



RENDER
FP7-ICT-2009-5
Contract no.: 257790
www.render-project.eu

RENDER

Deliverable D5.2.2

Diversified News Service

Editor:	Enrique Alfonseca, Google Inc.
Author(s):	Enrique Alfonseca, Google Inc., Mitja Trampus, JSI
Deliverable Nature:	Report (R)
Dissemination Level: (Confidentiality)	Public (PU)
Contractual Delivery Date:	September 2012
Actual Delivery Date:	September 2012
Suggested Readers:	Developers involved in the use cases.
Version:	1.0
Total number of pages:	26
Keywords:	Summarization, news, diversity

Disclaimer

This document contains material, which is the copyright of certain RENDER consortium parties, and may not be reproduced or copied without permission.

All RENDER consortium parties have agreed to full publication of this document.

The commercial use of any information contained in this document may require a license from the proprietor of that information.

Neither the RENDER consortium as a whole, nor a certain party of the RENDER consortium warrant that the information contained in this document is capable of use, or that use of the information is free from risk, and accept no liability for loss or damage suffered by any person using this information.

Full Project Title:	RENDER – Reflecting Knowledge Diversity
Short Project Title:	RENDER
Number and Title of Work package:	WP5
Document Title:	D5.2.2 – Diversified News Service
Editor (Name, Affiliation)	Enrique Alfonseca, Google
Work package Leader (Name, affiliation)	Javier Caminero, Telefonica
Estimation of PM spent on the deliverable:	15

Copyright notice

© 2010-2013 Participants in project RENDER

Executive Summary

This document describes the Google use case scenario, consisting in presenting news information to users highlighting different points of view, attitudes towards the event, and different entities (people, organizations, etc.) involved in the event reported in the news. The use case integrates different components from the RENDER project, such as

- The news crawler provided by JSI NewsFeed (D1.1.1), which provides current news clustered by topic, including the same event reported by different news agencies and publishers, and possibly reported from different points of view.
- The language analysis tools from Enrycher, from JSI, which includes the standard analysis annotations: token boundaries, sentence boundaries, part-of-speech tags, syntactic dependencies and named entity annotations.
- Summarization technologies, from Google, which allow the system to generate event summaries taking into account different scorers depending on the user's request. The possibilities that have been added to the demo are: summaries focused to certain entities, to some geographical location, or to some polarity (positive or negative).

The first version of the demo is now ready and allows users to browse current news and to explore them taking into account the diversity in the news sources. It will be further improved in the next six months in the work for deliverable D5.2.3. The evaluation, which will be reported at the end of the project in deliverable D5.2.4., is already ongoing.

List of authors

Organisation	Author
Google	Enrique Alfonseca
JSI	Mitja Trampus

Table of Contents

Executive Summary	3
List of authors	4
Table of Contents	5
List of Figures	6
Abbreviations	7
1 Introduction	8
2 Use case description	9
3 Diverse News Demo	11
3.1 User Interface	11
3.2 System Architecture	12
4 Summarization technology	14
4.1 Centrality models	15
4.2 Library usage	16
4.3 Examples	18
5 Evaluation	22
6 Conclusions and Future Work	24
References	25

List of Figures

Figure 1 Initial mock-up of the summarization user interface.....	10
Figure 2 The summarization user interface displaying a summary of articles about taxes, but with special focus on their relation to retirement and how it's viewed in the US.	11
Figure 3 The summarization interface, displaying the same set of articles as in Figure 2, this time with a focus on the European viewpoint with a negative connotation. The top articles are different due to automatic reordering.	12
Figure 4 Summarization Architecture.....	14
Figure 5 TopicSum plate diagram	16
Figure 6 Example code of how the summarizer can be used.....	18
Figure 7 Cluster of news with snippets, and summary generated from it.	19
Figure 8 Cluster of related news, and summary generated from it.	20
Figure 9 Cluster of related news articles and summary generated from it.	21
Figure 10 Evaluation template for summaries.	23

Abbreviations

KL	Kullback-Leibler
LDA	Latent Dirichlet Allocation
TAC	Text Analysis Conference

1 Introduction

Google collects articles from about 25,000 professional sources for Google News, and additionally also serves items provided by non-professional Web users such as blog items or tweets. This information is made available to our users through our products, including Web Search, Google News, Google+ (Sparks) and Google Alerts. In most of these products, users specify their information needs or preferences by means of a short textual query, that is used so as to choose the most relevant documents to be returned.

The focus of the RENDER project is the analysis and representation of diversity of opinions in the web. News articles on the web are also a source of diversity, as the same happening can be reported by different news sources from different viewpoints. We believe that users would benefit from having these viewpoints explicitly presented to them, making it easier for them to analyse the underlying events and to form an informed opinion about them. To this aim, a variety of existing technologies can be combined and further developed so that news textual sources can be analyzed and contrasted with each other, the writer's attitude and polarity towards the event can be coarsely classified, and the most relevant facts and viewpoints can be distilled from the collection of articles and presented to the user in a useful way.

