Models and Tools to Analyse and Predict Discussion Dynamics and Sentiment Towards Policy

Project acronym:	**SENSE4US**
Project full title:	**Data Insights for Policy Makers and Citizens**
Grant agreement no.:	**611242**
Responsible:	**Miriam Fernandez, Open University**
Contributors:	**Hassan Saif, Open University** **Harith Alani, Open University** **Timo Wandhoefer, GESIS** **Steve Taylor, IT Innovation**
Document Reference:	**D5.1**
Dissemination Level:	**PU**
Version:	**1.0**
Date:	**08/10/2014**

Project co-funded by the European Union under the Seventh Framework Programme

History

Version	Date	Modification reason	Modified by
0.1	02/09/2014	Initial draft	Miriam Fernandez
0.2	07/09/2014	Discussion Analysis	Miriam Fernandez
0.3	09/09/2014	Sentiment Analysis	Hassan Saif
0.4	21/09/2014	Quality check	Steve Taylor Timo Wandhoefer Harith Alani
1.0	8/10/2014	Final reviewed deliverable	Miriam Fernandez

Table of contents

Executive summary

Recognising if a policy is well or badly received by the citizens, what elements of the policy are more controversial, and who are the citizens discussing about that policy are key factors to support policy makers in understanding, not only the citizen's opinions about a policy, but also up to which level the social media dialogs represent public opinion and should be used to inform the policy making process. This deliverable describes the research and development conducted by WP5 regarding the analysis of policy discussions and sentiment in social media.

Sentiment is at the core of this WP. During the first year of the project, we investigated the use of contextual and conceptual semantics from Twitter posts for calculating sentiment [Saif et al., 2014b; Saif et al., 2014c, Saif et al. 2014d]. This involved running a comparison of the two types of semantics with respect to their impact on sentiment analysis accuracy. Results showed that using conceptual semantics (gleaned from term co-occurrence) improves sentiment accuracy over several baselines. Results also showed that adding conceptual semantics (entities extracted using AlchemyAPI) enhances this accuracy even further. Accuracy is key in the context of Sense4us since the project aims to provide trustable information in which policy makers can support their decisions. Following this goal we also studied the role of stop words on sentiment analysis [Saif et al., 2014a], showing that best results are achieved when using automatically generated dataset-specific set of stop words. Furthermore, we experimented with a new approach to automatically extend sentiment lexicons to render them more adaptable to domain change on social media, and generated and published a new gold-standard dataset for social media sentiment analysis [Saif et al., 2013].

The second main goal of this WP during this first year has been to analyse policy discussions in social media. Seeding with 42 policy related topics supplied by GESIS, we collected 17,790 tweets from 8,296 users that posted about those topics, and analysed the data to identify who the top contributors are, how many of them are citizens, news agencies, other organisations, etc. We also studied the geographic distribution of the users behind those tweets, to understand where the high concentrations are, and how they relate to policy topics [Fernandez et al., 2014]. Our goal with this study has been to identify the characteristics of those users discussing policy in social media and their main policy topics of interests.

This deliverable summarises the conducted research and the tools developed to track evolution of sentiment and discussions. The developed tools include methods for social media data collection, data pre-processing, discussions analysis, sentiment extraction and sentiment prediction.

List of figures

List of tables

List of abbreviations

<Abbreviation>	**<Explanation>**
PM	Policy Maker
TF	Term Frequency
IDF	Inverse Document Frequency

1 Sentiment Analysis

Making and implementing policy at any level of government is difficult, not due to a lack of information, but due to the difficulty of finding and aggregating the right data out of the sea of information which characterises our modern world. In particular, with the emergence of social networking sites fast and continuous streams of data are being generated by citizens about a variety of topics, including their opinions and arguments about policies.

Social media is currently consider an important source of information from where policy makers can obtain the citizen's opinions and gain insights about the policy discussions emerging in this social platforms [IBM, 2012]. For example Twitter, with over 270 million monthly active users and 500 million tweets sent daily,[1] has now become a goldmine for monitoring citizens' sentiment and opinions. However, summarising and extracting sentiment from these large and continuous streams of data constitutes a difficult and important research problem.

Most current approaches for identifying the sentiment of tweets can be categorised into one of two main groups: supervised approaches [Pak et al. 2010, Barbosa et al., 2010, Kouloumpis et al., 2011], which use a wide range of features and labelled data for training sentiment classifiers, and lexicon-based approaches [Thelwall et al., 2010, O'Connor et al., 2010, Bollen et al. 2011], which make use of pre-built lexicons of words weighted with their sentiment orientations to determine the overall sentiment of a given text. Popularity of lexicon-based approaches is rapidly increasing since they require no training data, and hence they are more suited to a wider range of domains than supervised approaches [Thelwall et al., 2012].

Nevertheless, lexicon-based approaches have two main limitations. Firstly, the number of words in the lexicons is finite, which may constitute a problem when extracting sentiment from very dynamic environments such as Twitter, where new terms, abbreviations and malformed words constantly emerge. Secondly and more importantly, sentiment lexicons tend to assign a fixed sentiment orientation and score to words, irrespective of how these words are used in the text. For example, the word "great" should be negative in the context of a "problem", and positive in the context of a "smile". The sentiment that a word express is not static, but depends on the semantics of the word in the particular context it is being used.

As part of WP5 we have been investigating novel sentiment analysis models that account for contextual semantics in order to enhance the accuracy of existing lexical-based sentiment classification methods. We will summarise these models in the following subsections as well as the conducted experimental evaluation. For concrete details of the theoretical and mathematical formalisation of the models, the reader is referred to the following publications [Saif et al., 2014b, Saif et al., 2014c, Saif et al., 2014d]

1.1 Existing Approaches for Sentiment Analysis

Most existing approaches to social media sentiment analysis focus on classifying the individual posts into subjective (positive or negative) or objective (neutral). They can be categorised as supervised approaches and lexicon-based approaches.

