EuroVis Workshop on Visual Analytics (2015)
E. Bertini and J. C. Roberts (Editors)

Visual-Interactive Text Analysis to Support Political Decision Making – From Sentiments to Arguments to Policies

T. Ruppert[1], J. Bernard[1], H. Lücke-Tieke[1], T. May[1], J. Kohlhammer[1,2],

[1]Fraunhofer IGD, Germany, [2]GRIS, TU Darmstadt, Germany

Abstract

Political decision making involves the evaluation of alternative solutions (so called policy models) to a given societal problem and the selection of the most promising one. Large amounts of textual information to be considered in decision making processes can be found on the web. This includes general information about policy models, individual arguments in favor or against these policies, and public opinions. Monitoring large text collections and extracting the relevant information is time consuming. In this approach we present a visual analytics system that supports users in assessing the results of automatic text analysis methods. The methods extract text segments from large document collections and associate them with predefined policy domains, policy models, and policy arguments. Moreover, sentiment analysis is applied on the text segments. Visualization techniques provide non-IT experts an intuitive access to the results. With the system, users can monitor public debates on the web. In addition, we analyze concepts that enable the user to give visual-interactive feedback on the text analysis results. This direct user feedback can help to improve the accuracy of individual text analysis modules and the credibility of the overall text analysis process. The system was tested with real users from the political decision making domain.

1. Introduction

Today, policy makers are requested to integrate large amounts of external knowledge and public opinions in their decision making processes. Large parts of the information to be considered are available on the web. However, the manual monitoring and analysis of this data is time consuming and therefore not applicable in real-world scenarios. Automatic methods for the mining and analysis of textual content exist. To make use of these powerful tools, some challenges need to be tackled. First, the methods need to be combined to a workflow and adapted to the addressed domain. Second, the results of the workflow need to be presented to policy makers in an intuitive way. Third, since the accuracy of text analysis methods is not necessarily satisfying the users' expectations, concepts for feedback loops need to be considered. Our contributions to tackle these challenges are:

1. As a baseline for our approach, we describe a text analysis workflow targeted on the domain of political decision making. The workflow applies text analysis methods that extract segments from a crawled document collection and associates them with predefined political concepts – policy models and arguments.
2. We present a visualization dashboard designed for the presentation of text analysis results. The dashboard helps policy makers to access the results in an intuitive way.
3. We extend the workflow with visual-interactive feedback concepts that enable the users to improve the accuracy of the text analysis results. As a result, we increase the credibility of the system.
4. We implemented our approach in a real-world environment during a European research project to proof its applicability for political decision making.

2. Related Work

Apart from the concept of supporting political decision making with data visualization [KNRB12] [RBLTK14] related work to our approach can be found in the fields of text analysis and visual text analysis.

Text Analysis. A general introduction of text mining (also including some visualization examples) is presented by Feldman and Sanger [FS06]. Liu provides a comprehensive work about data mining techniques for the extraction and analysis of textual data from the web. This includes topics like crawling, opinion mining, and sentiment analysis [Liu06]. A general introduction to opinion mining and

DOI: 10.2312/eurova.20151101

sentiment analysis is provided by Pang et al. [PL08]. Argumentation mining as a research field is relatively young. Relevant works can be found in Teufel's thesis [T*00], and further approaches presented by Palau et al. [PM09], and Feng and Hirst [FH11]. The state of the art report by Jones describes recent approaches about the automatic summarization of textual documents [SJ07].

Visual Text Analysis. During the design of our visualization dashboard we followed Few's suggestions for the visualization of quantitative data [Few09]. Prominent examples for visual analytics approaches related to sentiment analysis are proposed by Liu et al. [LHC05] and Chen et al. [CISSW06]. Oelke et al. visualize feature-based opinion clusters for Amazon products [OHR*09]. The Document Cards approach describes a technique for the summarization of single documents [SOR*09]. A comprehensive work about practical techniques for visualizing arguments is introduced by Kirschner et al. [KBSC03]. A visual analytics approach for analyzing social media content from Twitter and Youtube is presented by Diakopoulos [DNKS10]. Although these approaches are related to ours, most of them are mainly monitoring tools, feedback concepts are not considered.

3. Approach

Our approach tackles tasks that evolved during a European research project together with political decision makers. The tasks can be summarized as follows:

T$_1$: Extract the relevance of policy models, policy arguments, and policy terms in online discussions.
T$_2$: Analyze their relevance over time and per source.
T$_3$: Analyze the sentiment of extracted text segments.
T$_4$: Get access to the original textual content and sources.
T$_5$: Identify new arguments.