This document describes the first version of the Google use case application, developed in collaboration with JSI, for presenting recent, relevant news to users and allowing them to explore the content of those news from different points of view and highlighting different attitudes towards the event. It starts by describing the use case scenario, followed by a description of the application that we have built, and continues with a detailed description of the technologies underlying it. Finally, we describe the evaluation settings and metrics we are considering to measure the success of this research, which will be part of deliverables D5.2.3 and D5.2.4.

2 Use case description

The planned work for the use case based on news was the development of a diversity-aware visualization tool for the relevant information returned to the user. For example, for a user looking for news about [libya], the current result at the moment in Google News is a set of news, with those referring to the same story grouped together, about the general topic of the country Libya. But it is difficult for a reader to understand what are the differences between the different sources, and the points of view that are expressed in those news articles. The use case tool has been designed to allow the user to focus on a particular news story (reported by different sources, describing something that has happened recently) to dive, by means of automatic text summarization techniques, into different viewpoints and attitudes.


The main requisite for this tool, specified in Deliverable D5.2.1, is the following: this visualization tool should be able to process the news that refer to the same events, and to show to the users additional ways for analyzing the information, as for example:


- What is the most relevant information about this entity for the use case of the user? For example, a user might be interested only on the number and identities of French casualties in an air accident, whereas other user might want to know about all of them. This first problem is similar to the general query-guided summarization [Dang, 2006], where a user specifies a query containing an information need, and summarizer outputs are evaluated on several dimensions, including how informative they are to the user and how readable, coherent and grammatical they are. We make the assumption that the user interests can be expressed in terms of the most relevant Named Entities mentioned in the news.
- What are the different points of view in the different news that refer to the same event at different times? Different sources may report the same data with differences in bias or sentiment, or may simply highlight and omit different characteristics of the news. There is already substantial work combining sentiment analysis [Pang and Lee, 2008] and text summarization [Mani, 1999]. Some early approaches combining summarization with sentiment include [Stoyanov et al., 2004] and [Yu and Hatzivassiloglou, 2003], which divided the task into two subtasks: first finding the most relevant sentences in the corpus given the user query, and next classifying them as either factual or opinionated, in order to keep only the opinionated ones. The Text Analysis Conference in 2008 included a competitive evaluation on opinion mining, where most of the systems approached the problem by adapting their summarizers to include in sentence scoring a polarity or subjectivity score, or by filtering out non-subjective sentences. We will use a similar approach to drive the summarizer into generating summaries of the same event, expressing different polarities.

Figure 1 displays an earlier mock-up of the user interface, as included in Deliverable D5.2.1:



- A news cluster, previously selected by the user, and containing several news articles describing the same event, is central in the user interface. The UI will allow the user to work with the news stories included in the cluster.
- A text-based summary is generated from the different documents in the cluster, initially with neutral settings: including the most relevant information in the news stories according to how central it is, i.e. how many of them support it.
- Some controls would allow the user to tune the summary according to different criteria.

News cluster:

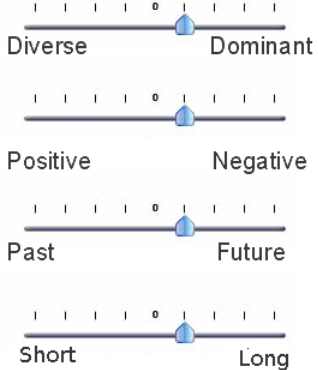


[Iran rejects Israeli charges on sending weapons to Gaza](#) 
Monsters and Critics.com - 1 hour ago
Army Commander Ataollah Salehi told IRNA that the charges made by Israel were lies and hoped that like the former leaders of **Egypt**, 'also the Zionist regime (Israel) will soon collapse.' On Tuesday, the Israeli military said that naval commandos had ...
[Iran exploiting Israel-Egypt coordination difficulties](#) Ynetnews
[Ship Bound for Egypt Is Seized by Israel](#) New York Times
[Israel Seizes Egypt Bound Ship With Arms On Board](#) RTT News
[Telegraph.co.uk - Ha'aretz](#)
[all 484 news articles »](#)

Cluster summary:

Iran rejected Israeli charges that a cargo ship intercepted in the Mediterranean was carrying weapons from Iran earmarked for Palestinian militants in Gaza, official news agency IRNA reported Wednesday. 
Army Commander Ataollah Salehi told IRNA that the charges made by Israel were lies and hoped that like the former leaders of Egypt, 'also the Zionist regime (Israel) will soon collapse.' 

Summary controls



The summary controls consist of four horizontal sliders, each with a blue knob and a '0' marker. The sliders are labeled as follows: 1. Diverse (left) to Dominant (right), with the knob positioned slightly past the center. 2. Positive (left) to Negative (right), with the knob positioned slightly past the center. 3. Past (left) to Future (right), with the knob positioned slightly past the center. 4. Short (left) to Long (right), with the knob positioned slightly past the center.

Figure 1 Initial mock-up of the summarization user interface.

As will be seen in the next sections the specific UI of this mock-up has been somewhat modified, although it remains quite similar to the original specification both in terms of appearance and functionality.

