

# Online Communities Support Policy-Making: The Need for Data Analysis

Roman Klinger<sup>1,\*</sup>, Philipp Senger<sup>1,\*</sup>, Sumit Madan<sup>1</sup>, and Michal Jacovi<sup>2</sup>

<sup>1</sup> Fraunhofer Institute for Algorithms and Scientific Computing (SCAI)  
Schloss Birlinghoven  
53754 Sankt Augustin, Germany  
{klinger,senger,madan}@scai.fraunhofer.de

<sup>2</sup> Social Technologies Group  
Information & Social Analytics Dept.  
IBM Haifa Research Lab  
Mt. Carmel, Haifa 31905, Israel  
jacovi@il.ibm.com

**Abstract.** Policy decisions in governmental models are often based on their perception and acceptance in the general public. Traditional methods for harvesting opinions like telephone or street surveys are time intensive and costly and direct interaction between a governmental member and the population is limited. Social media harbor the chance to easily get a high number of opinions and proposals in form of poll participation or interactive debate contributions.

Especially debates about political topics can generate data which are hard to interpret because of its length and complexity. We propose a collection of methods to support a decision maker in gaining an overview over textual debates coming from several social media to save time and effort in manual analysis. Our approach enables an efficient decision making process by a combination of automatic topic clustering, sentiment analysis, filtering, and search functionalities aggregated in a graphical user interface. We present an implementation and a use case proving the usefulness of the proposed methodologies.

## 1 Introduction

Decision-making processes for policies and their outcomes are often based on their perception and acceptance in the general public. An approach to gain a more representative opinion for a focused theme as well as to acquaint the public with the topic and present it from several perspectives, is to build a platform in which a policy maker (e. g. a member of government) solicits contributions from the public for a specific topic. The core functionalities expected from such a framework are: structured polls, moderated debates, and the ability to provide access from as many social media sites as possible. Structured polls should support questions with multiple choice or free text answers. Moderated debates should allow the public to provide more detailed feedback while also being exposed to others' thoughts,

---

\* The first two authors contributed equally to this work.

enabling decision makers to learn about aspects and perceptions not thought of before. In order to gain access to many participants, the framework needs to provide access to people in social media like Facebook, Blogger, Twitter, and others.<sup>1</sup>

As these social media are very well inhibited, polls or debate questions can lead to a huge number of contributions. A decision maker cannot be expected to consume all comments, and thus the framework also needs to provide support for efficient sifting through long debates. The need and a proposal for such framework, including several methods, are described in this paper. We focus on the analysis of textual data in debates.

Following is Sect. 1.1, in which we briefly cover related work. Sect. 1.2 then describes an implementation of a framework for policy-making support in online communities. The fundamentals as well as the combination of methods to support a decision maker by increasing the use of harvested data are presented in Sect. 2. Results are explained by means of a use case in Sect. 3, and a summary is given in Sect. 4.

## 1.1 Related Work

The analysis of weblogs (blogs) is a topic currently heavily investigated. As an example, the EU project *SynC3* [18] is aiming at structuring the information of personal blogs to combine them with news information. The application *BlogPulse* [6] was developed to discover trends in a set of 100,000 weblogs. The output are key person names, phrases, and paragraphs. Qiazhu Mei *et al.* focused on the extraction of spatio-temporal data together with subtopics [13]. Monitoring the development in blogs is described by [14]. They cover technical issues like website extraction and cleaning as well. The important challenge of analyzing trends of opinions and sentiments is addressed by [15]. Teufl *et al.* proposed a clustering and graph-based framework [20] to limit the need for manual analysis. Maragoudakis *et al.* [10] reviewed different opinion mining methods and developed a framework to use them. The impact of social media to elections (especially in the Netherlands) is proven by [4]. The way how members of parliaments use Twitter for online discussions is investigated by [19].

Next to this work, some governments already have online platforms to get into contact with the population. An example is the platform *ePetition*<sup>2</sup> in Germany allowing for signing a petition which forces the government to discuss the topic at a specific number of signatures. For each petition, a discussion forum is available as well. In Greece, the platform *OpenGov*<sup>3</sup> allows for discussion of current laws in development.

Research in these fields is often concerned with methods from the area of information extraction and natural language processing. A short overview of such approaches can be found in [11].