Our approach operates on individual text segments extracted from a large document collection. A text segment comprises one to several consecutive sentences. A query for extracting the documents and the segments from the collection is created from a concept graph representing user interests. This predefined graph relates policy domains (e.g. energy, transport, etc.), policy models (e.g. renewable energy directive, etc.), and arguments. The role of the concept graph is twofold. Firstly, it implicitly defines a document query by using search keywords from the political concepts. Secondly, it is used to structure the text segments, based upon the user's understanding of the policy domain. For more details about the graph and its editing process we refer to Spiliotopoulos et al. [SDK14]. The document collection is crawled from the web and comprises textual statements from news paper sites, social media platforms, blogs, etc.

4. Text Analysis Workflow

We present a text analysis workflow (see Fig. 1) that was designed and implemented to tackle the tasks described in the previous section. The individual text analysis modules that constitute the workflow are explained in the following. More details about the underlying linguistic pipeline are provided by Komourtzis et al. [KGP*14]. As described above a crawled document collection and a predefined concept graph serve as a prerequisite for our approach.

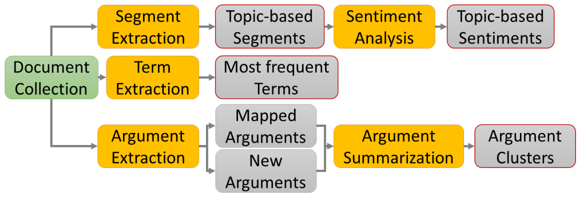

Figure 1: *Text analysis workflow including text analysis modules (orange) and intermediate results (grey).*

Segment Extraction. The segment extraction module analyses the document collection and extracts text segments related to the definitions provided by a given policy model, or argument (**T**$_1$). The module classifies text segments as associated or not associated with the given political concept.

Sentiment Analysis. The sentiment analysis module operates on two distinct modes. In the first case, the sentiment score of a textual document is calculated. In the second case, the module extracts the sentiment from individual segments with respect to the associated political concept. Hence, for an whole document the general sentiment scores are calculated. For the individual segments the topic-based sentiment scores are calculated (**T**$_3$). The applied sentiment analysis methods were introduced by Petasis et al. [PSTT14].

Argument Extraction. The argument extraction module identifies argumentative sentences (**T**$_1$,**T**$_5$). Text segments that pertain an argument structure, e.g., containing claims and premises, are extracted from the document collection and classified as arguments. In a second step the extracted arguments are mapped on the predefined arguments in the concept graph. The argument extraction methods applied in this approach are presented by Goudas et al. [GLPK14].

Argument Summarization. The main purpose of the argument summarization is to extract arguments that are not yet defined in the concept graph (**T**$_5$). It shall help the users to improve the existing policy models. The module forms argument clusters either based on (1) existing mappings of extracted arguments on predefined arguments (cf. argument extraction), or (2) based on textual similarities. In the latter case, a representative is chosen that summarizes the content of the new argument cluster.

Term Extraction. The term extraction module discovers the most frequent terms found in a document collection (**T**$_1$). The module is not restricted to terms directly relevant to the category (as these are more useful for classifying the content), but rather discovers and presents terms that are fre-