Supervised approaches are based on training classifiers from various combinations of features such as word n-grams [Bifet & Frank, 2010; Pak & Paroubek, 2010], Part-Of-Speech (POS) tags [Agarwal et al., 2011; Barbosa & Feng, 2010], and tweets syntax features (e.g., hashtags, retweets, punctuations, etc.) [Kouloumpis et al., 2011]. These methods can achieve 80%-84%

[1] https://about.twitter.com/company

in accuracy [Saif et al., 2011]. However, training data (i.e., data manually annotated by humans by assigning sentiment scores to texts) is usually expensive to obtain [Liu, 2010] especially for continuously evolving subject domains as in Twitter. Furthermore, classifiers trained on data on a specific domain (e.g., movie reviews) may produce low performance when applied to a different domain (e.g., camera reviews) [Aue & Gamon, 2005], or in the case of Sense4us, discussions around policy topics.

To avoid these problem, lexicon-based methods use the sentiment orientation of opinionated words (e.g., *great*, *sad*, *excellent*) found in a given text to calculate its overall sentiment [O'Connor et al., 2010, Bollen et al., 2011]. Instead of using training data to learn sentiment, lexicon-based methods rely on pre-built dictionaries of words with associated sentiment orientations [Taboada et al., 2011], (e.g., the word *happy* may have a +3 positive sentiment score while the word *sad* may have a -3 negative sentiment score in a given sentiment lexicon). Popular examples of these lexicons are SentiWordNet [Baccianella et al., 2010] or the MPQA subjectivity lexicon [Wilson et al., 2005]. Lexicon-based methods not only provide sentiment polarity (positive/negative), but also strength. For example, SentiStrength [Thelwall et al., 2010; Thelwall et al., 2012] computes the positive sentiment strength in the range from 1 (not positive) to 5 (extremely positive).

One limitation of lexicon-based approaches is that they consider **static sentiment values for terms, regardless of the contexts in which these terms are used and the semantics they convey**. In many cases however, the sentiment of a word is implicitly associated with the semantics of its context [Cambria, 2013]. For example, the word *good*, which is assigned a positive sentiment score in sentiment lexicons, conveys negative sentiment in the sentence "I will leave you for good".

With the aim of addressing these limitations and provide accurate sentiment analysis models that can adequately summarise sentiment towards policy, we present in this deliverable the SentiCircle approach. This approach is based on lexicons, and hence **can be applied to data of different domains**. This is relevant in the context of Sense4us, since policies may involve knowledge around a variety of domains. The proposed SentiCircle approach also **captures the contextual semantics of words** to update their sentiment orientation and strength, improving the accuracy of existing sentiment analysis methods.

1.2 Capturing and Representing Semantics for Sentiment Analysis

In this section we summarise the SentiCircle approach and explain how it captures contextual semantic information. As opposed to traditional lexicon-based approaches SentiCircles does not consider the sentiment of terms static, but depending on the context in which the terms are used. For example, existing lexicon-based sentiment analysis methods may classify the following post as positive *"#Syria. Execution continues with smile! #ISIS :("* because the word *smile* has positive sentiment score in sentiment lexicons. However, the word *smile* conveys a negative sentiment in the context of war and executions; i.e., the sentiment of a word is not static but depends on its contextual semantics.

To capture the words' contextual semantics, we follow the distributional hypothesis that words that occur in similar contexts tend to have similar meanings [Turney et al., 2010; Wittgenstein, 2011]. Therefore, the contextual semantics of a term *m* in our approach is computed from its co-occurrence with other terms. Note that a context is normally given by a collection of posts (in the case of Sense4us, a collection of posts representing social media discussions about a particular policy of interest).

In order to understand the semantics and the sentiment of a target word like "ISIS" (Islamic State in Iraq and Syria), our method relies on the words that co-occur with the target word in

a given collection of posts. These co-occurrences are mathematically represented as a 2d geometric circle. The target word ("ISIS") is at the centre of the circle and each point in the circle represents a context word that co-occurs with "ISIS" in the collection of posts (see Figure 1).

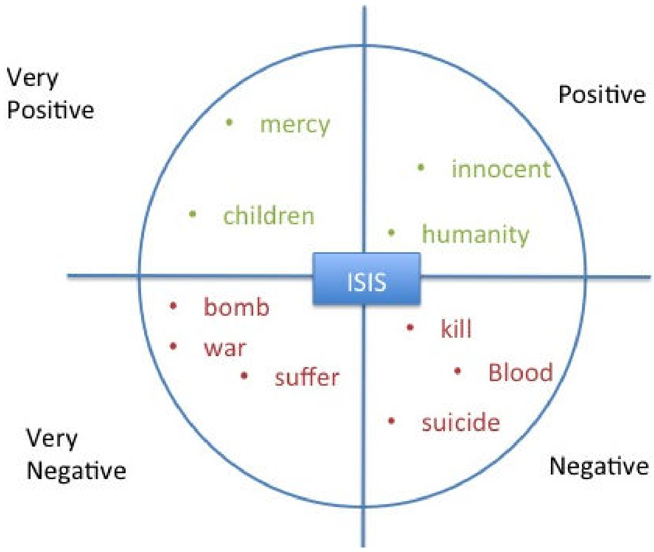

Figure 1: SentiCircle Representation for ISIS (Islamic State in Iraq and Syria)

The position of each context word (innocent, kill, children, etc.) in the circle determines its importance and sentiment toward the target term. The position of each context term within the circle is based on its angle (defined by its sentiment score extracted from a given lexicon), and its radius (defined by its co-occurrence with the target word in a given context).

For example the angle of the context word *kill* its calculated as, θ_{kill} = Prior_sentiment(*kill*) * π, where the prior sentiment is extracted from a given sentiment lexicon, and has a value between -1 and 1. Note that, by calculating the angle in this way, we are considering two independent semicircles. The upper semicircle, where the angle varies from 0 to π, capturing contextual words with positive sentiment, and the lower semicircle, where the angle varies from 0 to – π; capturing contextual words with negative sentiment. If the context term, in this case kill, appears in the vicinity of a negation (words like no, can't, etc.), its prior sentiment score is negated. The negation words are collected from the General Inquirer under the NOTLW category.[2] The degree of correlation between a context word, e.g., *kill*, and a target word, in this case *ISIS*, is computed based on the co-occurrence of these words within a given context, or set of posts.