---

<sup>1</sup> <http://www.facebook.com/>, <http://www.blogger.com/>,  
<http://www.twitter.com/>

<sup>2</sup> <https://epetitionen.bundestag.de/>

<sup>3</sup> <http://www.opengov.gr/>

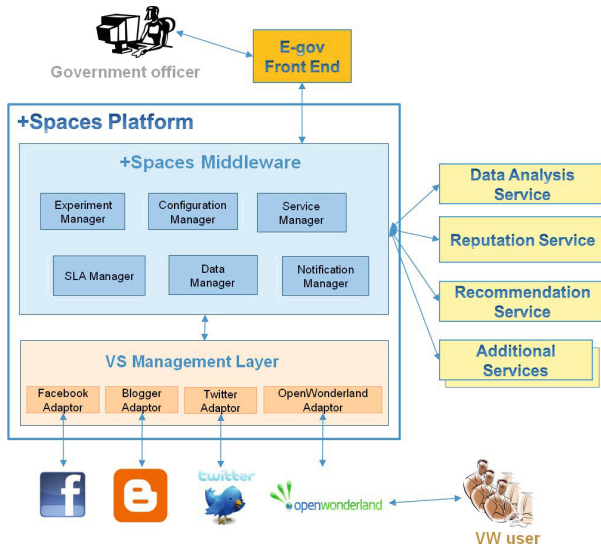


Fig. 1. Structure of the +Spaces platform

## 1.2 Policy-Making Support in Online Communities

The EC project +Spaces<sup>4</sup> (dubbed positive spaces) aims at developing a workflow to allow policy and decision makers to interact with the inhabitants of virtual spaces [22]. As presented in Fig. 1, a government officer accesses the platform through a single front end. She can *e. g.* assemble a poll consisting of several questions, each with several answer possibilities and/or a free text field. Another possibility is to initiate a discussion and provide an initial statement for that. While polls have the focus on getting feedback for a ‘closed’ question, the second possibility allows for learning about novel ideas and insights a government officer may not have thought of.

After designing the poll or debate statement (what we call experiment), the experiment is deployed to the virtual spaces using the middleware and the connected management layer. +Spaces is focusing on Facebook, Blogger, Twitter, and Open Wonderland<sup>5</sup>. The latter is an open source 3D world environment with similarities to Second Life<sup>6</sup> but the advantage of being deployable on self-administered servers and connectability due to source code availability.

Inhabitants of these worlds, namely people registered to Blogger, Twitter, or Facebook, can then participate in these experiments. Announcements are designed in a way such that viral dissemination is supported which is common in Facebook through ‘sharing’ and in Twitter through ‘re-tweeting’. This fact is supported by the numbers of participants of a pilot for poll experiments in

<sup>4</sup> <http://www.positivespaces.eu/>

<sup>5</sup> <http://www.openwonderland.org/>

<sup>6</sup> <http://secondlife.com/>

2011. Here, 77 participants made 473 contributions in different virtual spaces, 56 participants were recruited virally.

Through a notification mechanism, the data generated by the participants are propagated by the middleware to services to provide an analysis to the policy maker. Services as well as connectors to virtual spaces are modular and can be extended anytime. Figure 1 shows that one of the services developed in the +Spaces project is the proposed data analysis presented in this paper.

Our hypothesis is that such a way to prepare and present data, focusing on textual debates in this contribution, can support a policy maker by saving time and in getting an overview of the data. The fundamental idea is that by means of sentiments as well as main topics, debate contributions typically repeat themselves throughout a discussion—such characteristics need to be determined and shown to the user. In the following Sect. 2, the fundamentals of such methods and our adaptations are explained.

## 2 Methods

In the following, the methods for the analysis of semi-structured debates from online communities are described. They are designed to prepare the support of a policy maker.

### 2.1 Topic Modeling

The fundamental idea of modeling the topics of a debate is to present the main themes which are occurring. Additionally, the main words describing a topic are extracted. We follow two different strategies here, a  $k$ -means clustering [9] followed by an extraction of most important phrases; and alternatively a joint approach using latent Dirichlet allocation (LDA) [2]. The implementations from the MALLET toolkit [12] are used.

Typically for topic modeling,  $n$  grams or single tokens are used as input. To support a good understanding of the textual context, we are using noun phrases instead or optionally in addition. To limit the dimensionality, each token is previously transformed to its stem<sup>7</sup>. Stop-words are removed as well as URLs, email addresses, and numbers.