3 Diverse News Demo

The summarization component is showcased in a demo web service titled DiversiNews, available at <http://aidemo.ijs.si/diversinews/>. The service offers the user a fast overview of a set of news articles; it uses an interface similar to the one presented in the section above to allow the user to drill down into the diverse viewpoints represented by individual articles.

3.1 User Interface

See Figure 2 and Figure 3 for an example of how the interface can be used. In the figures, the user searched for news articles about taxes. A summary of results from recent news appears on top; just below it, the user is presented with individual articles that have been summarized, sorted by relevance. On the right, several panels allow the user to further specify which perspective on the news should be emphasized.

The first panel displays keywords automatically extracted from the analyzed articles. They are laid out in a way that puts related keywords closer together, segmenting the panel into “semantic regions” of sorts. By moving the red dot around, emphasis is put on the topics closer to the red dot. Similarly, the second panel displays a world map; by moving the red dot, higher importance is given to news reports written in the corresponding part of the world. Lastly, the slider at the bottom allows placing focus on news with either a positive or a negative outlook. Whenever one of the red dots or the slider gets moved, the summary is updated accordingly and the articles reordered.

Please note that the screenshots have been created using a small development database, affecting the quality somewhat.

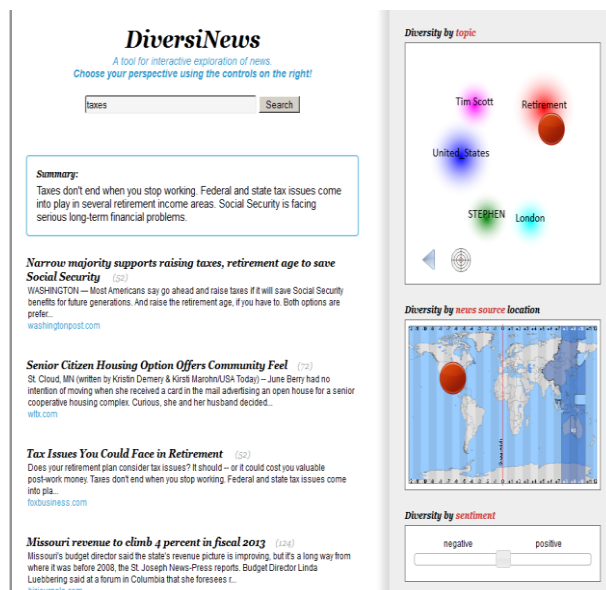


Figure 2 The summarization user interface displaying a summary of articles about taxes, but with special focus on their relation to retirement and how it’s viewed in the US.

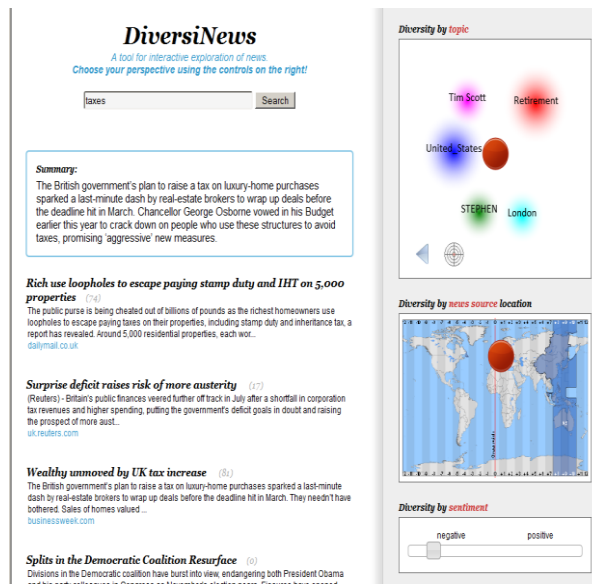


Figure 3 The summarization interface, displaying the same set of articles as in Figure 2, this time with a focus on the European viewpoint with a negative connotation. The top articles are different due to automatic reordering.

3.2 System Architecture

The demo uses the JSI Newsfeed, described in deliverable D1.1.1 Initial Data Collection, as the data source. The newsfeed includes articles from across the world, including a subset of articles collected by the Google News service.

A crawler collects web articles and associated metadata, e.g. the publisher, and cleans them of redundant HTML elements, e.g. menus and navigation. The articles are then passed through Enrycher (Štajner et al., 2010), a service which augments them with keywords, identifies the named entities appearing in the text, and performs a sentiment analysis. Additionally, we have extended Enrycher to include full constituency parses of input sentences. We use a wrapped version of the Stanford parser (Klein and Manning, 2003) for the task.

The processed articles are stored in a database with inverted indices for efficient retrieval. When the user submits a query, the articles are first filtered by the query terms. Each article is then assigned a relevancy score for each of the viewpoint controls from the user interface in an intuitive fashion:

- For topicality, the cosine similarity between the article and the keywords of the »semantic region« around the red UI dot. The keyword data is not as sparse as it might appear from the figures; only the most prominent keywords are displayed to the user, but the underlying model is denser.
- For source geography, we simply take the euclidean distance between the red UI dot and the publisher location.
- For the sentiment, we map both the sentiment score given by Enrycher and the position of the UI slider to the [-1,1] interval; the product of the two numbers is taken as the relatedness measure.