Terms in the two upper quadrants of the SentiCircle have a positive sentiment (θ between 0 and π), with upper left quadrant representing stronger positive sentiment since it has larger angle values than those in the top right quadrant. Similarly, terms in the two lower quadrants have negative sentiment values (θ between 0 and - π). Points representing context terms the circle have different radii (0 ≤ ri ≤ 1), which reflect how important a context term is to the target term. The larger the radius, i.e., the distance to the origin, the more important the context term is to the target term. For specific mathematical details about the model the reader is referred to the following publication [Saif et al., 2014b]

[2] http://www.wjh.harvard.edu/~inquirer/NotLw.html

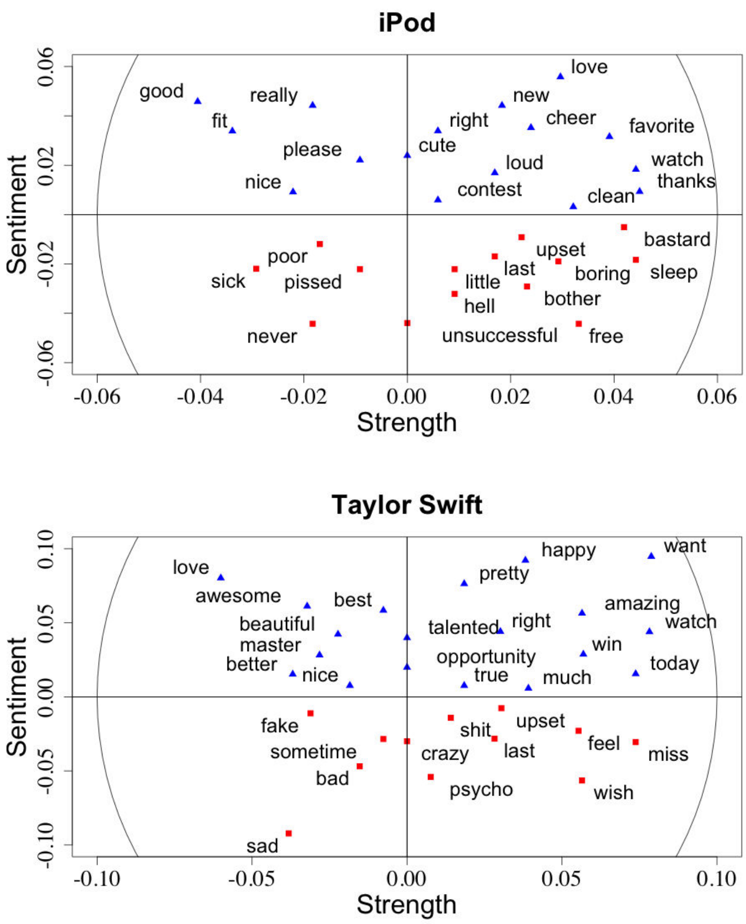

Figure 2: Example SentiCircles for "iPod" and "Taylor Swift". We have removed points near the origin for easy visualisation. Dots in the upper half of the circle (triangles) represent terms bearing a positive sentiment while dots in the lower half (squares) are terms bearing a negative sentiment.

Figure 2 shows two more examples of the SentiCircles generated for the target words "iPod" and "Taylor Swift". Remember that terms (i.e., points) inside each circle are positioned in a way that represents their sentiment scores (angle) and their importance to the target word (radii). For example, "Awesome" in the SentiCircle of "Taylor Swift" has a strong positive sentiment; hence it is positioned in the "*Very Positive*" quadrant. The word "Pretty", in the same circle, also has positive sentiment, but not as strong as "Awesome", hence it is positioned in the "*Positive*" quadrant. We can also notice that there are some words that appear in both circles, but in different positions. For example, the word "Love" has stronger positive sentiment strength with "Taylor Swift" compared to "iPod", although it has a positive sentiment in both circles.

In addition to the contextual semantics captured by SentiCircles, our model also incorporates conceptual semantics. Conceptual semantics refer to the semantic concepts (e.g., "person", "company", "city") that represent entities (e.g., "Steve Jobs", "Vodafone", "London") appearing in posts. The rational behind enriching the SentiCircle with conceptual semantics is that certain entities and concepts tend to have a more consistent correlation to terms of positive or negative sentiment. To extract these entities from posts we make use of semantic annotators, and in particular of AlchemiAPI.[3] We add the concepts into the SentiCircle representation using the Semantic Augmentation method [Saif et al., 2012].

[3] www.alchemyapi.com

1.3 Using Semantics to Measure Sentiment

The SentiCircles representation is used for detecting the positive and negative sentiment expressed on social media. SentiCircles can be used in the following two sentiment analysis tasks:

- **Entity-level Sentiment Analysis**: which aim to detect the sentiment of a given named entity (e.g., "Obama", "David Cameron", "ISIS")
- **Post-level Sentiment Analysis**: which aim to detect the overall sentiment of a given post (e.g., "*#Syria #ISIS. Execution continues with smile! :(*")

To compute sentiment at entity-level we use the geometric median of the SentiCircle (or **Senti-Median**). The Senti-Median is a point within the circle, capturing the overall sentiment and the sentiment strength of the target entity or term (which seats at the centre of the SentiCircle) by using all the previously positioned contextual terms. The geometric median of the SentiCircle (or Senti-Median) is calculated as the point where the Euclidean distances[4] to all the points within the circle is minimum.

To compute sentiment and tweet level we propose three different methods:

- **The Median Method**: The method calculates the sentiment of a post by capturing the SentiCircle of each term within the post and computing the Senti-Median of each of these SentiCircles. The final sentiment score of the tweet is computed as the median of all the previously computed Senti-Medians.
- **The Pivot Method:** This method favours some terms in a post over others, based on the assumption that sentiment is often expressed towards one or more specific targets, which we refer to as "Pivot" terms. For simplicity, we assume that the pivot terms are those having the Part of Speech (POS) tags: {Common Noun, Proper Noun, Pronoun} in a post.
- **The Pivot-Hybrid Method**: This last method is a combination of two previous ones. In the cases where the Pivot-Method fails to find a pivot term (because the post is too short, or it contains many ill-formed words) the Median Method is applied to compute the sentiment of the post.