**$k$ -means Clustering.** The clustering method  $k$ -means [9] is an iterative approach to assign instances  $d_i$  ( $0 \leq i < n$ , with number of instances  $n$ ) to a given number  $k$  ( $k \in \mathbb{N}^{>0}$ ) of clusters. All instances in one cluster should have a high similarity with respect to some metric  $m(d_q, d_r)$ . Each instance is corresponding with a debate contribution in our case.

In short, the Voronoi iterations to find  $k$  clusters work as follows: Randomly, the instances  $d_i$  are assigned to  $k$  clusters. With respect to the metric  $m$ , the cluster centers are computed. Then, each instance is assigned to the closest cluster

---

<sup>7</sup> Using the Snowball Stemmer <http://snowball.tartarus.org>

center, followed by re-computation of the cluster centers. This iterative algorithm stops when the clusters are stable or after a specified number of iterations.

The cosine similarity measure on the tf-idf-weighted term vector space is used as metric  $m$  here. Each instance is represented by a vector of weights  $\mathbf{w}_l$  for the occurring tokens in the  $l$ th instance. The weight  $w_{k,l}$  for the  $k$ th token in the  $l$ th instance is  $w_{k,l} = \text{tf}_{k,l} \cdot \text{idf}_k$ , where  $\text{tf}_{k,l}$  is the frequency of the  $k$ th token (*term frequency*) in the  $l$ th instance normalized by the frequency of the most frequent term in that instance, which is a local measure. The global measure *inverse document frequency* is the logarithm of the number of instances by the number of instances with the  $k$ th term [1].

The cosine similarity measure  $m(d_q, d_r) := \frac{\mathbf{w}_q \cdot \mathbf{w}_r}{|\mathbf{w}_q| \cdot |\mathbf{w}_r|}$  is the degree between two instances  $d_q$  and  $d_r$  with weight vectors  $\mathbf{w}_q$  and  $\mathbf{w}_r$  in that vector space [1]. To detect the most important terms representing all instances in each of the  $k$  clusters, we use the highest ranked terms with respect to tf-idf. Note that these tf-idf values are not the same as in the clustering: In contrast to measuring the similarity of contributions, similarities of clusters are taken into account here. Therefore, all terms in one cluster are handled equivalently and  $0 \leq l < k$ .

**Latent Dirichlet Allocation.** The basic idea of latent Dirichlet allocation is that instances are represented as random mixtures over latent topics, where each topic is characterized by a distribution over terms [2]. Again, we allow to use stemmed noun phrases. LDA combines the two steps of clustering and extraction of keywords presented in the previous section in a joint fashion. As all documents are assigned to several topics, we report the most probable topic only. The values of the parameters of the LDA implementation in the MALLET toolkit are adapted as described by [17].

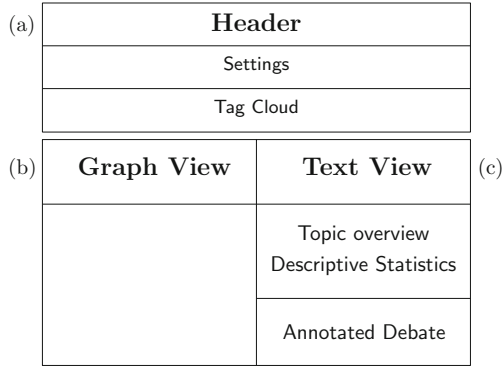
## 2.2 Sentiment Analysis

Sentiment analysis is the assignment of an expressed sentiment to a text fragment. Typically, the classes *positive* and *negative*, and *neutral* are used [5,7,8]. Most systems incorporate dictionary-based features; in the most straight-forward case, string matching with word lists with positive and negative connotation. An example for such system incorporating dictionaries next to other methods applied on online debate data has been presented by [16].

We apply a dictionary-based approach using the word lists of 6859 words (4818 negative, 2041 positive) provided by Hu and Liu<sup>8</sup> [7]. Let  $D^+$  be the set of positive words and  $D^-$  the set of negative words. A sentiment score  $\text{sent}(d_i)$  for a textual contribution  $d_i$  (where  $t_{ik}$  is the  $k$ th token in  $d_i$ ) is  $\text{sent}(d_i) = \sum_k s_{ik}$ , where  $s_{ik} = 1$  if  $t_{ik} \in D^+$  and  $s_{ik} = -1$  if  $t_{ik} \in D^-$ .