The three numbers are combined in a linear fashion using experimentally determined weights to form a final score. This score is used

- to reorder the articles on the web page, showing the most relevant first, and
- as (partial) input to the summarizer for the guided summarization task described in the next section

All the computation is done server-side; whenever one of the UI controls change, a new request is issued in the background. The server returns the summary and a short list of the most relevant articles, both to be displayed to the user.

4 Summarization technology

The core processing component in this use case is the summarizer with the ability to generate news summaries describing the same event according to different viewpoints.

According to the use case specification, we needed a summarizer able to adapt to different situations:

- The generation of generic summaries of the news, before the user has specified explicitly any particular viewpoint or bias with which to guide the summary.
- The generation of query-focused summaries where some named entities are given more important, if the user wants to have more information about how these are involved in the event or how journalists perceive their involvement.
- The generation of polarity-biased summaries where either more positive or more negative attitudes are primed in the generation of the summary.

A summarizer has been built complying with these requirements. Furthermore, in order to facilitate its integration with the other modules of the RENDER consortium and to maximise external impact, the summarizer has been fully open-sourced and made available at <http://code.google.com/p/summarizer>.

Figure 4 shows the architecture of the summarizer:

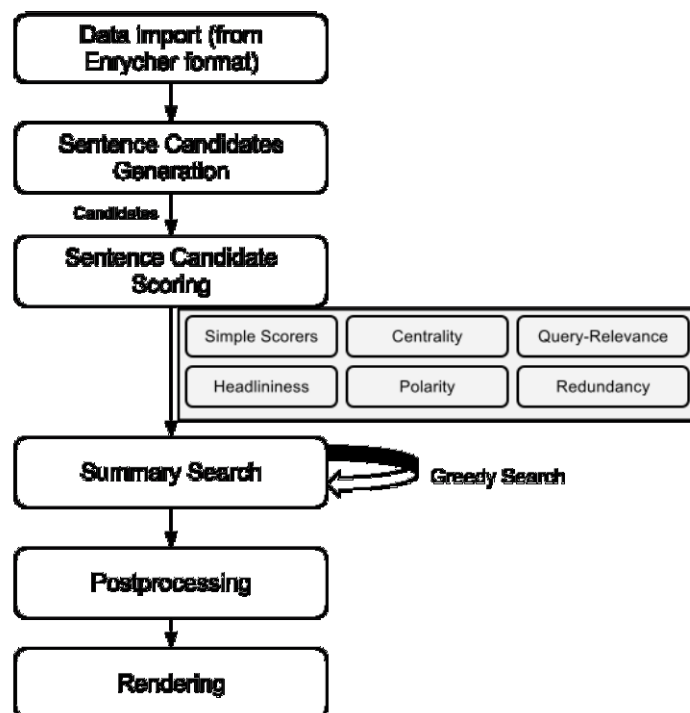


Figure 4 Summarization Architecture

The generation of a summary works in the following way:

- Initially, the output from Enrycher (provided in XML) needs to be parsed and transformed into the summarizer's input format. The data structured used for input and output in the summarizer are

all based on protocol buffers, which allow for very efficient serialization, extensibility, and are language and platform neutral, facilitating portability across applications.

Document collections, documents and sentences are represented using the protocol buffers defined in:

<https://code.google.com/p/summarizer/source/browse/src/summarizer/document.proto>

Term distributions to identify the most central topics of news collections are represented using the format in

<https://code.google.com/p/summarizer/source/browse/src/summarizer/distribution.proto>

Finally, the summarization options are represented using the protocol buffer defined in

<https://code.google.com/p/summarizer/source/browse/src/summarizer/summarizer.proto>

- Once the input has been translated into the summarizer's input format, it starts by storing the set of sentences from the original documents. Using these, an extractive summary will be constructed.
- The scorers are initialized by analyzing the document collection, from which statistics are extracted, e.g. identifying terms that represent the central topic which the news are reporting. This step needs to be run just once, when the user selects a news collection of interest.
- For every summarization request, the collection sentences are scored and a non-redundant summary is generated. Particular summarization options, e.g. whether the summary needs to reflect more positive or more negative opinions, are taken into account during sentence selection. A post-processing step removes unnecessary fragments of the selected sentences, may attempt to reorder them, and returns them to the client. These steps are performed whenever the user requests a new summary, e.g. by changing the options about polarity or the entities of interest.

4.1 Centrality models

This section describes the main scoring mechanism used in order to identify which sentences contain the most central information in the collection of news.