1.4 Evaluating the SentiCircle approach

As mentioned in Section 1.2 the contextual semantics captured by the SentiCircle representation are based on terms co-occurrence from the context or corpus and an initial set of sentiment weights extracted from a sentiment lexicon. We propose an evaluation set up that uses three different corpora (collections of tweets) and three different generic sentiment lexicons. This enables us to assess the influence of different corpora and lexicons on the performance of the SentiCircle approach.

1.4.1 Datasets

In this section, we present the three Twitter datasets used for the evaluation; OMD, HCR and STS-Gold. We use the OMD [Diakopoulos & Shamma, 2010] and HCR [Speriosu et al., 2011] datasets,[5] to assess the performance of our approach at the post level only since they provide

[4] http://en.wikipedia.org/wiki/Euclidean_distance
[5] https://bitbucket.org/speriosu/updown

human annotations for posts but not for entities (i.e., each tweet is assigned a positive, negative or neutral sentiment label).

Due to the lack of gold-standard datasets for evaluating entity-level sentiment, we have generated an additional dataset (STS-Gold) [Saif et al., 2013]. This dataset contains both; post and entity sentiment ratings and therefore, we use it in this work to assess the performance of SentiCircles at both entity- and post-level. Numbers of positive and negative tweets within the three datasets are summarised in Table 1:

Table 1: Twitter datasets used for the evaluation

Dataset	*Tweets*	*Positive*	*Negative*
OMD	1,081	393	688
HCR	1,354	397	957
STSGold	2,034	632	1,402

Obama-McCain Debate (OMD): This dataset was constructed from 3,238 tweets crawled during the first U.S. presidential TV debate in September 2008 [Diakopoulos & Shamma, 2010]. Sentiment ratings of these tweets were acquired using Amazon Mechanical Turk, where each tweet was rated by one or more voter as either positive, negative, mixed, or other.

Health Care Reform (HCR): The HCR dataset was built by crawling tweets containing the hashtag “#hcr” (health care reform) in March 2010 [Speriosu et al., 2011]. A subset of this corpus was manually annotated with three polarity labels (positive, negative, neutral).

Standford Sentiment Gold Standards (STS-Gold): We constructed this dataset as a subset of the Stanford Twitter Sentiment Corpus (STS). It contains 2,034 tweets (632 positive and 1402 negative) and 58 entities manually annotated by three different human evaluators. The complete description of this dataset is available at [Saif et al., 2013]

1.4.2 Sentiment Lexicons

As describe in Section 1.2, initial sentiments of terms in SentiCircle are extracted from a sentiment lexicon (prior sentiment). We evaluate our approach using three external sentiment lexicons in order to study how the different prior sentiment scores of terms influence the performance of the SentiCircle representation for sentiment analysis. The aim is to investigate the ability of SentiCircles in updating these sentiment lexicon scores based on the contextual semantics extracted from different corpora. We selected three state-of-art lexicons for this study: (i) the SentiWordNet lexicon [Baccianella et al., 2010], (ii) the MPQA subjectivity lexicon [Wilson et al., 2005] and, (iii) Thelwall-Lexicon [Thelawall et al., 2012].

1.4.3 Baselines

We compare the performance of our propose SentiCircle representation when being used for tweet and entity sentiment analysis against the following baselines:

Lexicon Labelling Method: This method uses the MPQA and the SentiWordNet lexicons to extract the sentiment of a given text. If a tweet contains more positive words than negative ones, it is labelled as positive, and vice versa. For entity-level sentiment detection, the sentiment label of an entity is assigned based on the number of positive and negative words that co-occur with the entity in its associated tweets. In our evaluation, we refer to the

method that uses the MPQA lexicon as *MPQA-Method* and to the method that uses the SentiWordNet lexicon as *SentiWordNet-Method*.

SentiStrength: SentiStength [Thelwall et al., 2010, 2012] is a state-of-the-art lexicon-based sentiment detection approach. This method overcomes the common problem of ill-formed language on Twitter and the like, by applying several lexical rules, such as the existence of emoticons, intensifiers, negation and booster words (e.g., absolutely, extremely).[6] For entity-level sentiment detection, the sentiment of an entity is assigned based on the total number of positive, negative tweets in which the entity occurs.

1.4.4 Results

We report the performance of our proposed approaches in comparison with the baselines in both the entity- and tweet-level sentiment detection tasks. For entity-level sentiment detection, we conduct experiments on the STS-Gold dataset, while for tweet-level sentiment detection we use the OMD, HCR and STS-Gold datasets.

Table 2: Entity-level sentiment analysis results

Subjectivity Classification (Subjective vs. Objective)										
Methods	Accuracy	Subjective			Objective			Average		
		P	R	F1	P	R	F1	P	R	F1
MPQA-Method	63.79	67.27	92.50	77.89	0	0	0	33.64	46.25	38.95
SentiWordNet-Method	63.79	67.27	92.50	77.89	0	0	0	33.64	46.25	38.95
SentiStrength	62.07	64.15	91.89	75.56	40.00	9.52	15.38	52.08	50.71	51.38
Senti-Median (SentiWordNet)	81.03	90.91	78.95	84.51	68.00	85.00	75.56	79.45	81.97	80.03
Senti-Median (MPQA)	77.59	90.00	72.97	80.60	64.29	85.71	73.47	77.14	79.34	77.03
Senti-Median (Thelwall-Lexicon)	79.31	84.85	80.00	82.35	72.00	78.26	75.00	78.42	79.13	78.68
Polarity Classification (Positive vs Negative)										
Methods	Accuracy	Positive			Negative			Average		
		P	R	F1	P	R	F1	P	R	F1
MPQA-Method	72.5	80	92.31	85.71	71.43	45.45	55.56	75.71	68.88	70.63
SentiWordNet-Method	77.50	88.00	88.00	88.00	75.00	75.00	75.00	81.50	81.50	81.50
SentiStrength	85.00	95.65	81.48	88.00	70.59	92.31	80.00	83.12	86.89	84.00
Senti-Median (SentiWordNet)	87.50	89.29	92.59	90.91	83.33	76.92	80.00	86.31	84.76	85.45
Senti-Median (MPQA)	85.00	86.21	92.59	89.29	81.82	69.23	75.00	84.01	80.91	82.14
Senti-Median (Thelwall-Lexicon)	82.50	85.71	88.89	87.27	75.00	69.23	72.00	80.36	79.06	79.64