As described in Sect. 2.1, the contributions are clustered into topics of similar content to provide the user with an overview what the debate is about. To

<sup>8</sup> <http://www.cs.uic.edu/~liub/FBS/opinion-lexicon-English.rar>, accessed 1st Dec 2011.



**Fig. 2.** Schematic structure of the graphical user interface for the interactive debate visualization

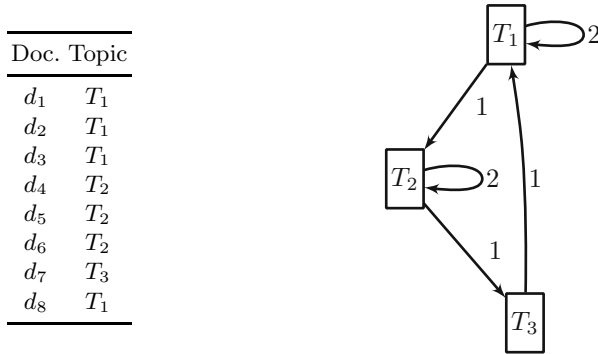
enhance that with the associated sentiment, we assign the sentiment score of the documents of topic  $T_i$  by  $\widehat{sent}(T_i) = \frac{1}{|T_i|} \sum_{d_i \in T_i} sent(d_i)$ , where  $\frac{1}{|T_i|}$  is a normalization factor. The combination of sentiments for LDA could have been implemented by taking the probabilities of tokens representing a specific topic into account as well. A drawback would be the limited transparency of the approach to the user.

### 2.3 Concepts of User Interaction

In Sect. 2.1 and 2.2, the methods to support a user in analyzing a textual debate for a specific theme have been introduced and explained. The results of these analyzes need to be shown to a reader or decision maker in an intuitive way. We implemented a web service based on a relational database harvesting the necessary data. This database acts as a temporary storage of the clustering results and updates the content in real time when a new debate contribution is included. Figure 2 shows a schematic overview of the web-based interface, divided into two main sections: The header of the page (a) acts as a common part for automatic summarization of the debate and the possibility to parametrize topic modeling and the sentiment analysis. The lower part is divided into the graph view (b) and the textual view (c) onto the debate.

An introduction of the debate content is given by a tag cloud of the main phrases in the debate and a depiction of demographic data of the participants. The selection of methods and parameters presented in Sections 2.1 and 2.2 is user-specified.

In the graph view section (Fig. 2(b)), a graph consisting of vertexes  $v_i$  for each cluster  $T_i$  and directed edges  $e_j = (v_k, v_l)$  is shown. Such edge is introduced if and only if  $\exists d_m \in T_k$  and  $\exists d_n \in T_l$  such that  $d_m$  and  $d_n$  are directly succeeding in the debate thread. A weight denoting the number of succeeding contribution pairs is additionally attached to that edge. In that way, the graphical depiction



**Fig. 3.** Example debate thread with clustering and graph depiction

summarized the linear structure of the thread by means of topics. It makes loops in contributions of same topic or between different topics clearly observable. An example of a debate thread together with its graph view is shown in Fig. 3. Here, topic  $T_1$  as well as  $T_2$  consist of three contributions, while two are answering without changing the topic. The contribution with topic  $T_3$  is leading back to  $T_1$ , while there is only one transition between all different clusters of topics.

The textual view (Fig. 2(c)) consists of two different sections. Firstly, the found topics are presented together with additional information, *i. e.* the top words describing the topic of the cluster, the number of debate contributions  $|T_i|$  in the  $i$ th topic and a normalized sentiment score  $\widehat{sent}(T_i)$ . The second subsection of the textual view consists of the annotated text itself. The information added to the text is the annotation of words associated with a sentiment, the sentiment score of each contribution and a highlighting of the topic-distinguishing words of each cluster.

This view needs to be highly interactive by means of different filtering possibilities: The user can select a topic to limit the shown debate distribution to those in the topic. Additionally, it needs to be possible to show only contributions associated with a positive or negative sentiment or to filter the contributions by a specific word, be it freely specified or selected from the tag cloud.

