



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

RENDER

Deliverable D5.2.3

Refined Diversified News Service

Editor:	Daniele Pighin, Google Inc.
Author(s):	Enrique Alfonseca and Daniele Pighin, Google Inc. Mariana Damova, Ontotext.
Deliverable Nature:	Prototype (P) (due to DoW)
Dissemination Level: (Confidentiality) ¹	Public (PU)
Contractual Delivery Date:	March 2013
Actual Delivery Date:	March 2013
Suggested Readers:	Developers involved in the use cases.
Version:	1.1
Keywords:	Summarization, news, diversity

¹Please indicate the dissemination level using one of the following codes: • **PU** = Public • **PP** = Restricted to other programme participants (including the Commission Services) • **RE** = Restricted to a group specified by the consortium (including the Commission Services) • **CO** = Confidential, only for members of the consortium (including the Commission Services) • **Restreint UE** = Classified with the classification level "Restreint UE" according to Commission Decision 2001/844 and amendments • **Confidentiel UE** = Classified with the mention of the classification level "Confidentiel UE" according to Commission Decision 2001/844 and amendments • **Secret UE** = Classified with the mention of the classification level "Secret UE" according to Commission Decision 2001/844 and amendments

Disclaimer

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

In case of Public (PU):

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

In case of Restricted to Programme (PP):

All RENDER consortium parties have agreed to make this document available on request to other framework programme participants.

In case of Restricted to Group (RE):

The information contained in this document is the proprietary confidential information of the RENDER consortium and may not be disclosed except in accordance with the consortium agreement. However, all RENDER consortium parties have agreed to make this document available to <group> / <purpose>.

In case of Consortium confidential (CO):

The information contained in this document is the proprietary confidential information of the RENDER consortium and may not be disclosed except in accordance with the consortium agreement.

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 – Diversity case studies
Document Title:	D5.2.3 – Refined Diversified News Service
Editor (Name, Affiliation)	Daniele Pighin, Google
Work package Leader (Name, affiliation)	Javier Caminero, Telefonica

Copyright notice

© 2010-2013 Participants in project RENDER

Executive Summary

This document describes the latest technology incorporated in the Google use case scenario, which presents collection of related news to users highlighting:

- different points of view, attitudes towards the event;
- different entities (people, organizations, etc.) involved in the event reported in the news.

The use case integrates different components from the RENDER project, namely:

- JSI's news crawler NewsFeed (D1.1.1, D1.1.3), 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 provided by Enrycher (JSI) which includes the standard analysis annotations: token boundaries, sentence boundaries, part-of-speech tags, syntactic dependencies and named entity annotations.
- Summarization technologies, (Google and JSI) which allow the system to generate event summaries taking into account different scorers depending on the user's request.
- The RENDER Data Layer (described in D.1.1.3) and in particular Ontotext Web API and FactForge (D.1.3.1) which we query to retrieve sets of entities which are related to the news currently displayed.

With respect to D5.2.2 - "Diversified News Service", this deliverable introduces two main novelties that enhance the integration of different RENDER components into the use case: relevant entities retrieved through FactForge and a new summarizer developed at JSI. The deliverable also includes an evaluation of the related entities retrieved through Ontotext Web API.

The *DiversiNews* platform is now considered stable, and we are not planning to add any new features. During the next months, most efforts will be devoted to polishing the integration of the different components and to the evaluation of the DiversiNews service. The complete evaluation of the service, in which both its quality and usability will be assessed, will be reported at the end of the project in deliverable D5.2.4.

Table of Contents

1	Introduction.....	6
2	Use Case Description.....	7
3	The DiversiNews Platform.....	8
3.1	User Interface.....	8
3.2	System Architecture.....	9
3.3	Summarization technology.....	11
3.4	Related Entities Panel.....	12
4	Evaluation.....	14
4.1	Evaluation of the Related Entities Feature.....	14
4.2	Final Evaluation Plan.....	16
5	Conclusions and Future Work.....	18
	References.....	19

List of Figures

The search view of the DiversiNews service.....	8
The main view of the DiversiNews platform.....	9
Integration of the DiversiNews platform.....	10
Component diagram of the DiversiNews platform.....	11
Component diagram of Google summarization technology.....	12
Results of the first round of related entities annotation.....	14
Results of the second round of related entities annotation.....	15
A template for the evaluation of summarization technology.....	17

1 Introduction

This document describes the final version of the Google use case application, developed in collaboration with JSI and Ontotext, for presenting recent, relevant news to users and allowing them to explore the content of those news from different points of view, highlighting different attitudes towards the event and showing related entities, not necessarily present in the news currently displayed, that might be of interest for the reader.