If we assume that we can obtain a language representing the most central event in a news collection, $P_D(\cdot)$, then sentences can be scored using Kullback-Leibler divergence, a widely used metric to measure the similarity between two probability distributions. A greedy, incremental algorithm to search for the summary that overall minimizes the KL divergence with respect to P_D is described in (Haghighi and Vanderwende, 2009). Ideally, every possible extracted summary S would be evaluated as

$$Score(Sum) = KL(P_D || P_S)$$

where KL is the Kullback-Leibler divergence, and P_S is the empirical unigram distribution from the summary. Because the space of possible summaries is combinatorial in the total number of sentences in the collection, evaluating all of them is intractable in practice and a greedy approach is commonly followed by adding to the summary, at each iteration, the sentence that minimizes the divergence. This algorithm is called KLSum.

An open question in KLSum is how to estimate the collection distribution $P_D(\cdot)$. A well-known system for this is TopicSum (Haghighi and Vanderwende, 2009), that uses a simple LDA-like topic model (Blei et al., 2003) to better estimate the collection topic model used in KLSum.

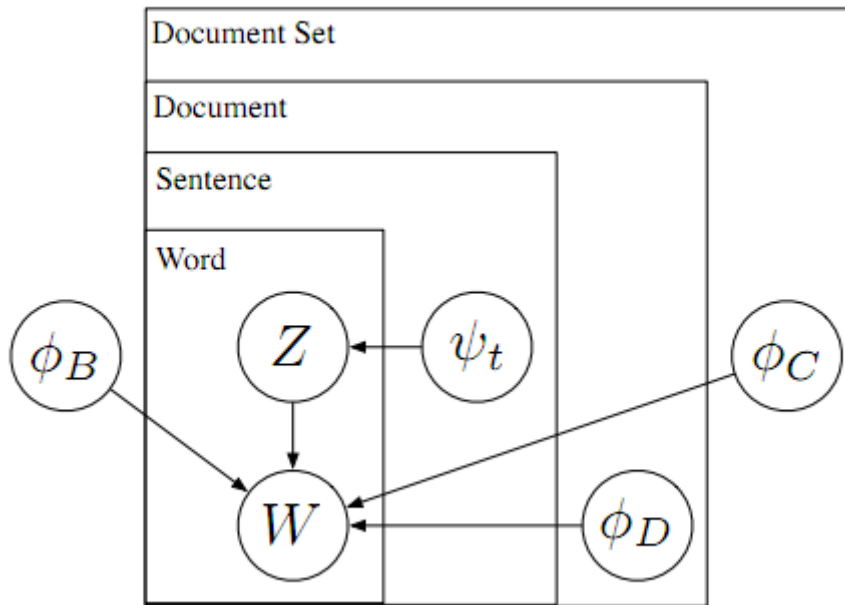


Figure 5 TopicSum plate diagram

Figure 5 shows the plate diagram. The objective is to be able to identify which words belong to (a) a background distribution ϕ_B containing general words that can appear in any document about any topic, such as stopwords and auxiliary verbs; (b) per-document distributions ϕ_D containing the most characteristics words in that document; and (c) per-collection distributions ϕ_C containing the most characteristics words in each collection. Given a particular document collection \mathcal{C} to be summarized, its distribution ϕ_C will be used as the collection distribution in KLSum. In the diagram, W refers to a word, and Z is the selector variable that decides whether the word is chosen from the document distribution, the collection distribution or the background distribution. The mixture of topics inside a document that determines the topic probabilities that govern Z is ψ_t .

4.2 Library usage

Figure 6 has an example code of how the summarizer can be used. Basically, it is necessary to invoke two functions, `Init()`, which initializes the sentence scorers and infers the collection distribution, and `Summarize()`, which produces a summary according to the summary options provided.


```
#include "summarizer/summarizer.h"

#include <fstream>
#include <iostream>
#include <streambuf>
#include <string>
#include <vector>

using namespace std;
using namespace topicsum;

int main(int argc, char** argv) {
    // Read article.
    ifstream article_file("testdata/article_14319162.xml");
    if (!article_file) {
        cerr << "File not found\n";
        return 1;
    }

    // Read file content (STL magic).
    string article((istreambuf_iterator<char>(article_file),
        istreambuf_iterator<char>()));

    // Create options for summarizer.
    SummarizerOptions summarizer_options;

    // Create KLSum options.
    KLSumOptions* klsумoptions = summarizer_options.mutable_klsум_options();
    // (Optional) Weight of a sentence score from a summarizer.
    // Default: 1.0
    klsумoptions->set_summary_weight(0.5);
    // (Optional) Weight of a prior score of a sentence.
    // Default: 0.0
    klsумoptions->set_prior_weight(0.5);

    // Add one document to collection.
    Article* document = summarizer_options.add_article();
    // Set the content of the article.
    document->set_content(article);
    // Add scores for sentences. Scores must be in same order as sentences in
    // the document.
    const int kNumberOfSentences = 17;
    for (int i = 0; i < kNumberOfSentences; i++) {
        // Prior score of a sentence.
        document->add_score(0.5);
    }

    // Create summary options.
    SummaryOptions summary_options;

    // Set the expected length of the summary.
    SummaryLength* sum_length = summary_options.mutable_length();
    // Set desired unit (SENTENCE or TOKEN).
```

```
sum_length->set_unit(SummaryLength::SENTENCE);
// Set the expected length of the summary.
sum_length->set_length(2);

// Initialize summarizer with collection and options. Return true on success.
TopicSummarizer summarizer;
if(!summarizer.Init(summarizer_options)) {
    // Print last error.
    cerr << summarizer.last_error_message();
    return 1;
}

// Sumarize the collection with desired options. Return true on success.
string summary;
if (!summarizer.Summarize(summary_options, &summary)) {
    // Print last error.
    cerr << summarizer.last_error_message();
    return 1;
}

// Print the summary.
cout << summary << endl;

return 0;
}
```