### 3 Results and Discussion

#### 3.1 Experimental Setup

As described in Sect. 1.2, one aspect of the +Spaces project is to create a platform for civil servants and policy makers which provides an easy to use and clearly arranged graphical user interface for practical usage within a political context. To reach this goal, an example debate has been annotated in cooperation with the policy makers for analyzing the usefulness of the presented approaches.

A debate about the current Greek financial crisis was selected [21] and annotated by two members of a focus group. This group has the function to evaluate

debt americans civil democracy europe germans government  
 junta papandreou people political retirement taxes work  
 greece

Fig. 4. Frequency based tag cloud of the evaluation debate

Show	Topics	Most relevant words	Number of contributions in this topic	Specificity of relevant words in topic	Averaged sentiment of topic
<input checked="" type="checkbox"/>	Topic 1	working, germans, people, weeks, figures, products, vacation, timents, average, good, don, laws, job, receipts, europe, europeans, book, sector, comments, paid, retirement, fact, germans, starting, businesses	16	64.1%	1.19
<input checked="" type="checkbox"/>	Topic 2	greek, greece, countries, americans, days, problem, mania, taxes, while, government, place, credit, speaks, years, truth, public, person, debts, paying, developers, mentality, make, veri, home, question	15	72.3%	0.0
<input checked="" type="checkbox"/>	Topic 3	papandreou, years, junta, constitution, grandfather, power, economy, fires, civil, tomaski, prime, george, gap, crisis, michael, return, coup, parliament, democraci, mans, taking, parti, good, war, ministri	9	63.9%	-1.11
<input checked="" type="checkbox"/>	Topic -	contribution not classified!	-	-	-

Fig. 5. LDA clustering results of the evaluation debate as list

and discuss the results coming from the +Spaces consortium and the members have a strong political background. The chosen political debate has a length of 39 contributions coming from 24 different participants. The total number of words in this debate is 3994, the average contribution length is 102, the standard deviation is 90.

The task of the annotators was to underline the most important terms, as well as to assign a sentiment. For simplicity, one of three sentiment classes (*positive*, *negative*, *neutral*) were attached to each contribution. This task turned out to be complex for the special case of political debates; the inter-annotator agreement determined via Cohen’s kappa [3] for the sentiment annotation is not substantial.

### 3.2 Header Section

The top part of the graphical user interface contains a tag cloud providing a first overview of the debate’s content (*cf.* Fig. 4). The main tags (size is coupled logarithmically to frequency) show that the content of the debate is a political discussion concerning Greece and Greek people, the government and taxes. The tag “debt” provides an indication that the debate is about a financial topic.

### 3.3 Textual View Results

Figure 5 exemplifies the structured results of LDA applied to the debate introduced in Sect. 3.1. Each row in the table corresponds to a detected topic. The first column allows the user for filtering the debate contributions only showing the ones from the specified topic. The third and main column shows representative words and phrases together with the color used to highlight them in the debate text. The fourth column shows the number of debate contributions in this topic. This gives the user an impression of how dominant a specific topic



**Debate description**

Greece's fiscal problems is not exactly my great area of expertise, but when you're invited to meet an actual head of state, you go. And so I was one of about 17 journalists who met with George Papandreou this morning at the Center for American Progress. Here's the Reuters news write-up of the session: WASHINGTON, March 10 (Reuters) - Greek Prime Minister George Papandreou said on Wednesday his country was working to get its fiscal house in order and wanted markets to recognize that. Speaking at the end of his first visit to the United States as Greece's new leader, Papandreou said his government was not trying to "scapegoat" its problems by blaming them on market speculators. Papandreou said European leaders such as German Chancellor Angela M [...] more

Moderator: +Spaces Moderator, Posted at: 2011-11-15 16:15:06

**Written posts from Facebook and/or Blogger**

ah the joys of pedantry! the Greek head of state is Karolos Papoulias. **george papandreou** is the **prime minister**. Perhaps your next quiz ought to be on the **grecian**<sup>1</sup> form of titles and honorifics?