In Section 2 we briefly review the main aspects of the Google use case, which has been discussed in more detail in D5.2.2. In Section 3 we present the latest version of the DiversiNews platform, which implements the use case. In particular, we focus on the differences with respect to the version of the platform that we presented in D5.2.2. In Section 4 we report on the results of the evaluation of the latest addition to the DiversiNews platform, i.e., the FactForge powered related entities panel, and detail our plans for a vertical evaluation of the platform that will be at the core of the coming deliverable D5.2.4. Finally, in Section 5 we draw our conclusions.

2 Use Case Description

The use case is based on the idea of providing diversity-aware visualization tool for users browsing recent news. The overwhelming amount of news that users are exposed to whenever they search for news on some specific topic makes it very difficult for readers to understand what are the differences between news coming from different sources, originating from different locations, expressing different points of view or representing different interests. The use case tool has been designed to allow users to focus on a particular news story (reported by different sources, describing something that has happened recently) and to dive, by means of automatic text summarization techniques, into different viewpoints and attitudes.

The visualization tool is able to process the news that refer to the same events, and to show to the users additional ways for analysing the information, including:

- **What are the most relevant topics covered in the news from the user's point of view?** We make the assumption that the user interests can be expressed in terms of the most relevant topics mentioned in the news, and allow users to selectively select the news within a collection that are mostly relevant with respect to the selected topics.
- **What are the different points of view in the different news that refer to the same event?** Different sources may report the same data with differences in bias or sentiment, or may simply highlight and omit different characteristics of the news. We isolate two main components to characterize news polarization: the geographical source of the news and the emotional polarity of the words used to describe an event.

These two dimensions of analysis were already available in the version of the DiversiNews platform presented in D5.2.2. In the current version of DiversiNews, we have added a further dimension along which users can explore the content of news:

- **What other topics should the user explore to form a more informed opinion on a set of news?** To put some news in perspective, it might be necessary to acquire more information about related topics which are not directly mentioned in the news. For example, a user reading about the bleeding-edge features of a new smart-phone may benefit for reading the reaction of the press to another model released some weeks before by another company. Providing links to news clusters discussing *smart-phones* or other companies' products would help the user getting a clearer picture.

To make the large amount of information easily accessible, we employ automatic summarization technology to present to the users a compact representation of the most relevant information that can be retrieved according to the specified preferences. Whenever the user decides to look at the news from a different angle, the most relevant news for that perspective are selected, and a summary containing the most salient information is generated. These aspects have already been discussed in D5.2.2. As the summarizer may itself introduce a bias in the information received by users, in the latest version of the platform we have also included the possibility to make use of different summarization technologies, which users may select based on their preferences.

The next section describes how all these aspects are integrated in a unified interface for diversity-aware access to news. Unlike with D5.2.2, we especially focus on how the DiversiNews service integrates technologies provided by different partners of the RENDER consortium.

3 The DiversiNews Platform

The DiversiNews platform is accessible as a web application at <http://aidemo.ijs.si/diversinews/>. The service offers users a fast overview of a set of news articles and allows them to drill down into the diverse viewpoints represented by individual articles. In the remainder of this section, in 3.1 we will describe the latest version of the interface that is presented to the users; in Section 3.2 we will discuss the integration of all the components with respect to the RENDER technical architecture; in Section 3.3 we will review the summarization technologies currently employed by DiversiNews. Finally, in Section 3.4 we will discuss the related entities panel that we recently added to the service. The quality of the entities shown in the panel will be the subject of the evaluation that we will carry out in Section 4.

3.1 User Interface

The interface of DiversiNews consists of two main views. In the first view, shown in Figure 1, users can either insert a query in a search box to automatically retrieve a collection of related news. As an alternative, they can click on the title of some recent news, listed at the bottom of the panel, to summon the related collection of news.