Figure 6 Example code of how the summarizer can be used

4.3 Examples

Figure 7, Figure 8 and Figure 9 show example summaries generated with these systems for real clusters of news articles, together with a title and snippet for a selection of the articles. These have been generated using TopicSum, by extracting the sentences that are most similar to the collection language models.

[Dodgers wrap up set against Cards without Kershaw](#)

Consecutive victories have the Los Angeles Dodgers in position to take over sole possession of the second wild card position in the National League from the St. Louis Cardinals. However, they will be without their ace Clayton Kershaw on Sunday ...

[Dodgers lose 5-2 to Cardinals in 12 innings](#)

AP The Dodgers had the Cardinals in their ballpark for four days. All they have to show for it is a split, which leaves them in the same position they were in when the defending World Series champions came to Los Angeles - a game behind St. Louis for ...

[Ethier reaches Waino for 2-run homer](#)

LOS ANGELES -- Little mistakes continue to grow into large ones for the Cardinals. Seemingly in control of the Dodger lineup, Adam Wainwright committed his first gaffe by issuing a leadoff walk of catcher Matt Treanor to begin the third inning.

[Dodgers Lose to Cardinals in 12, Fall One Game Back in Wild Card Hunt](#)

The sheen from ElyMania 2010 has long lost its luster. Locked in a 2-2 stalemate for much of the game, Ely once again melted down in extra innings for the Dodgers allowing the St. Louis Cardinals to escape town with a 5-2 victory.

[NL roundup: Astros score 4 in 7th to rally past Phillies](#)

Los Angeles Dodgers' Dee Gordon, bottom, looks back after being tagged out by St. Louis Cardinals second baseman Daniel Descalso while trying to steal second during the seventh inning of their baseball game on Sunday, Sept. 16, 2012, in Los Angeles.

[Clutch Dodgers beat Cards again](#)

LOS ANGELES – Luis Cruz came up with his second clutch hit in two nights, a tying two-out RBI double in the ninth inning, and pinch-hitter Juan Rivera followed with a run-scoring single to lift the Los Angeles Dodgers to a 4-3 victory over the St.

[St. Louis Cardinals Top Los Angeles Dodgers 5-2, Regain Wild Card Lead](#)

The St. Louis Cardinals' Wild Card lead had vanished Sunday morning, after Jason Motte's blown save on Saturday night led to an in-series tie with the Los Angeles Dodgers. Sunday, in the 12th inning, it came back decisively on a Jon Jay double into the ...

Summary[Dodgers left game back of Cardinals for wild card after 5-2 loss](#)

[ESPNLosAngeles.com The Los Angeles Dodgers' playoff hopes took a big, slow-motion punch with Sunday's 5-2 loss to the St. Louis Cardinals in 12 innings.](#)

[The Cardinals jumped out to a 2-0 lead in the first when Matt Holliday singled with two outs and Craig followed with his 21st home run of the season.](#)

Figure 7 Cluster of news with snippets, and summary generated from it.

[Kim Kardashian Takes Home a Kitten and Names Her 'Mercy'](#)

Kardashian, 31, Tweeted Sunday alongside a series of snapshots of her new addition (who may make the perfect playmate for sister Kourtney's 2-month-old Penelope). So what's in a name? Good news for the rapper's ego: The kitty's moniker seems to be a ...

[Kim Kardashian, Kanye West dating, the move for her.](#)

(ThyBlackMan.com) Yeah Kim Kardashian is still dealing with the rather negative after-glow of her highly publicized relationship with Kris Humphries, which included an over-the-top, fairy-tale, made-for-TV wedding that sadly ended in divorce and not so ...

[News: Kim Kardashian Says Oh Mercy, Names Her Kitty After Kanye West](#)

Recently, Yeezy raided Kardashian's wardrobe with his personal style assistant to help give Kim a fashion makeover. In a new teaser clip for "Keeping Up With The Kardashians," Kanye and his stylist take over Kim's massive closet, as he outfits her with ...

[Kim Kardashian Tweets Photos of New Cat](#)

... with a collage of photos. Kim named her new Persian kitten "Mercy" after the title of her boyfriend Kanye West's latest hit song from his album "Cruel Summer". Kardashian recently appeared on "The View" and opened up about her relationship with West, 35.

[Kim Kardashian Has Mercy: Meet Her Adorable New Kitten](#)

Kim Kardashian got herself a new and absolutely adorable little sidekick, who was spotted with the star while in Miami, but Kimmy made sure to formally introduce her new friend to all of the Twitterverse. "Meet Mercy, " she wrote. Meet Kendall and Kylie's new ...