Contribution 1 posted by user 16159 in Facebook at: 2011-11-15 16:20:28, Topic 3, Reputation: ★★☆☆☆, Sentiment Score: 1 (1, 0)

Indeed. Apparently Mr. Tomasky should have begun with "Greece's HISTORY is not exactly my **area**<sup>1</sup> area of expertise" (either). But why **bother**<sup>1</sup> reading up before launching into a relevant article. I'm 32, born in the US, but still consider **greece** very much "home", and **while** I'm no fan of the **lack**<sup>1</sup> of financial disciplines of this and previous **greek** administrations (and certainly no fan of the lineage of Papandreou - in my opinion they shouldn't be let anywhere near Athens, but that's another **rant**<sup>1</sup>). While the rest of Europe was attempting to rebuild, **greece** was driven by internal and external forces rove directly into Civil War. The resulting tensions came to be known as the Cold War. To **ignore**<sup>1</sup> Mr. Papandreou's tepidness in the context of these historical sentiments of the **greek** populous against the Germans, Russians, Brits, and **americans** is short-sighted on your part (though admittedly many a book has been written on the matter, so a short article isn't the **pleist** to lay **assess**<sup>1</sup> to such content). I don't have the **decent**<sup>1</sup> at the moment but you should really look into the post-WW II events of **greece**. You just glanced over it **leading**<sup>1</sup> into the 1960's junta that drove hundreds of thousands of **greek**s into the international diaspora. In a very **biased**<sup>1</sup> way, articles **like**<sup>1</sup> this **make** me **cringe**<sup>1</sup>.

Contribution 2 posted by user 16159 in Facebook at: 2011-11-15 16:20:41, Topic 2, Reputation: ★★☆☆☆, Sentiment Score: -3 (4, 7)

Apparently it's s.o.p. for **businesses** not to write **people** **receipts**, that way the transaction is off the **books** for all involved. When that's your **starting** point, you've got a **job** ahead of you. I think if that's your **starting** point it really helps. When Italy started making **people** write **receipts** a few years ago revenues shot up overnight(as did restaurant prices). It's **good**<sup>1</sup> to have low hanging fruit.

Contribution 3 posted by user 16159 in Facebook at: 2011-11-15 16:20:50, Topic 1, Reputation: ★★☆☆☆, Sentiment Score: 1 (1, 0)

M, **lucky**<sup>1</sup> you to attend such a function! Papandreou was **educated**<sup>1</sup> abroad, which has influenced him. He **speaks** several languages. So far so **good**<sup>1</sup>. Don't forget that speaking softly is **polite**<sup>1</sup> in most societies, and by speaking softly people actually will have to be more **able**<sup>1</sup> to listen and to concentrate. This can be seen as **modesty**<sup>1</sup> or as **clever**<sup>1</sup>. He has a huge task on his shoulders and he needs an extremely **sharp**<sup>1</sup> carving **knife**<sup>1</sup> to surgically alter the mindset of the average **greek** **paying** **taxes**, keeping **precise**<sup>1</sup> books **correctly**<sup>1</sup> and contributing your dues should be standard **issue**<sup>1</sup> and that is not exactly the case in **greece**. The **greek** national sport is to avoid **taxes** **like**<sup>1</sup> the **plague**<sup>1</sup> and everybody is in on it. The collection system of **taxes** is not in working order, and sales of anything does get excused under the counter (no receipts or trace of transaction). It is an **impossible**<sup>1</sup> situation, and the **mentality** has nothing to do with the colour of the leader of any political party. The **problem**<sup>1</sup> is the **mentality** of the **greek** people. You should hear what **developers** are up to in **greece**!!!!!! Causing fires, so that vast tracts of land and property gets set alight and voila, the **developers** turn up to **make** a bid and take over the **piece**. **While** development is officially **prohibited**<sup>1</sup>. Strict laws should be created and the policing of these laws should be done by non-Greeks. No backhanders, no bribes. Those **days** should be gone. Nor the Germans, nor the French will bail out the **greek**s. Papandreou has been told this, **while** touring Europe. The **greek**s have to solve this **mess**<sup>1</sup> themselves, a **mess**<sup>1</sup> of their own making. It has nothing to do with the **credit** crunch or the collapsing banking system or a political party. The root of the **problem**<sup>1</sup> is the **mentality** with regards **taxes** and keeping a **precise**<sup>1</sup> record.

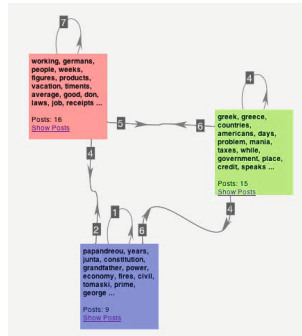
**Fig. 6.** Automatically annotated text of the evaluation debate [21] with the additional pseudonymized information about the user, the virtual space, the topic membership, the time stamp, and the sentiment score

is in comparison to the others. The fifth column shows the ratio of the number of found relevant words in this topic and in the whole debate. This value can be understood as specificity of the topic.<sup>9</sup> The last column shows the overall sentiment score of the contributions in the specific topic.