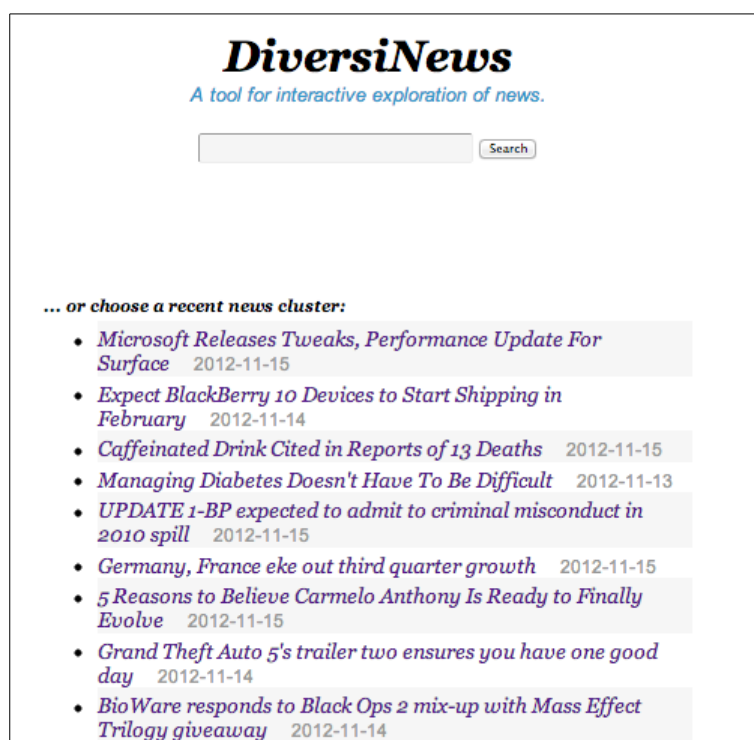


Figure 1: The search view of the DiversiNews service.

After selecting the news cluster of interest, users are presented with the interface depicted in Figure 2. In this case, the user selected a set of news related to the recasting of Daniel Craig in the role of James Bond. A summary of results from recent news appears on top; just below it, the user is presented with the individual articles that have been summarized, sorted by relevance. At the top of the summary box, the user can select what summarizer should be used. On the right, interactive widgets allow the user to further specify which perspective on the news should be emphasized.

The first panel displays keywords (topics) automatically extracted from the analysed 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. The slider below allows placing focus on news with either a positive or a negative outlook. Finally, the box at the

bottom shows other entities which are related to, yet not necessarily mentioned in, the set of news. Whenever one of the red dots or the slider gets moved, the summary is updated accordingly and the articles reordered. When the user clicks on one of the related entities, a new collection of news is displayed using the selected entity as the search term.

DiversiNews
A tool for interactive exploration of news.

Search: Search

««« ... or return to the cluster listing

Summary:
Choose summarization algorithm: *Type1* (current) *Type2*

The movie site's Pete Hammond has revealed that there are murmurs about 'Skyfall' being so good, as in the "best ever", and so could really be in with a chance of security a 'Best Picture' nod at the Academy Awards, as well as possibly for its supporting cast.

I didn't anticipate needing to arrive at work at 8 a.m. to edit one of our monthly publications, but I never said that going through with this plan was a particularly smart decision.

"It is the first I can recall getting truly serious Oscar talk, not only for Best Picture but even in acting categories for supporting players Javier Bardem as the villainous Silva and Judi Dench with a particularly meaty turn as M this time around," he wrote.

Bond is back, and better than ever
If this year is any indication, I often make poor choices when it comes to seeing movies. As a superfan of properties like Batman, I find midnight showings ridiculously enticing. There's something about the prospect of seeing a movie before everyone...
[lyndentribune.com](#) (4874 69952709 eng +0.462 [48.9,-122.5])

Bond Takes to the Stratosphere in Skyfall
This Bond is different. Anyone who has seen either of Daniel Craig's previous portrayals of the secret agent already knows this. Craig's Bond is more visceral, action oriented and serious than previous versions of the famed super spy. This is no...
[synapse.ucsf.edu](#) (9117 70113244 eng +0.448 [0.0,0.0])

Roger Moore: Daniel Craig is now my favorite James Bond
This year not only marks James Bond's 50th anniversary onscreen, but also 40 years since Sir Roger Moore accepted the role that Sean Connery defined. Moore's first outing as the famed 007 came in 1973's "Live and Let Die." He app...
[entertainmentnbcnews.com](#) (9822 70131474 eng +0.443 [38.0,-97.0])

Could Skyfall Be An Oscar Winner?
The latest James Bond movie 'Skyfall' could become the first of the 007 franchise to be in with a serious chance of taking home an Oscar for 'Best Picture', according to reports. Deadline.com has claimed that a recent screening of the...
[entertainmentwise.com](#) (4882 69952723 eng +0.406 [54.0,-2.0])

James Bond is Back and Better than Ever
Iconic, masculine and sophisticated, the film, much like James Bond, delivers a self-validating

Diversity by topic

AGENT COMMENTS SILVA
NATIONAL PETRAEUS WRIGHT
ACADEMY OSCARS PICTURE
MILLION GROSS HIGHEST

Diversity by news source location
Enable world map ☐

Google

Diversity by sentiment
negative positive

Search related

Halle Berry
James Bond
Ian Fleming
Pierce Brosnan
the United States
James Bond (ornithologist)
Daniel Craig
Javier Bardem
Sam
Royale

Figure 2: The main view of the DiversiNews platform.

3.2 System Architecture

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

Figure 3 highlights with a red box the components of the RENDER infrastructure on which DiversiNews insists. The web 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] to be annotated with keywords, named entities appearing in the text and sentiment scores.