[Kim Kardashian names new kitten after Kanye West song](#)

Leave it to Kim Kardashian to use a new kitten to help promote her boyfriend's music album "Cruel Summer." What do you think of her name choice? Kim Kardashian. Photo credit: Getty Images. Here family can't stay out of the headlines this week at all.

Summary

[Kim Kardashian Wows in Black as She Names New Kitten After Kanye West's Song 'Mercy'](#)

[Kim Kardashian has named her new kitten after one of her boyfriend Kanye West's songs. During the interview, Kimmie was asked whether she felt pressure to televise her wedding for her reality show ``Keeping Up With the Kardashians."](#)

[Arriving at the sister-operated Dash boutique in Miami on Sunday, the reality star stepped out with a little white kitten named Mercy.](#)

Figure 8 Cluster of related news, and summary generated from it.

[Herald News: Calculating calories](#)

WE ARE slowly realizing, as a nation, what an obesity epidemic means. It's taken too long, and cost too many lives, but at least now we are starting to make important changes. One of them is paying more attention to what we eat, and how much.

[McDonald's adds calorie counts to menu](#)

McDonald's announced Wednesday it will start to include calorie information on its menus. A Big Mac is 550 calories, and a 10-piece McNuggets is 470 calories, according to the McDonald's website. Jessica Nascimento went to McDonald's for her favorite ...

[McDonalds will post calorie info – thanks to the ACA](#)

Today, the McDonalds restaurant chain announced that beginning on Monday it will post the calories in every item on their menus, including menu boards and drive-thru boards. And herein lies an interesting story. A few years back when New York City and ...

[McDonald's to post calories for its menu](#)

First they began offering salad, then they added fruit. On Wednesday, McDonald's announced it would put something else on its menu to help customers watching their waistlines: calorie counts. The fast food giant said that starting next week, ...

[McDonald's new menus likely won't make us thinner](#)

By Tim Carman Mickey D's likely has nothing to worry about by listing the calorie counts on its menu. (Tim Carman/The Washington Post) Earlier this week, McDonald's decided it would be among the first to take a bullet and announce the roll out of new ...

[McDonald's unveiling new calorie count menu board](#)

McDonald's is unveiling a new menu board that will show the calorie count of each item. Nutritionist and registered dietitian Cynthia Sass, author of "Sass Yourself Slim," speaks to the "CBS This Morning: Saturday" co-hosts about the change.

[Calorie Counts Coming to McDonald's Menus](#)

McDonald's announced last week that starting this month, the company will list calorie information on restaurant and drive-thru menus nationwide, including here in Paulding County, to further inform and help customers and employees make ...

Summary[McDonald's adds calorie counts to menu](#)

[Starting Monday, McDonald's will display calorie information on restaurant and drive-thru menus nationwide, such as a New York McDonald's has for its breakfast menu.](#)

[But that same study found that fast food calorie counts were more accurate than the counts posted at sit-down restaurants.](#)

Figure 9 Cluster of related news articles and summary generated from it.

5 Evaluation

During development, to allow for a fast turnaround and a short development cycle, we have been using an automatic evaluation framework, consisting of previous test sets from the competitive summarization evaluations in the Text Analysis Conference (TAC). Although we were able to obtain state-of-the-art results on datasets from previous years, some of which were published in (Delort and Alfonseca, 2012), we came to the conclusion that these evaluations are run in very controlled settings which do not always reflect real usage scenarios. For example, in these competitions the news collections are always of the same size, and the documents are carefully selected to make sure that they are all about the same topic. These are two assumptions that hardly ever happen in real life, where automatic news clustering techniques are noisy and sometimes include off-topic news, and for some events there may be only a few articles available.

In order to have a more realistic evaluation for our task, we have implemented some task-specific evaluation templates. Figure 10 shows a template example for evaluating, at the same time, the quality of the summary (informativeness), whether it reflects the expected polarity (positive or negative) and whether it shows information focused on the relevant entities selected by the user. This, applied to a random sample of news and possible user selections, allows us to track progress over time and ensure that the quality of the system is improving.

The final evaluation results will be described in D5.2.4.

Please make yourself familiar with the news below:

[Oil prices rise ahead of key European bank meeting](#)

AP The price of oil rose above \$96 a barrel on Thursday ahead of a meeting of the European Central Bank which is expected to announce a plan to help ease the eurozone's debt crisis. The ECB is expected to announce a bond-buying program to reduce high ...

[Oil price rises ahead of key European bank meeting](#)

2 hours ago • AP BANGKOK (AP) — Oil prices rose Thursday ahead of a meeting of European central bankers who are expected to announce a plan to help financially strapped countries. The European Central Bank is expected to announce a bond-buying program ...

[Oil Rises a Second Day on U.S. Supply Drop, ECB Plan Optimism](#)

Oil rose for a second day in New York amid signs of a reduction in US crude supplies and as European Central Bank President Mario Draghi prepared to outline his plan to stem the region's debt crisis. Futures gained as much as 1.3 percent after the ...