LDA automatically identified three different topics in this case, while a fourth cluster contains contributions that are unspecific (no important words found). 24 of 74 unique relevant words defined by annotators to be of interest (32.43 %) are detected. The *k*-means algorithm with tf-idf ranking determinates 15 of the relevant words (20.27 %). The three topics give a good insight in the main parts of the discussion: Topic 1 may be called *Political Context in Europe*, the 2nd *Financial and Business Issues* and the 3rd *Greek Opinions*.

With the knowledge of the main topics, the user may want to have a closer look at the detailed results of the topic modeling and the sentiment analysis according to the debate's text. Figure 6 exemplifies the contribution-based layout of this section. The important information provided to the user are the time when the contribution was written, the user information, the assigned topic id and the

<sup>9</sup> If the ratio is close to 1, nearly all relevant words are only mentioned in the topic specific contributions.



**Fig. 7.** Screenshot of the graph view of the evaluation debate with LDA

sentiment score of the particular contribution. The text itself is annotated with relevant words using the appropriate color. The topic assignment is stated by the topic id and the colored bar on the left border of the contribution.

In addition to the annotation of all relevant words of each topic, the textual view is enriched with the sentiment score. The annotations in the text are presented with a colored underscore and a  $[+]$  for a word with positive and  $[-]$  for a word with negative sentiment.

### 3.4 Graph View Results

Figure 7 shows the visualization of the resulting graphs based on LDA. As described in Sect. 2.3, the number of nodes is equivalent to the number of topics and the labeled edges encode the topics and their transition within the succeeding contributions. In  $k$ -means clustering, the number of topics is exactly as specified by the user, in contrast, LDA uses three clusters and leads to a well-arranged graph. The colorization of the nodes is associated with the colors in Fig. 5. The most relevant words of each topic are presented inside of each node. A link to the textual view of the contributions within this cluster is provided. The graph view enables the user to analyze the structure of the debate in a very condensed way.

## 4 Summary and Future Work

Analysis of textual debates from online communities and presenting them in a way that is clear and valuable for policy makers is a challenging task. In this paper we presented our approach to this challenge, incorporating topic modeling and sentiment analysis, and a web-based implementation of innovative visualizations for presenting the results in an easily perceivable way.

An overview of a debate is presented as a frequency weighted tag cloud. The determination of the main topics along with their most important words allows policy makers to get deeper understanding into a long debate, especially

in the graphical view and the annotated text with different colors. Each topic is additionally assigned with the average sentiment score to not only show most important phrases but the associated emotions as well.

The manual annotation of data is a difficult task, as shown by the limited agreement of two annotators. Nevertheless, we could show on an example debate that despite of the comparatively low compliance with manual annotation, the automated approach can lead to understandable and helpful results. How to evaluate such unsupervised methods is still question of research: While an annotator may find some clustering meaningful, another automatically detected one can as well be helpful while being less obvious.

Future work will focus on the execution of a pilot of the whole +Spaces platform, specifically for debates. We assume to retrieve a large amount of real world data and will optimize our approach on such contributions coming from Facebook, Blogger, and Twitter. Presumably, the language used in political debates performed on such platforms differs from texts on other platforms. The same holds for other topics, depending on the participating users and technical limitations, like the limited lengths of texts.

**Acknowledgements.** This publication is funded in the context of the Positive-Spaces project (<http://www.positivespaces.eu/>) by the European Community's Seventh Framework Programme [FP7/2007-2011] under grant agreement no. 248726. We thank all other partners in the consortium of the project as well as the focus group in the Hellenic Parliament.