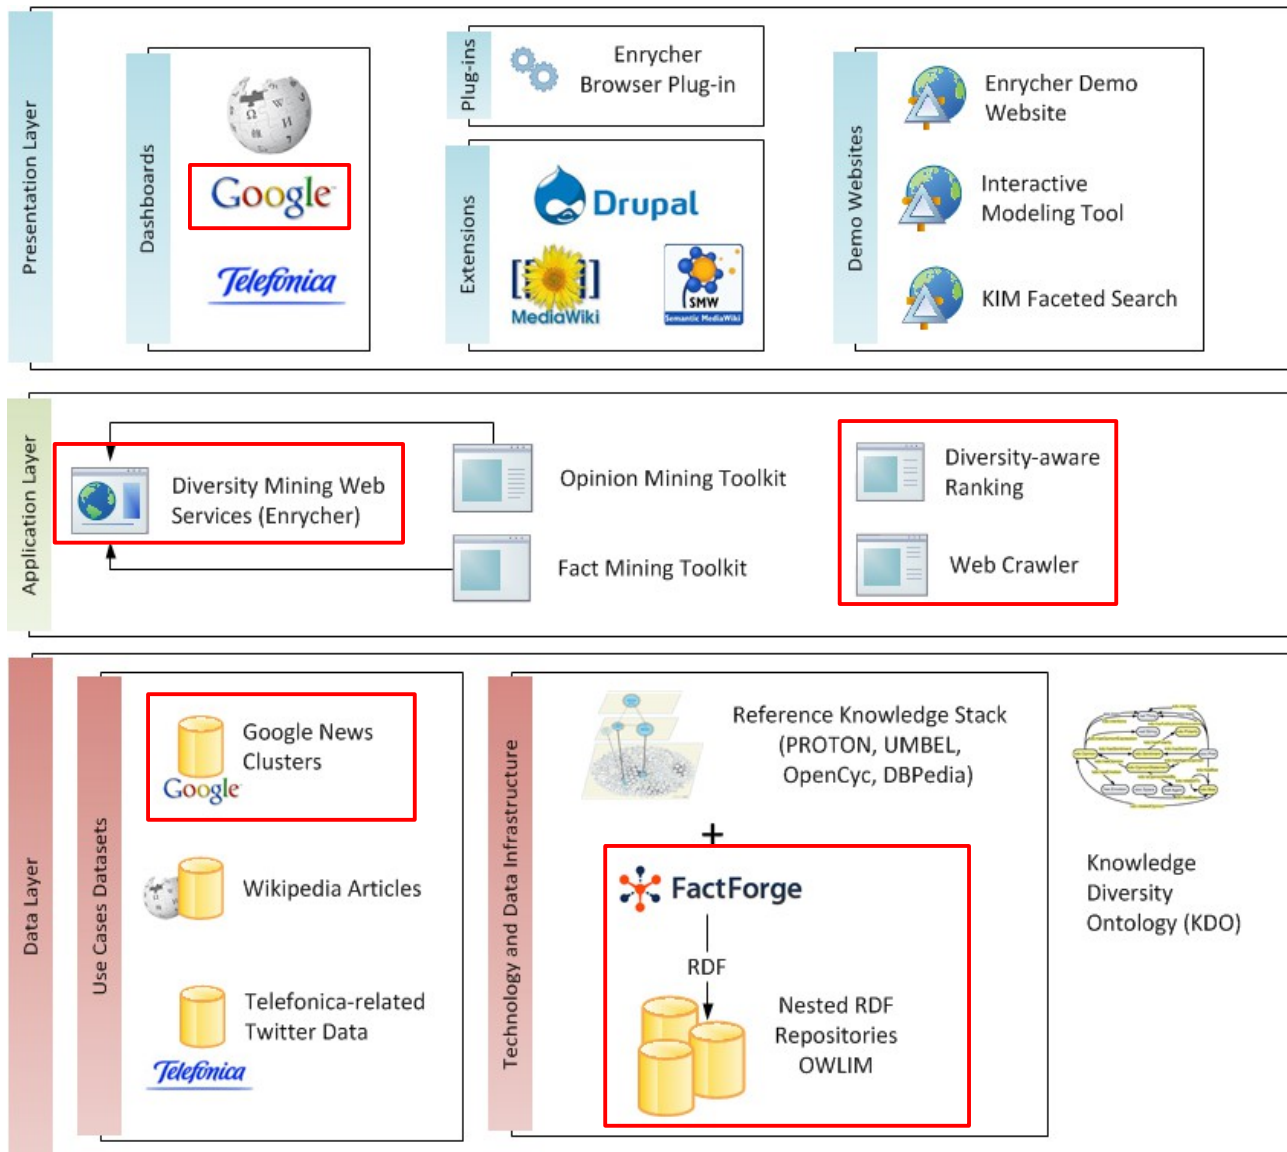


Figure 3: Integration of the DiversiNews platform.

Each retrieved 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.