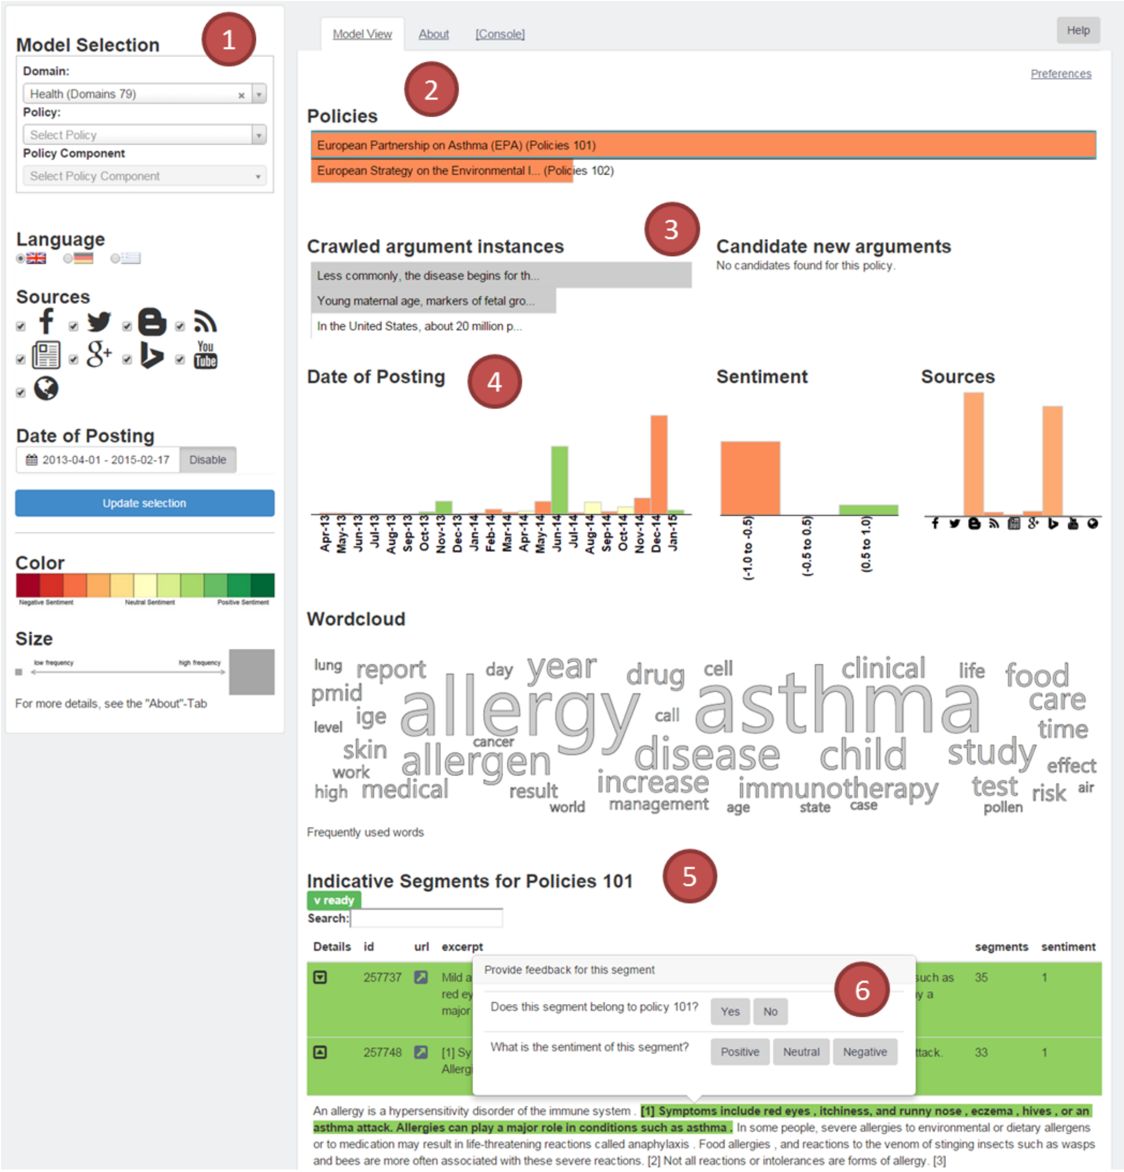

Figure 2: *Visualization dashboard.*

quently used within the context. (e.g. 'wind farm' is a term related to the domain 'energy', while 'noise', or 'efficiency' are terms that are common in discussions under the 'energy' category, thus they denote issues that must be taken into account when constructing a policy).

In summary, the text analysis workflow extracts text segments from a large document collection and associates these segments with the predefined policy models, and arguments. Moreover, argumentative segments are identified, mapped on predefined arguments, or clustered and marked as potentially new arguments. For all documents the general sentiment scores are calculated. For extracted text segments the sentiment score towards the associated policy model, or argument is calculated. Finally, for all subgroups of the document collection the most frequent terms are calculated.

5. Visualizing Text Analysis Results

The visualization module of our visual analytics system was designed as a dashboard with the goal to present the text analysis results to the users in an intuitive way (see Fig. 2). Since most of the users do not have an IT background, we chose familiar visualization techniques.

The dashboard is divided into three areas: a navigation panel (Fig. 2(1)), a statistics panel (Fig. 2(2)(3)(4)), and a text segment panel (Fig. 2(5)). As denoted by the legend, in all views the color reflects the sentiment score (green = positive sentiment, yellow = neutral sentiment, red = negative sentiment) ($\mathbf{T}_3$). The size of the visual objects reflects the number of extracted segments, hence, the concept's relevance ($\mathbf{T}_1$). The navigation panel consists of a hierarchical topic selection menu. The menu represents the political concept graph described in Section 3. The user can select (a) a policy domain to get details about underlying policy models, (b) a policy model to get details about underlying policy components (parts of a policy model), and (c) a policy component to get details about underlying arguments. These details are shown in the statistics panel. The underlying policy models, policy components, or arguments are displayed in a sorted bar chart (Fig. 2(2)). This enables the user to get a quick overview about the relevance of the political concepts ($\mathbf{T}_1$) and the overall sentiment ($\mathbf{T}_3$). An additional bar chart shows the extracted argument clusters separated into predefined (left) and potentially new clusters (right) (Fig. 2(3))($\mathbf{T}_5$). Further statistical information includes the temporal distribution of underlying text segments, the distribution per web source, and a sentiment distribution (Fig. 2(4))($\mathbf{T}_2$,$\mathbf{T}_3$). An additional word cloud provides users an idea about the discussed textual content. A tabular view provides the original textual documents including the extracted and highlighted segments to the user. (Fig. 2(5))($\mathbf{T}_4$). Finally, the queries can be refined based on language, web source, and date of posting filters.