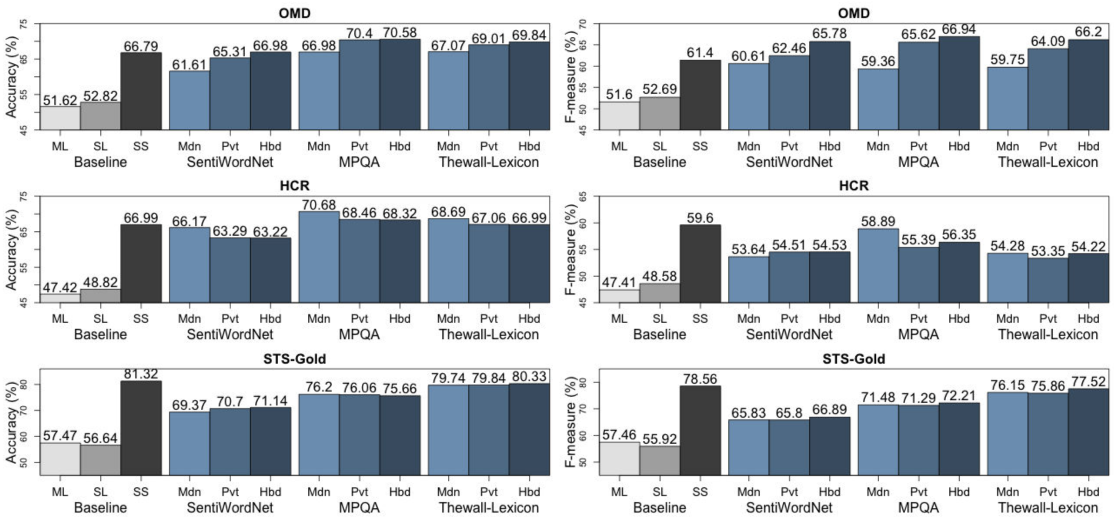

Figure 3: Tweet-level sentiment detection results (Accuracy and F-measure), where ML: MPQA-Method, SL: SentiWordNet-Method, SS: SentiStrength, Mdn: SentiCircle with Median method, Pvt: SentiCircle with Pivot method, Hbd: SentiCircle with Pivot-Hybrid method.

[6] http://sentistrength.wlv.ac.uk/documentation/ SentiStrengthJavaManual.doc

Table 2 and Figure 3 show the potential of using SentiCircles for sentiment detection at the entity and the tweet levels respectively.

Results at entity level show that merely using MPQA or SentiWordNet for sentiment labelling fails to detect any neutral entities. This is expected since words in both lexicons are oriented with positive and negative scores, but not with neutral ones. SentiCircles, on the other hand, was able to amend sentiment scores of words in both lexicons based on contexts - hence achieving a much higher performance in detecting neutral entities than all the baselines.

At the tweet-level, the evaluation was performed on three Twitter datasets and using three different sentiment lexicons. The results showed that our SentiCircle approach outperforms significantly the MPQA-Method and SentiWordNet-Method. Compared to SentiStrength, the results were not as conclusive, since SentiStrength slightly outperformed SentiCircles on the STS-Gold dataset, and also yielded marginally better F-measure for the HCR dataset. This might be due to the different topic distribution in the datasets. STS-Gold dataset contains random tweets, with no particular topic focus, whereas OMD and HCR consist of tweets that discuss specific policy topics, and thus the contextual semantics extracted by SentiCircle are probably more representative in these datasets than in STS-Gold.

In conclusion, our proposed SentiCircle approach, which captures the semantics of words from their context and update their sentiment orientations and strengths accordingly, outperformed other lexicon labelling methods and overtake the state-of-the-art SentiStrength approach in accuracy, with a marginal drop in F-measure. Unlike most other lexicon-based approaches, **SentiCircle was able to update the sentiment strength of many terms dynamically based on their contextual semantics**. This increase in accuracy is very relevant to provide trustable information to Policy Makers about the opinions that citizen's are expressing in social media.

1.5 Analysing and Predicting Sentiment Dynamics and Evolution: Implementation Details

Based on the previously presented research we have developed a set of modules to analyse and predict sentiment around social media conversations. These modules are able to monitor and compute sentiment over time for both, English and German. The key set of provided functionalities is described in Table 3.

Table 3: Functionalities of the Sentiment Analyser

Function Name	*Input*	*Output*	*Description*
computeSentimentForText	String text	sentiment score	Calculates the sentiment for a given post
computeSentimentForTextList	ArrayList<String> textList	ArrayList< sentimentScores >	Calculates the sentiment for a list of posts
computeSentimentForTimePeriod	Date initDate Date finalDate	sentiment score	Calculates the average sentiment for the subset of posts created during the specified time period
computeSentimentFor	String userID	sentiment score	Calculates the average

User			sentiment for the subset of posts created by the specified user
computeSentimentForUserBetweenTimePeriod	String userID Date initDate Date finalDate	sentiment score	Calculates the average sentiment for the subset of posts created by the specified user during the specified time period
ComputeSentimentOverTime	Date initDate Date finalDate Int period	ArrayList< sentimentScores >	Calculates sentiment over time considering the time periods of length period between the intialDate and finalDate
ComputeSentimentForUserOverTime	String userID Date initDate Date finalDate Int period	ArrayList< sentimentScores >	Calculates sentiment over time for the user userID considering the time periods of length period between the intialDate and finalDate
CreateSentimentPredictionModel	Date initDate Date finalDate Int period	SentimentPredictionModel (composed by regression models capturing the evolution of positive, negative and neutral sentiment)	Creates three regression models for sentiment prediction considering the evolution of positive, negative and neutral sentiment between initDate and finalDate at period level of granularity.
PredictSentiment	SentimentPredictionModel sentimentScore	sentimentScore	Given a sentiment prediction model previously generated by considering the evolution of sentiment over time, and the sentimentscores of the last time period it returns the predicted sentiment score

Please note that **the functionalities *computeSentimentOverTime* and *computeSentimentforUserOverTime* allow us to track sentiment dynamics**, for specific conversations around a policy topic, or for a specific user. To compute the over time sentiment we consider a time period (per hour, per day, per week, etc.). Figure 4 shows the evolution of sentiment over time between the 20th of April and the 25th of April for social media conversations around the topic "energy". The chosen period in this case is evolution per day. For each day, our module returns a sentiment score containing the level of positivity, negativity and neutrality of the conversations around energy.