These three values are combined in a linear fashion. The three partial scores are weighted equally. To make them more directly comparable, we normalize each of them prior to computing the sum: each of the three partial scores is adjusted linearly so that its mean is zero and its standard deviation one across the set of articles being sorted. The final score is used to reorder the articles on the web page, showing the most relevant first, and as input to the summarizer that will use this information to select the most relevant sentences. The related entities shown in the bottom left panel are retrieved through Ontotext Web API, as detailed in Section 3.4.

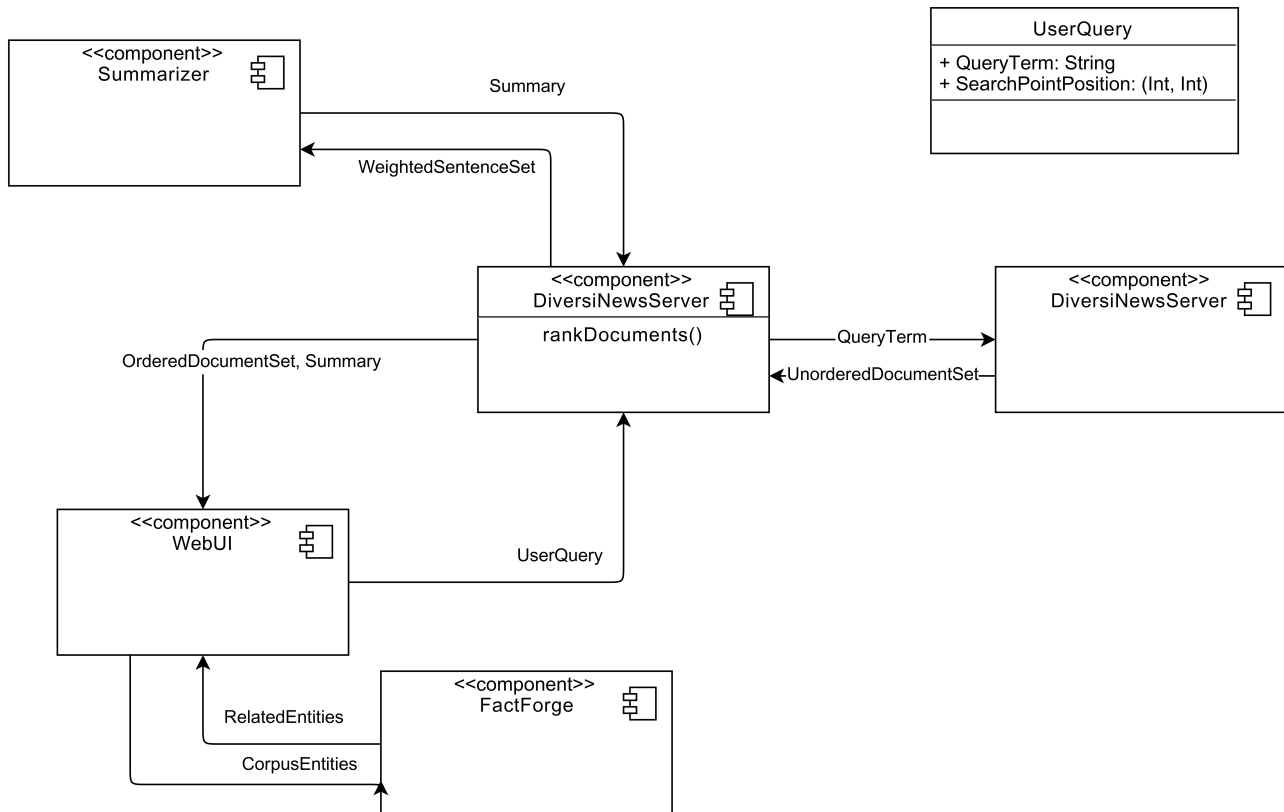


Figure 4: Component diagram of the DiversiNews platform.

Figure 4 describes the interaction among the components of the DiversiNews platform in more detail. When the interface receives a user query, it generates a request to the DiversiNewsServer, which responds with the set of ranked documents for that query and an automatically generated summary. The summary is generated by invoking the summarizer selected by the user. The most relevant entities identified in the collection of analyzed news are used to query FactForge and retrieve the set of related entities, which are displayed in the related entities panel. A lower level description of this last step will be provided in Section 3.4.

3.3 Summarization technology

One of the core design principles of DiversiNews is that information should be easily accessible. Simply filtering the news based on the user preferences might be insufficient, as salient differences in the news might be lost due to the large amount of news available. In this respect, the summary box in the top left corner of DiversiNews interface is supposed to present to the user a distilled selection of the available information that best represents the selected point of view.