6. User Feedback

In general, it cannot be assumed that the results of text analysis processes perfectly match with human understanding of a domain. To mitigate this problem, the user is able to refine the results by giving incremental feedback on policies, arguments, segments, sentiments and their proposed relations. Feedback is generally triggered from a sub-menu by selecting a corresponding visual representation. For all modules feedback can be given in at least two ways: (1) a general validation or approval of the implied relationship. (2) a manual correction. In any case, the user feedback is collected in a database for the refinement of the analytical models. Because the text corpus is too large for interactive adaption, the actual modification of the models is done in an offline process on a regular basis. For the ongoing session, the feedback only changes the specified modules. In the following, we will describe the type of feedback for every text analysis module presented in Section 4. An exemplary visual-interactive feedback concept is shown in Fig. 2(6).

Segment Extraction. For the validation of the segment extraction, a user can indicated that a text segment is, in fact, relevant for a policy model (cf. Fig. 2(6)). For a correction, a user may attach the text segment to another policy model.

Sentiment Analysis. For the sentiment analysis module the user is able to feed back whether the sentiment scores are correct or not. If this is not the case, the user may adjust the scores (cf. Fig. 2(6)). Corrected sentiment scores for documents or segment-topic-pairs are included into the training corpus of the sentiment analysis module.

Argument Extraction. Concerning the argument extraction module, a user has three options for feedback: Because not all segments might in fact be arguments, the feedback includes the validation whether a specific segment can be accounted for an argument at all. Arguments are identified by their similarity to predefined 'template' arguments. A user may specify whether this association is valid or possibly suggest another predefined argument. Finally, the user can rephrase a new argument and add it to the corpus of predefined arguments to capture a new aspect or to better distinguish between different predefined arguments.

Argument Summarization. With respect to the argument summarization module the user may approve the grouping of extracted arguments or remove outliers from their respective groups. In addition, the user may associate an similarity-based argument cluster with an existing predefined argument, or phrase a new argument, that describe the argument cluster and add it to the corpus of predefined arguments.

Term Extraction. As a possible user feedback, terms that are automatically extracted from a textual corpus can be excluded from the display. This might be feasible for terms that are obvious for a given domain and should not be highlighted anymore. As an example the term 'energy' does not provide any helpful insights in the energy domain, while the term 'efficiency' would. Therefore, the exclusion of terms from the most frequent term list would be a valuable user feedback that can improve the quality of the term extraction module. From a technical perspective, the terms to be deleted could be added to a user-defined stop word list.

7. Conclusion

In this approach we presented a visual text analysis system applied to the political decision making domain. The system extracts text segments from the web and associates them with predefined policy models and arguments. It combines a text analysis workflow with a visualization dashboard with the focus to facilitate the access to text analysis results. In addition, we introduced concepts that enable the user to provide direct feedback on results. These concepts help to improve the accuracy of individual text analysis modules and increase the credibility of our system. Our approach serves as a starting point for further research in visual text analysis applied to the political decision making domain.

Acknowledgements. This research has been partially supported by the EU 7th Framework Programme under grant agreement no. 611964 (EU Community).

References

[CISSW06] CHEN C., IBEKWE-SANJUAN F., SANJUAN E., WEAVER C.: Visual analysis of conflicting opinions. In *Visual Analytics Science And Technology, 2006 IEEE Symposium On* (2006), IEEE, pp. 59–66. 2

[DNKS10] DIAKOPOULOS N., NAAMAN M., KIVRAN-SWAINE F.: Diamonds in the rough: Social media visual analytics for journalistic inquiry. In *Symposium on Visual Analytics Science and Technology (VAST)* (2010), IEEE, pp. 115–122. 2

[Few09] FEW S.: *Now You See it: Simple Visualization Techniques for Quantitative Analysis*. Analytics Press, 2009. 2