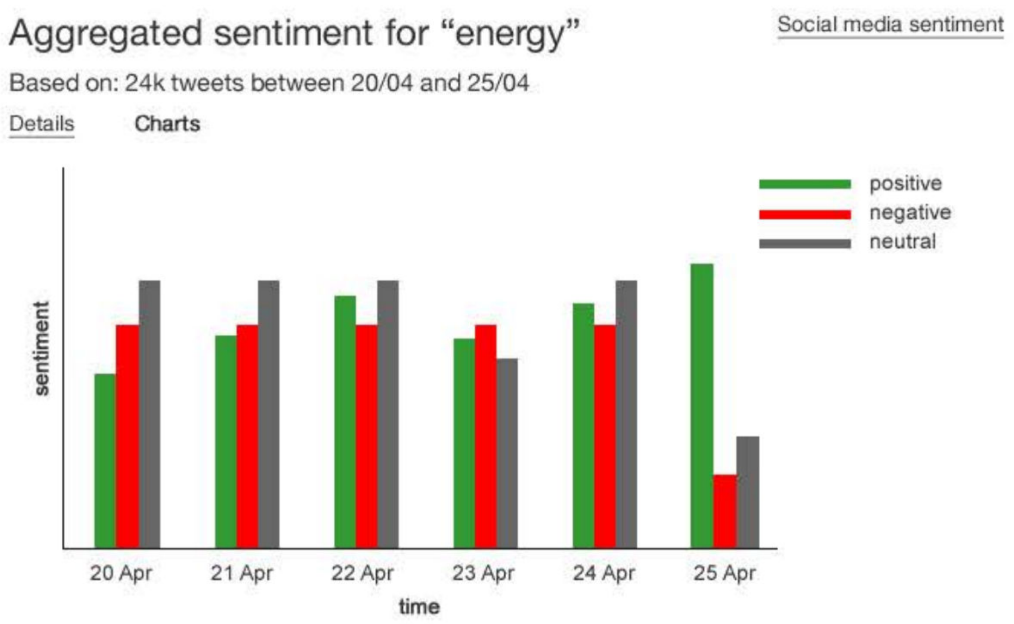

Figure 4: Sentiment over time for the term "energy"

Evolution over time is considered for sentiment prediction. The function **CreateSentimentPredictionModel** generates three linear regression models to capture the evolution of positive, negative and neutral sentiment over time. The time period and granularity level used to generate these three regression models are based on the initDate, finalDate and period parameters input within the function. Once the regression models have been generated, the function **PredictSentiment** is used to predict future sentiment levels.

Table 3 presents the high level functionalities. **All these functionalities are based on the computation of sentiment per entity or per post, which use the sentiment analysis methods previously presented**.

Based on our experiments, reported in the previous section, our core modules to compute sentiment are based on a combination of two approaches: SentiStrength [Thelwall et al., 2012] and our proposed SentiCircles. For those cases where a lexicon adaptation is not needed (e.g., fast sentiment analysis for recently collected posts, or sentiment analysis for posts spamming a variety of domains) our system calls SentiStrength for computing the corresponding sentiment scores. Otherwise, the sentiment computation is performed using our proposed SentiCircle.

A description of the SentiStrength key functionalities is described in (http://sentistrength.wlv.ac.uk/). A detailed architectural description of our proposed SentiCircle is described below:

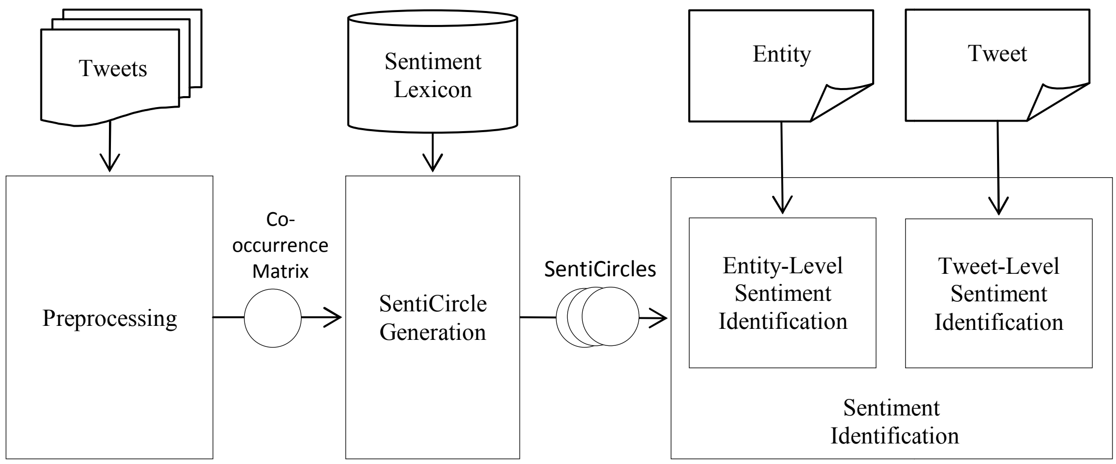

Figure 5: SentiCircle Architecture

1.5.1.1 SentiCircle Pre-processing Module

This module takes as input a collection of raw posts and applies several syntactical pre-processing procedures on the terms found in the collection as described in the table below. The output is a two dimensional array (co-occurrence matrix) that represents the occurrence frequency between each two unique terms in the collection of posts. Key functions used in this module are described below.