[Oil Gains a Second Day as US Stockpiles Drop to Five-Month Low](#)

By Ben Sharples on September 06, 2012 Oil rose for a second day in New York after an industry report showed stockpiles shrank to the lowest in more than five months in the US, the world's biggest crude consumer. Futures gained as much as 0.9 percent ...

Imagine that a user has specified that he is interested to know about the European Central Bank, and wants to read a summary of these news articles that highlights the most positive relevant information included in these news.

The user is presented with this summary:

ECB expectations boost crude oil prices

Crude oil futures settled slightly higher Wednesday, awaiting guidance from Thursday's US oil inventory data and signals from the European Central Bank's policy meeting. US gasoline stockpiles probably fell 3 million barrels last week, according to the median estimate of nine analysts in the Bloomberg survey before the Energy Department report.

Does this summary accurately reflect the main news event reported in the original articles?

Strongly disagree Disagree No opinion Agree Strongly agree

Does this summary include a positive view on the event?

Strongly disagree Disagree No opinion Agree Strongly agree

Is this summary relevant about the European Central Bank?

Strongly disagree Disagree No opinion Agree Strongly agree

Figure 10 Evaluation template for summaries.

6 Conclusions and Future Work

A system for automatically generating summaries of news stories according to different viewpoints has been built. It has been a close collaboration between Google, responsible of the summarization technology, and JSI, which has contributed the language technologies and the frontend. The system is currently being iterated upon and being evaluated.

References

- [1] [Blei et al., 2003]
Blei, D. M., Ng, A. Y. and Jordan, M. I., *Latent dirichlet allocation*. JMLR (2003)
- [2] [Dang, 2006]
Trang Dang, H., *DUC 2005: Evaluation of question-focused summarization systems*. In Proceedings of the COLING-ACL'06 Workshop on Task-Focused Summarization and Question Answering, pages 48–55, Prague, 2006.
- [3] [Delort and Alfonseca, 2012]
Delort and Alfonseca, E., J.-Y., *A Topic-Model for Update Summarization*, in proceedings of EACL-2012.
- [4] [Haghighi and Vanderwende, 2009]
Haghighi, A. and Vanderwende, L., *Exploring content models for multi-document summarization*. In Proceedings of Human Language Technologies: The 2009 Annual Conference of the North American Chapter of the Association for Computational Linguistics, 2009.
- [5] [Klein and Manning, 2003]
Klein, D., and Manning, C. D. (2003). Accurate unlexicalized parsing. Proceedings of the 41st Annual Meeting on Association for Computational Linguistics (pp. 423–430).
- [6] [Mani and Maybury, 1999]
Mani, I. and Maybury, M. T., *Advances in automatic text summarization*. The MIT Press, 1999.
- [7] [Nenkova and Vanderwende, 2005]
Nenkova, A. and Vanderwende, L., *The impact of frequency on summarization*. Technical report, Microsoft Research (2005).
- [8] [Pang and Lee, 2008]
Pang, B. and Lee, L., *Opinion Mining and Sentiment Analysis*. In Foundations and Trends in Information Retrieval, volume 2, issue 1-2., January 2008.
- [9] [Stoyanov et al., 2004]
Stoyanov, V., Cardie, C., Litman, D. and Wiebe, J., *Evaluating an opinion annotation scheme using a new multi-perspective question and answer corpus*. In Proceedings of the AAAI Spring Symposium on Exploring Attitude and Affect in Text: Theories and Applications (2004).
- [10] [Štajner et al., 2010]
Štajner, T., Rusu, D., Dali, L., Fortuna, B., Mladenčić D., Grobelnik, M. (2010). A service oriented framework for natural language text enrichment. *Informatica (Ljubljana)*, 2010, 34:3, pp. 307-313.
- [11] [TAC, 2011]
TAC 2011 Summarization Track. <http://www.nist.gov/tac/2011/Summarization/> Proceedings to appear at the Text Analysis Conference (TAC) 2011 Workshop, November 14-15, 2011, National Institute of Standards and Technology, Gaithersburg, Maryland USA
- [12] [Titov and McDonald, 2008]
Titov, I. and McDonald, R., *A joint model of text and aspect ratings for sentiment summarization*. In Proceedings of ACL-2008.
- [13] [Yu and Hatzivassiloglou, 2003]

Yu, D. and Hatzivassiloglou, V., *Towards answering opinion questions: Separating facts from opinions and identifying the polarity of opinion sentences*. In Proceedings of EMNLP-2003.

[14] [Yue et al., 2009]

Yue, L., Zhai, C. and Sundaresane, N., *Rated aspect summarization of short comments*. In Proceedings of WWW-2009.

[15] [Zhu et al., 2009]

Zhu, J., Zhu, M., Wang, H. and Tsou, B. K., *Aspect-based sentence segmentation for sentiment summarization*. In Proceeding of the 1st international CIKM workshop on Topic-sentiment analysis for mass opinion, 2009.