Google summarizer, which has been in DiversiNews since its inception, has already been described in D5.2.2, where the component and its API are fully described and documented. Here, we would just like to provide a more detailed view of the components of Google summarization technology, to complement the information already available in D5.2.2. The interaction between the components of the summarizer is shown in Figure 5. The high level API (SummarizerInterface) invokes the TopicSum summarizer on a collection of related news with a set of user-configurable parameters. These parameters control, among other things, what filters and post-processors will

be invoked in order to generate the final summary. Within DiversiNews, these parameters are controlled by the DiversiNewsServer, which is also responsible of re-scoring the sentences in a news collection in response to changes in the diversity widgets in DiversiNews interface. By comparing the word distributions of common English texts (background distribution) against the news collection's and each document's specific distribution, TopicSum selects the sentences which are most likely to contain relevant information. Redundancy checks and post-processors decide which sentences should be included in the final summary.

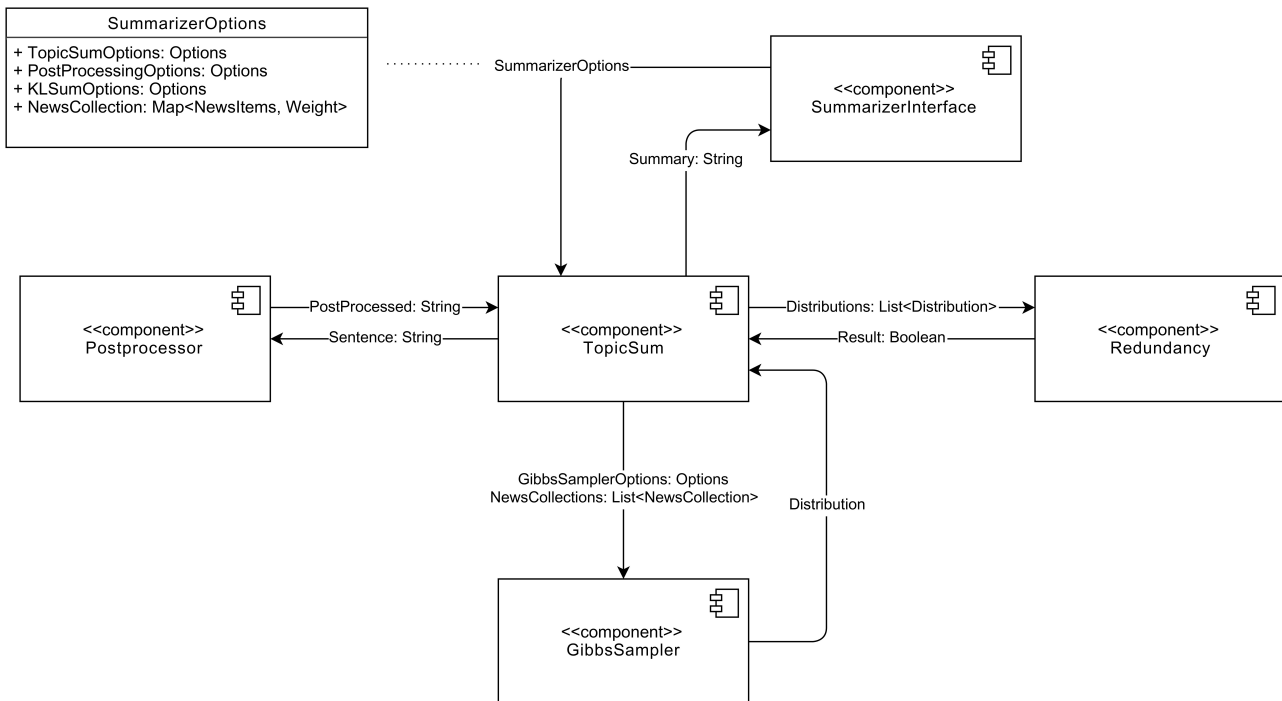


Figure 5: Component diagram of Google summarization technology.

Very recently, the JSI team has extended the DiversiNews service by adding a second summarizer that can be selected as an alternative to Google's. This alternative summarizer explores the idea of generating extractive summaries by relying on semantic representations of the input sentences which, ideally, should alleviate data-sparsity and information redundancy. More details about the new summarization technology are provided in D3.2.1. The final evaluation of the Google Use Case (to be part of D5.2.5) will also involve a side-by-side comparison of the two summarization technologies.

3.4 Related Entities Panel

The related entities panel is meant to offer users an orthogonal perspective with respect to accessing and exploiting diversity in news. The related entities panel is supposed to answer the question: *what other collections of news should I read into if I really want to understand the one that I am looking at?* To this end, the related entities panel should present the user a set of entities which are related with the news currently displayed, but not necessarily mentioned in the news. By clicking on any of the displayed entities, the user can browse on news for which those entities are central. This feature provides the connection between the news retrieved by the news crawler and the basic RENDER data layer described in D1.1.3.