Table 4: Functionalities of the SentiCircle pre-processing module

Function Name	*Input*	*Output*	*Description*
LoadPosts()	CSV File DB Connection	List of tweets messages	Loads posts into memory
TonknisePosts()	Post	List of terms	Tokenises the post into individual terms and assign a part-of-speech tag to each term
ProcessSlang()	Term	Term	Reverts slangs to their original English form (e.g., luv -> love). This functionality is not available for German
ProcessNegation()	List of terms	List of terms	Removes negation terms from tweet and mark those in the window of negation (e.g., not happy -> $happy$)
CalcFrequency()	List of terms	Co-occurrence Array	Finds the co-occurrences between each two terms in a

			given list.
ExportCoMatrix()	Co-occurrence Matrix	JSON File	Saves the co-occurrence matrix of terms into JSON-formatted file.

1.5.1.2 SentiCircle Generation Module

This module takes as input a co-occurrence matrix of terms and generates a SentiCircle for each unique term. Key functions used in this module are described in the table below.

Table 5: Functionalities of the SentiCircle Generation module

Function Name	*Input*	*Output*	*Description*
LoadCoMatrixFile()	JSON File	Co-occurrence matrix	Loads the co-occurrence matrix into memory
CalcTermCorreltion()	Co-occurrence matrix	Correlation Matrix	Calculates the term degree of correction (T-Doc) between each two terms in the co-occurrence matrix
AssignPriorSentiment	- Co-occurrence matrix - Sentiment Lexicon	Prior_Sentiment Matrix	Assigns terms in the co-occurrence matrix with their prior sentiment orientation using a given sentiment lexicon
GenerateSentiCircle()	- Term - Correlation Matrix - Prior_Sentiment Matrix	SentiCircle (XY-Matrix)	Generates the SentiCircle of a given term
ExportSentiCircles()	List of SentiCircles	JSON File	Saves the SentiCircles of terms into JSON-formatted file

1.5.1.3 Sentiment Identification Module

This module is responsible for the sentiment identification process in our system. It consists of two sub sentiment identification modules: tweet-level sentiment identification module and

entity-level sentiment identification module. In the following we describe each of these modules separately along with the functions used within them.

3.1 Entity-level Sentiment Identification Module

This module takes as input a named-entity (e.g., Obama, iPhone, etc.) and calculates its sentiment orientation based on its corresponding SentiCircle.

Table 6: Functionalities of the SentiCircle Entity sentiment identification module

Function Name	*Input*	*Output*	*Description*
LoadSentiCircle()	- Entity - SentiCirlce File	The SentiCircle of the given entity	Retrieves the SentiCircle of the entity from the given SentiCircle JSON File
CalcEntitySentiment()	- Entity - SentiCircle	Sentiment Orientation	Calculates the sentiment orientation of the queried entity (i.e., positive, negative, neutral)
ExportEntitySentiment()	Sentiment Orientation	File	Saves the sentiment orientation of the entity to a CSV-formatted file.
CalcAllEntities()	- Entity File - SentiCircle File	File of entities with their sentiment orientation	Calculates the sentiment of a collection of entities (imported from a given CSV file) and export the results to a CSV file.

3.2 Post-level Sentiment Identification Module

This module takes as input a social media post and calculates its overall sentiment orientation based the SentiCircles of the terms within the tweet.

Table 7: Functionalities of the SentiCircle Tweet sentiment identification module

Function Name	*Input*	*Output*	*Description*
LoadTermsSentiCircle()	- Post - SentiCirlce File	List of SentiCircles	Retrieves the SentiCircle of each term in the post using the provided SentiCircle

			file.
CalcPostSentimentMedianMethod()	List of SentiCircles	Sentiment Orientation	Uses *the Median Method* to calculate the sentiment orientation of the given post.
CalcPostSentimentPivotMethod()	List of SentiCircles	Sentiment Orientation	Uses *the Pivot Method* to calculate the sentiment orientation of the given post.
CalcPostSentimentHybridMethod()	List of SentiCircles	Sentiment Orientation	Uses *the Pivot-Hybrid Method* to calculate the sentiment orientation of the given post.
ExportPostSentiment()	Sentiment Orientation	File	Saves the sentiment orientation of the post to a CSV-formatted file.
CalcSentimentAllPosts()	- List of Posts - File of SentiCircles	Posts with their sentiment orientation	Calculates the sentiment of a collection of posts and export the results to a CSV file.

2 Analysis of Policy Discussions

Understanding who are the users discussing policy in social media and how policy topics are debated could help Policy Makers assessing how their views and opinions should be weighted and considered to inform policy making.

In this first year of the project, WP5 has focused on developing methods to monitor and analyse policy discussions in social media. The software modules encapsulating these methods are described in section **Error! Reference source not found.**. To exemplify how these modules can be exploited to obtain insights into policy discussions this deliverable presents an analysis of Twitter discussions around 42 different policy related topics. Sixteen Policy Makers, who are members of different political institutions in Germany, selected these topics. This research has been published in [39].

The rest of the section is structured as follows: Subsection 2.1 describes state of the art intended to characterize users and policy discussions in social media. Subsection 2.2 describes the data collection process and the final dataset used for this study. Subsection 2.3 explains the analyses performed over the data and the extracted insights. Subsection 2.4 presents our conclusions. The developed software modules are presented in section 2.5.

2.1 Characterising users and policy debate in social media

Statistics about e-participation are studied regularly. These statistics are computed globally [E-Gov Survey, 2012], at EU level,[7] and for individual countries.[8] While such studies highlight the benefits of e-participation platforms, they also indicate that participation via specific online government services is generally low. The last report of the United Nations [E-Gov Survey, 2012] points out that within the 27 EU countries, only 32% of individuals aged 16 to 74 have used the Internet for interacting with public authorities. These reports also emphasize the need of using social media to improve public services, reduce costs and increase transparency.

Several studies have been conducted that investigate the characteristics of users participating in social media [Madden, 2010; Poblete et al., 2011; Honigman 2012; Beevolve 2012]. Regarding Twitter, the SNS selected for this study, Beevolve concluded that: (i) there is 6% more of women than men in Twitter, (ii) 75% of users fall between 15 to 25 years of age and, (iii) the average Twitter user follows 102 users, is followed by 50 users and post 307 times during her Twitter life. While these works extract important insights and demographics *they aim to characterise the average social media user, and not those particular users engaged with policy topics*. A deep review of the use of social media for eGoverment can be found in [Magro, 2012]. While this review includes a historic overview of the use of social media for eGoverment, none of the works referenced in this study investigates the characteristics of those users participating in policy discussions.

