and Naval Warfare Systems Center, San Diego, and by the National Science Foundation under Grant IIS-0081962. This paper does not necessarily reflect the position of the U.S. Government. We would like to thank our colleagues at New York University, who provided useful suggestions and discussions, including, Prof. Ralph Grishman, Mr. Kiyoshi Sudo and Mr. Yusuke Shinyama. Also, we thank the three annotators to do the tedious job.

References

- (ACE-homepage)
<http://www.nist.gov/speech/tests/ace/index.htm>
- (DUC-homepage)
<http://www-nlpir.nist.gov/projects/duc/>
- (TSC-homepage) <http://lr-www.pi.titech.ac.jp/tsc/>
- (McKeown et al. 2001) K.R. McKeown, R. Barzilay, D. Evans, V. Hatzivassiloglou, M. Yen Kan, B. Schiffman, S. Teufel, “Columbia Multi-Document Summarization: Approach and Evaluation”, Proceedings of the Document Understanding Conference (DUC-2001), 2001
- (Grishman and Sundheim 1996) R. Grishman, B. Sundheim, “Message Understanding Conference - 6: A Brief History”, Proceedings of the 16th International Conference on Computational Linguistics (COLING ’96), 1996
- (Mani 2001) I. Mani, “Automatic Summarization”, John Benjamins Publishing Company, 2001
- (Riloff 1996) Ellen Riloff, “Automatically Generating Extraction Patterns from Untagged Text”, Proceedings of the 13th National Conference on Artificial Intelligence, 1996
- (Sudo et al. 2001) Kiyoshi Sudo, Satoshi Sekine and Ralph Grishman, “Automatic Pattern Acquisition for Japanese Information Extraction”, Proceedings of Human Language Technologies (HLT ’01), 2001
- (Yangarber 2000) Roman Yangarber, Ralph Grishman, Pasi Tapanainen and Silja Huttunen: “Automatic Acquisition of Domain Knowledge for Information Extraction”, Proceedings of the 18th International Conference on Computational Linguistics (COLING 2000), 2000

Appendix

Original data is all in Japanese. The example data here is English translation.

Sample Data (Article Set #mai9899.189565)

<Axis> Annotator 1: s-event, s-person, Annotator 2: s-person, s-event, Annotator 3: s-person

<Summary by annotator 1>

Mr. Ikuo Kashima, 69, a yacht sailor in Hannnann-shi, Osaka-fu, departed a yacht harbor in Misaki-machi, Osaka, on September 15, with his yacht "Koraasa 70" to achieve a solo sailing around the world without calling at any port as a world's oldest challenger. The original plan was to sail across the Pacific, the Atlantic, and the Indian Ocean in about 330 days, but he called at Honolulu, Hawaii, on October 26 due to the breakdown of fresh-water-generator. He repeatedly departed from and returned to Honolulu because of the engine and generator trouble on November 16, the leak and the breakdown of the helm in December, and he gave up on his plan after the fourth trial on January 2. In July 99, he left Misaki-machi to make another attempt to sail around the world, and the news of passing through the Cape Horn, the south most point in South America came on December 30. He is expected to arrive at the harbor in Misaki-machi next July if the voyage goes smoothly.

<Summary by annotator 2>

To achieve a solo sailing around the world without calling at any port by the world's oldest challenger, Mr. Ikuo Kashima left Tannowa Yacht Harbor in Misaki-machi, Osaka-fu, on September 15, 1998. Around noon on October 26, however, he had to call at Honolulu Harbor on Oahu Island, Hawaii, to repair the fresh-water-generator in his yacht "Koraasa 70," which he was using for this voyage. Although he changed his plan and made Hawaii his starting point for his around-the-world solo sailing and started sailing again on October 31, he had to return to Honolulu again on November 16, due to the breakdown of the engine and the generator. Another trial started on December 4, again, did not succeed. Mr. Kashima made his fourth trial on January 2, but again, he had to return to Honolulu on January 11 for the leak in the yacht and the trouble in the helm this time. Later, Mr. Kashima, aiming at a solo sailing around the world again, left Tannowa Yacht Harbor on July 20. On December 30, he succeeded in passing through the most difficult navigation point, Cape Horn. He is planed to be back to Misaki-machi, Osaka, in July 2000.

<Table summary by annotator 1> Suitability for table: Suitable

DATE	PLACE	EVENT	PURPOSE
9/15/98	Misaki-machi, Osaka	Departure	A non-stop solo sailing around the world
10/26	Honolulu, Hawaii	Calling	To repair fresh-water-generator
10/31	Honolulu, Hawaii	Departure	Second trial to sail around the world
11/16	Honolulu, Hawaii	Calling	Troubles in the engine and the generator
12/4	Honolulu, Hawaii	Departure	Third trial to sail around the world
?	Honolulu, Hawaii	Calling	Leak and the breakdown of helm
1/2/99	Honolulu, Hawaii	Departure	Fourth trial to sail around the world
7/20	Misaki-machi, Osaka	Departure	Another trial for the sailing around the world
12/30	Cape Horn, the south most point in South America	Passing	