Deployment-wise, the related entities panel is generated by a JavaScript function which is hosted on Ontotext servers and which is loaded by the DiversiNews interface. The DiversiNews interface invokes the function with the set of the most relevant entities that Enrycher identified in the news collection. In turn, the JavaScript function generates a SPARQL query to retrieve the related

entities throw the Web API, and populates the information in the panel by adding the entities retrieved through FactForge and those initially identified by Enrycher.

More in detail, the DiversiNews interface queries the RENDER data layer with the DBpedia URIs of the 5 most relevant topics identified in the current news collection. A JavaScript modifies a predefined SPARQL query, executes it on the basic RENDER data layer, and pushes the result to the client side, where it is displayed in the related entities panel. The SPARQL query selects all the entities which are related with the first (most relevant) of the topics which appear in the request. Here, relatedness is defined as the property of co-occurring in at least one of the RDF triples stored in the data layer and having the same type (e.g., an input entity of type “person” would only activate entities which are also labeled as “person”). Then, an algorithm inspired by Google's PageRank² is used to assign a *popularity* score to each candidate answer: the *RDF-Rank* of an entity is a score in the range [0,1] which is based on the interconnectedness of the entity in the knowledge base. The retrieved entities are sorted based on their popularity. A post processing step decides which is the “preferred label” to be assigned to an entity, and stores the labels selected for the five most popular entities that have been retrieved. The concatenation of these five entities with the five topics in the original query will be used to generate the contents of the related entities panel.

² <http://en.wikipedia.org/wiki/PageRank>

4 Evaluation

This section is divided into two parts. In Section 4.1 we report on the evaluation of the perceived quality, in term of *relatedness* with a news collection, of the entities listed in the related entities panel; In Section 4.2 we lay out our plans for the final evaluation of the DiversiNews service, which will be included in D5.2.4.

4.1 Evaluation of the Related Entities Feature

After adding the related entities panel to DiversiNews, we carried out an evaluation to understand if the if it would be possible to further integrate the RENDER data layer within the Google use case. Initial plans involved, for example, the possibility of using the related entities to slant or characterize the summaries, or to generate SPARQL queries to select relevant events to be considered for the ranking and compression of the news.

For this task, two annotators have annotated the relevance of the related entities displayed by the DiversiNews interface for 31 different news collections. The annotators were asked to judge the quality of the related entities from the point of view of a user who has an interest in the currently selected collection of news, and who would like to continue browsing related news that might present a different angle on similar matters. Each set of entities has been evaluated on a 5-points *Likert* scale, where a score of 1 means “completely irrelevant” and a score of 5 means “completely relevant”. In the cases in which no entities were displayed, the annotators were instructed to label the example as 0, so as to differentiate this case from the cases in which they were dissatisfied with the results.

The evaluation has been carried out in two rounds. During the first round evaluation, the annotators were presented with an earlier version of the system, which differs from the one described in Section 3.4 for two aspects:

- the related panel only displays entities retrieved through the RENDER data layer, i.e., the most relevant entities identified by Enrycher in the news collection are *not* shown in the related entities panel;
- the entities in the related panel are separated with a comma, and they are not listed on separate lines as shown in Figure 2.

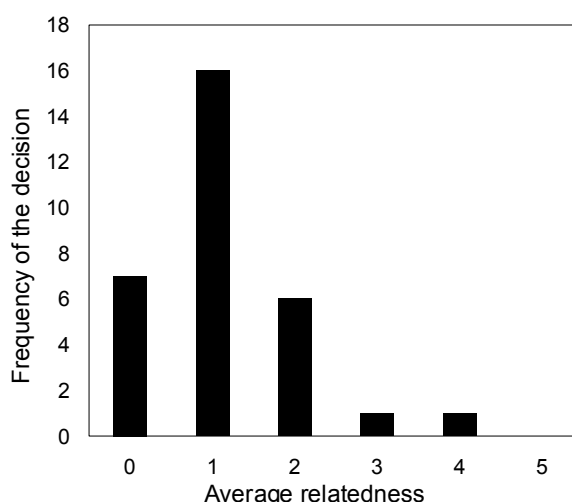


Figure 6: Results of the first round of related entities annotation.

Figure 9 Shows the results of the first round of evaluation. In seven cases out of 31, the related entities panel did not show any result. In the remaining cases, in the majority of the cases (16/24) the annotators labeled the set of entities as completely unrelated. The second most frequent choice (6/24) is mostly unrelated. Only in two cases the entities are labeled as acceptably or mostly related.