[FH11] FENG V. W., HIRST G.: Classifying arguments by scheme. In *Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies - Volume 1* (Stroudsburg, PA, USA, 2011), HLT '11, Association for Computational Linguistics, pp. 987–996. 2

[FS06] FELDMAN R., SANGER J.: *Text Mining Handbook: Advanced Approaches in Analyzing Unstructured Data*. Cambridge University Press, New York, NY, USA, 2006. 1

[GLPK14] GOUDAS T., LOUIZOS C., PETASIS G., KARKALETSIS V.: Argument extraction from news, blogs, and social media. In *Artificial Intelligence: Methods and Applications - 8th Hellenic Conference on AI, SETN 2014, Ioannina, Greece, May 15-17, 2014. Proceedings* (2014), Likas A., Blekas K., Kalles D., (Eds.), pp. 287–299. 2

[KBSC03] KIRSCHNER P., BUCKINGHAM-SHUM S., CARR C.: *Visualizing Argumentation: Software Tools for Collaborative and Educational Sense-Making*. Computer Supported Cooperative Work. Springer, 2003. 2

[KGP*14] KIOMOURTZIS G., GIANNAKOPOULOS G., PETASIS G., KARAMPIPERIS P., KARKALETSIS V.: NOMAD: Linguistic Resources and Tools Aimed at Policy Formulation and Validation. In *Proceedings of the 9th International Conference on Language Resources and Evaluation* (2014), European Language Resources Association (ELRA), pp. 3464–3470. 2

[KNRB12] KOHLHAMMER J., NAZEMI K., RUPPERT T., BURKHARDT D.: Toward visualization in policy modeling. *Computer Graphics and Applications, IEEE 32*, 5 (Sept 2012), 84–89. 1

[LHC05] LIU B., HU M., CHENG J.: Opinion observer: analyzing and comparing opinions on the web. In *Proceedings of the 14th international conference on World Wide Web* (2005), ACM, pp. 342–351. 2

[Liu06] LIU B.: *Web Data Mining: Exploring Hyperlinks, Contents, and Usage Data (Data-Centric Systems and Applications)*. Springer-Verlag New York, Inc., Secaucus, NJ, USA, 2006. 1

[OHR*09] OELKE D., HAO M., ROHRDANTZ C., KEIM D., DAYAL U., HAUG L., JANETZKO H.: Visual opinion analysis of customer feedback data. In *IEEE Symposium on Visual Analytics Science and Technology, (VAST)* (2009), pp. 187–194. 2

[PL08] PANG B., LEE L.: Opinion mining and sentiment analysis. *Foundations and Trends in Information Retrieval 2*, 1-2 (2008), 1–135. 2

[PM09] PALAU R. M., MOENS M.-F.: Argumentation mining: the detection, classification and structure of arguments in text. In *Proceedings of the 12th international conference on artificial intelligence and law* (2009), ACM, pp. 98–107. 2

[PSTT14] PETASIS G., SPILIOTOPOULOS D., TSIRAKIS N., TSANTILAS P.: Sentiment analysis for reputation management: Mining the greek web. In *Artificial Intelligence: Methods and Applications - 8th Hellenic Conference on AI* (2014), Likas A., Blekas K., Kalles D., (Eds.), vol. 8445 of *LNCS*, Springer, pp. 327–340. 2

[RBLTK14] RUPPERT T., BERNARD J., LÜCKE-TIEKE H., KOHLHAMMER J.: Towards a tighter coupling of visualization and public policy making. In *Visual Analytics Science and Technology (VAST), 2014 IEEE Conference on* (Oct 2014), pp. 271–272. 1

[SDK14] SPILIOTOPOULOS D., DALIANIS A., KOUROUPETROGLOU G.: Accessibility driven design for policy argumentation modelling. In *Universal Access in Human-Computer Interaction. Design for All and Accessibility Practice*, Stephanidis C., Antona M., (Eds.), vol. 8516 of *Lecture Notes in Computer Science*. Springer, 2014, pp. 101–108. 2

[SJ07] SPÄRCK JONES K.: Automatic summarising: The state of the art. *Information Processing and Management 43*, 6 (Nov. 2007), 1449–1481. 2

[SOR*09] STROBELT H., OELKE D., ROHRDANTZ C., STOFFEL A., KEIM D., DEUSSEN O.: Document cards: A top trumps visualization for documents. *IEEE Transactions on Visualization and Computer Graphics 15*, 6 (2009), 1145–1152. 2

[T*00] TEUFEL S., ET AL.: *Argumentative zoning: Information extraction from scientific text*. PhD thesis, University of Edinburgh, 2000. 2