<Table summary by annotator 2> Suitability for table: Suitable

DATE	PLACE	ACTION	CAUSE
9/15/1998	Tannowa Yacht Harbor, Misaki-machi, Osaka-fu	Departure	To achieve a non-stop solo sailing around the world as the world's oldest challenger.
10/26/1998	Honolulu, Oahu Island, Hawaii	Calling	To repair fresh-water-generator in "Koraasa70"
10/31/1998	Honolulu, Oahu Island, Hawaii	Departure	To challenge a non-stop solo sailing around the world starting from Hawaii.
11/16/1998	Honolulu, Oahu Island, Hawaii	Calling	To repair "Koraasa 70"'s engine and generator.
12/4/1998	Honolulu, Oahu Island, Hawaii	Departure	To aim at a non-stop solo sailing around the

			world starting from Hawaii.
1/2/1999	Honolulu, Oahu Island, Hawaii	Departure	To aim at a non-stop solo sailing around the world starting from Hawaii.
1/11/1999	Honolulu, Oahu Island, Hawaii	Calling	For the leak and the troubles in the helm of "Koraasa 70."
7/20/1999	Tannowa Yacht Harbor, Misaki-machi, Osaka-fu	Departure	To achieve a non-stop solo sailing as a world's oldest challenger
12/30/1999	Cape Horn	Passing	
July 2000	Tannowa Yacht Harbor, Misaki-machi, Osaka-fu	Expected to arrive	Arrival from the non-stop solo sailing around the world as an oldest challenger.

Sample Data (Article Set #mai9899.106495)

<Axis> Annotator 1: s-person, Annotator 2: s-person, Annotator 3: s-person

<Table summary by annotator 1> Suitability for table: Not suitable

DATE	COHEN, DEFENSE SECRETARY
3/27/98	Meeting with Mordechai, Israeli Minister of Defense. Promised to cooperate in Israel's missile defense development.
11/6/98	Visited Middle Eastern countries. Asking Iraq to withdraw its suspension on cooperating UNISCOM
1/11/99	Visited Japan. Discussed the matters on US-Japan security guideline-related bills.
2/22/99	Planned to visit Japan and China in April. (The plan was postponed due to the prolonged air raid in Yugoslavia.)
7/9/99	Planned to visit Japan and South Korea in late-July.

<Table summary by annotator 2> Suitability for table: Suitable

DATE	PLACE	MEETING WITH	PURPOSE	TIME OF VISIT
3/27/98	Israel	Mordechai, Israeli Minister of Defense	To cooperate in Israel's missile defense development	3/98
11/6/98	Middle East		To discuss the issues on Iraq	11/98
1/7/99	Japan	Komura, Japanese Foreign Minister Noroda, Japanese Chief of Defense	To discuss the situation in Korean Peninsula and an outlook on passing US-Japan security guideline-related bill in Japan's Diet	1/11-1/14/98
1/7/99	South Korea			1/99
2/2/99	Japan		To explain about the US-China defense cooperation and to discuss the Guideline-related bill and its approval in the Diet.	4/99
2/2/99	China			4/99
7/9/99	Japan	Obuchi, Japanese Prime Minister	To discuss the issues on North Korea's ballistic missile and on relocation of Futenma Airport in Okinawa.	Late July
7/9/99	South Korea			Late July

Sample Data (Article Set #mai9899.141141)

<Axis> Annotator 1: s-location, Annotator 2: m-event, Annotator 3: s-product

<Table summary by annotator 1> Suitability for table: Suitable

DATE	EVENT	KOSOVO LIBERATION ARMY	YUGOSLAVIA/ SERBIAN GOVERNMENT	CONTACT GROUP
2/6/99	Peace Talk began	Claimed for Kosovo's independence	Disagreed to Kosovo's independence	Presented a peace plan including Kosovo's auton-

				omy
2/10	Conflict in Peace talk	Demanded Serbs to sign on agreement for immediate cease-fire	Requested Kosovo to sign the peace plan	Aimed at reaching an agreement
2/19	A day before the negotiation deadline		Disagreed to NATO's presence in Kosovo	
2/20	Agreement deadline. Talk continued.		Rejected NATO peace force's presence in Kosovo	Foreign Ministers arrived
2/23	Final day for the Peace Talk	Rejected to sign the peace plan. Requested the vote for Kosovo independence. Rejected to dissolve Kosovo Liberation Army.	Rejected to sign the peace plan, and refused to have NATO's peace force.	Acknowledged the temporary agreement on Kosovo's autonomy.
3/15	Second Peace Talk began			
3/16	Second day of Peace Talk	Expressed the willingness to sign the peace plan	Presented the disagreed items in peace plan in writing	Rejected to change the peace plan
3/18		Signed the peace plan		Requested Yugoslavia to accept the plan by 24 th .

<Table summary by annotator 2> Suitability for table: can-be

DATE	NO.	PERSON	ACTION
2/6/99	1	Representative of "Kosovo Liberation Army"	Arrived in Paris
2/10/99	1	Yugoslavian government representative and groups of Albanian residents	Each group handed in its own requirements and caused conflicts
2/19/99	1	Cook, Foreign Minister, U.K.	Arrived in Paris
2/20/99	1	Albright, Secretary of State, U.S.A Fischer, Foreign Minister, Germany	Arrived in Paris
3/1/99	2	Albanian residents	Expressed their willingness to sign the peace plan
3/16/99	2	Contact Group of US, Europe, and Russia	Pursued Serbian Republic representatives to accept the peace plan
3/16/99	2	Representatives of Serbian Republic	Presented objection to the peace plan in writing
3/16/99	2	Hill, Special US envoy	Rejected the big change in the peace plan
3/18/99	2	Representatives of Serbian Republic	Announced to sign the peace plan
3/18/99	2	Contact Group of US, Europe, and Russia	Requested Yugoslavia to accept the final peace plan by March 24.