The agreement between the annotators is very high, as between their decisions we measured a Pearson correlation of 0.93 and a Root Mean Square Error (RMSE) of 0.36. After the annotation, we collected the impressions of the annotators about the general quality of the related entities. These are the comments that we collected:

1. Separating the related entities with commas is confusing, especially considering that person names are presented as "<surname>, <name>". The resulting sequence becomes difficult to parse when the names of several people are in the list, e.g. ``Martin, Steve, Murphy, Eddie, Burton, Tim". Using a different separator or putting each entity on a separate line might be a better choice.
2. Most results appear to be outdated and only marginally related to the currently selected news. For example, for news clusters about sports most related entities are retired players of the sport mentioned in the news.
3. When no related entities are retrieved, it would be preferable not to display the related entities panel at all, instead of showing an empty box.




Figure 7: Results of the second round of related entities annotation.

Based on the comments of the annotators, we have focused on improving both the user experience and the quality of the results:

- We have changed the layout of the related entities in the panel: now, each entity is listed on a separate line (addresses reviewers remark 1)
- We decided to display in the related entities panel both the topical entities identified in the news collection (i.e., those which appear in the request that the DiversiNews interface sends to Ontotext) and the related entities retrieved by Ontotext (addresses reviewers remarks 2 and 3).

The result of this process is the current implementation of DiversiNews. We ran a second round of evaluation, under the same conditions as the first one, to assess the improvement of the related entities service after these additions. The results are shown in Figure 7. The histogram shows that now the panel always displays some entities (no case has been assigned a score of 0), and the perceived appropriateness of the displayed results is largely improved. Though in many cases the annotators are still very or somewhat dissatisfied (average relatedness in [1,2]), in 14 cases out of 31 (i.e., 45%) the entities are found to be sufficiently related. In half of these cases (7/14), the set of entities is annotated as at least very related.

Also for this second round of annotation, the agreement among the annotators is quite high. We have measured a Pearson correlation of 0.85 between the two annotators, and an RMSE between their decisions of 0.62. As in the previous case, we have asked the annotators to comment on the general quality of the presentation and the quality of the results. This is the feedback that we gathered:

- In some cases, the same entity appears more than once in the panel with a different surface form, e.g. “N.Y.” and “New York”;
- Some of the entities which are retrieved are too general and they do not appear to be especially useful for improving news understanding. In particular, it is very frequent to see names of countries or cities as related entities. If the reader is browsing a cluster of news about, say, “economy”, then browsing to a cluster related to “Portugal” or “France” would lead her outside her domain of interest;
- The fact that the panel displays a variable number of entities (between five and ten) is perceived as unsettling, as it compromises the uniformity of the browsing experience.

One conclusion that we can draw is that the selection of related entities based on their *popularity* in a pool of generic and encyclopedic knowledge does not necessarily align with the commonly accepted notion of *relevance*. On the other hand, it is expected that the diversity-aware ranking algorithm, developed in WP3, will be able to offer an alternative way of selecting related entities, which will be based again on identification of clusters inside the basic RENDER data layer. Integrating this method into the search for related entities might provide an interesting venue for evaluating relevance and relatedness from a user stand point, and for better comparing the quality of the results perceived by human evaluators.

4.2 Final Evaluation Plan

As we already discussed in D5.2.2, we came to the conclusion that evaluations carried out on synthetic benchmarks (e.g., [TAC, 2011]) do not always reflect real usage scenarios. For example, 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 designed some task-specific evaluation templates that will allow us to directly measure the quality of the generated summaries in the real-world setting provided by DiversiNews automatically categorized news. Figure 8 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. Different summarizers, including Google's, will be evaluated in order to select the best tool for the task. This evaluation will be carried out in a crowd-sourcing setting.

In parallel, we will carry out field studies with real users who will be asked to use DiversiNews and rate the experience in terms of perceived utility, ease of use and responsiveness of the interface to their commands.

The final results of this large-scale evaluation 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 ...

Now, consider this summary, and answer the following questions:

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 8: A template for the evaluation of summarization technology.

5 Conclusions and Future Work

We have discussed the latest additions to the DiversiNews platform, which has currently reached a stable state. DiversiNews now allows users to use different summarizers to generate a compact view of the currently selected news collection, and offers them the possibility to explore related news collections by means of the newly added “related entities” panel, in which relevant topics identified in the news are listed with “popular” related entities retrieved through RENDER data layer. In its current implementation, DiversiNews is the result of a close collaboration between Google, responsible of the summarization technology, JSI, which maintains the language technologies and the user interface, and Ontotext, which is responsible for the retrieval of related entities.

We have presented the results of an evaluation activity aimed at assessing the quality of the entities retrieved through the data layer and assess the possibility of pushing further the integration with the knowledge base. The evaluation has shown that better retrieval strategies need to be devised before deeper levels of integration can be explored.

We have outlined our plans for an in-depth evaluation of the quality of the generated summaries and their adherence to the user-specified constraints. The results of this evaluation, as well as field studies that will measure the usability and ergonomics of DiversiNews, will be the main topics of the coming D5.2.4.

References

[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.

[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.