## References

1. Baeza-Yates, R.A., Ribeiro-Neto, B.: Modern Information Retrieval. Addison-Wesley Longman Publishing Co., Inc., Boston (1999)
2. Blei, D.M., Ng, A.Y., Jordan, M.I.: Latent dirichlet allocation. *Journal of Machine Learning Research* 3, 993–1022 (2003)
3. Cohen, J.: Weighted kappa: Nominal scale agreement provision for scaled disagreement or partial credit. *Psychological Bulletin* 70(4), 213–220 (1968)
4. Effing, R., van Hillegersberg, J., Huibers, T.: Social Media and Political Participation: Are Facebook, Twitter and YouTube Democratizing Our Political Systems? In: Tambouris, E., Macintosh, A., de Bruijn, H. (eds.) ePart 2011. LNCS, vol. 6847, pp. 25–35. Springer, Heidelberg (2011)
5. Esuli, A., Sebastiani, F.: Determining term subjectivity and term orientation for opinion mining. In: Proceedings of the Conference of the European Chapter of the Association for Computational Linguistics, EACL (2006)
6. Glance, N.S., Hurst, M., Tomokiyo, T.: Blogpulse: Automated trend discovery for weblogs. In: WWW 2004 Workshop on the weblogging ecosystem: Aggregation, Analysis, and Dynamics. ACM (2004)
7. Hu, M., Liu, B.: Mining and summarizing customer reviews. In: Proceedings of Knowledge Discovery in Databases Conference. ACM, Seattle (2004)
8. Kim, S.-M., Hovy, E.: Determining the sentiment of opinions. In: Proceedings of the COLING conference, Geneva, Switzerland (2004)
9. MacKay, D.: Information Theory, Inference and Learning Algorithms. In: An Example Inference Task: Clustering, ch. 20, pp. 284–292. Cambridge University Press (2003)

10. Maragoudakis, M., Loukis, E., Charalabidis, Y.: A Review of Opinion Mining Methods for Analyzing Citizens' Contributions in Public Policy Debate. In: Tambouris, E., Macintosh, A., de Bruijn, H. (eds.) ePart 2011. LNCS, vol. 6847, pp. 298–313. Springer, Heidelberg (2011)
11. McCallum, A.: Information extraction: Distilling structured data from unstructured text. *Queue* 3, 48–57 (2005)
12. McCallum, A.K.: MALLET: A Machine Learning for Language Toolkit (2002), <http://mallet.cs.umass.edu>
13. Mei, Q., Liu, C., Su, H., Zhai, C.: A probabilistic approach to spatiotemporal theme pattern mining on weblogs. In: Proceedings of the 15th International Conference on World Wide Web, WWW 2006, pp. 533–542. ACM, New York (2006)
14. Nanno, T., Fujiki, T., Suzuki, Y., Okumura, M.: Automatically collecting, monitoring, and mining japanese weblogs. In: Proceedings of the 13th International World Wide Web Conference on Alternate Track Papers & Posters, WWW Alt. 2004, pp. 320–321. ACM, New York (2004)
15. Pang, B., Lee, L.: Opinion mining and sentiment analysis. *Found. Trends Inf. Retr.* 2, 1–135 (2008)
16. Somasundaran, S., Wiebe, J.: Recognizing stances in ideological on-line debates. In: Proceedings of the NAACL HLT 2010 Workshop on Computational Approaches to Analysis and Generation of Emotion in Text, Los Angeles, California, USA (June 2010)
17. Steyvers, M., Griffiths, T.: Probabilistic topic models. In: *Latent Semantic Analysis: A Road to Meaning*. Lawrence Erlbaum (2007)
18. sync3: Eu project sync3 (2011), <http://www.sync3.eu/>
19. Sæbø, Ø.: Understanding Twitter<sup>TM</sup> Use among Parliament Representatives: A Genre Analysis. In: Tambouris, E., Macintosh, A., de Bruijn, H. (eds.) ePart 2011. LNCS, vol. 6847, pp. 1–12. Springer, Heidelberg (2011)
20. Teufel, P., Payer, U., Parycek, P.: Automated Analysis of e-Participation Data by Utilizing Associative Networks, Spreading Activation and Unsupervised Learning. In: Macintosh, A., Tambouris, E. (eds.) ePart 2009. LNCS, vol. 5694, pp. 139–150. Springer, Heidelberg (2009)
21. The Guardian: Greece is the word (2011), <http://www.guardian.co.uk/commentisfree/michaeltomasky/2010/mar/10/greece-usa-breakfast-with-papandreou>
22. Tserpes, K., Jacovi, M., Gardner, M., Triantafillou, A., Cohen, B.: +spaces: Intelligent virtual spaces for e-government. In: Proceedings of International Conference on Intelligent Environments (IE 2010), Guanajuato, Mexico (June 2010)