Some works have studied the dynamics of policy discussions in social media but in the context of concrete political events, such as elections [Adamic & Glance, 2005; Tumasjan et al., 2010; Conover et al., 2011] or revolutions [Aday, 2010; Bhuiyan, 2011]. While these works focus on analysing debates around a particular event our goal is to provide an overview of how policy topics are discussed; which topics are more interesting for the general public and what is the level of positive and negative sentiment expressed about those topics.

---

[7] http://epp.eurostat.ec.europa.eu/statistics_explained/index.php/E-government_statistics#Publications

[8] http://www.hansardsociety.org.uk/wp-content/uploads/2012/10/Digital-Citizens-and-Democratic-Participation-2010.pdf

2.2 Data collection and processing

Policy makers need to know citizens' reactions to specific policies and policy topics, especially when formulating the policy – citizens' reactions provide valuable feedback about society's reaction to the policy if it were enacted. Following this premise, we contacted 16 Policy Makers, all members of different political institutions in Germany: the German Bundestag, the State Parliament North Rhine-Westphalia, the state Chancellery of the Saarland and the cities Cologne and Kempten. Each of these Policy Makers indicated four or five topics that were of particular interest to them, generating a total of 76 policy-related topics including issues such as nuclear power, unemployment or immigration. We filtered 34 out of the 76 initial topics, remaining with a total of 42. The purpose of the filtering process was to discard very generic topics such as "women", which led to the collection of Twitter discussions not related to policy topics. This filtering process allowed us to reduce the noise of the collected data sample. Table 8 shows the filtered list of topics. Please note that these topics were selected by German Policy Makers and therefore all of them are expressed in German. The English translation is provided for convenience.

Table 8: Filtered topics and their corresponding English Translation

Topics	*English Translation*	*Topics*	*English Translation*
Betreuungsgeld	Care Benefit	Nichtraucherschutz	Non Smoking Protection
Bildungspolitik	Education Policy	NPD-Verbot	NPD Ban
Bürgerrechte	Civil Rights	Open Government	Open Government
Castorbehälter	Castor Containers	Parteispenden	Political Donations
Datenschutz	Privacy Policy	Praxisgebühr	Practice Fee
Energiepolitik	Energy Policy	Rauchverbot	Smoking Ban
Europapolitik	European Policy	Rechtsextremismus	Right-wing
Finanzpolitik	Fiscal Policy	Schuldenbremse	Debt Brake
Fracking	Fracking	Schulreform G8	School Reform G8
Frauenquote	Women's Quota	Solidarpakt West	Solidarity Pact West
Generationengerechtigkeit	Intergenerational Equity	Sozialpolitik	Social Policy
Gentechnik	Genetic Engineering	Sozialticket	Social Ticket
Gleichstellung	Equality	Studiengebühren	Tuition
Harz4	Fourth law to reform the rendition of services on the job market	Tempolimit	Speed Limit
Innenpolitik	Domestic Policy	Verbraucherpolitik	Consumer Policy
Kohlekraftwerk Datteln	Coal Power Plant Dates	Verkehrspolitik	Transport Policy
Kommunale Grundversorgung	Municipal Primary Care	Verteidigungspolitik	Defence Policy
Linksextremismus	Left-wing Extremism	Umweltpolitik	Environmental Policy
Migranten	Migrants	Urheberrecht	Copyright
Mindestlohn	Minimum Wage	Volksbegehren	Referendum
Netzpolitik	Network Policy	Europäische	European

2.2.1 Obtaining policy related tweets

To investigate the characteristics of users discussing policy topics in social media we monitored a sample of the Twitter population. We collected users and posts via the Twitter search API[9] using as queries the topics described in Table 8. We restricted the collection to the German language to avoid gathering noisy information. The sample was collected during a week, starting on 4th of January 2014 and finishing on 12th of January 2014. The collected dataset consists of 17,790 posts from 8,296 different users. For both users and posts, we extracted the set of features provided by Twitter and computed an additional set of features to conduct the analyses presented in this deliverable. The complete set of features is listed below:

User features

To analyse the characteristics of each particular user and his role in the conversations around policy topics we extracted the following features:

- *Number of posts*: number of posts that the user *u* has written since his registration on Twitter
- *Post rate*: Number of post per day created by the user *u* since his registration on Twitter
- *Number of policy posts*: number of posts generated by the user *u* in our sample dataset
- *Initiations*: number of conversations that the user *u* has initiated in our sample dataset
- *Contributions*: number of conversations in which the user *u* has participated (reply) in our sample dataset
- *Followers*: number of users who follow the user *u* (a high number of followers indicates high popularity)
- *Friends*: number of users that the user *u* follows (a high number of friends indicates high engagement)
- *Location*: location that the user *u* specifies in his Twitter profile
- *Description*: description that the user *u* specifies about himself in his Twitter profile

Note that demographic information such as age, or gender is not available via the Twitter API.

Content features

To analyse the characteristics of Twitter content around policy topics we extracted the following features:

- *Sentiment*: sentiment polarity and strength of the post *p*
- *Mentions*: the users that are mentioned within the tweets (mentions are identified by the symbol @)
- *Hashtags*: the topics that are explicitly mentioned within the tweets (hashtags are identified by the symbol #)

2.3 Data Analysis

The following section presents the analyses performed over the collected data. The first analysis studies the characteristics of Twitter users discussing policy related topics. The second analysis investigates the debates around policy topics including topic popularity and users’ sentiment in relation to these topics.

[9] https://dev.twitter.com/docs/api/1/get/search

2.3.1 Characterising users contributing to policy discussions

The purpose of this analysis is to characterise those users discussing policy topics in Twitter. Figure 6 shows the distribution of users per number of posts, which presents a long-tail pattern. Users appearing in the head section of this distribution are responsible of 36% of the generated content of our data sample. We will refer from now on to this part of the population as “top contributors” for the rest of our analysis.

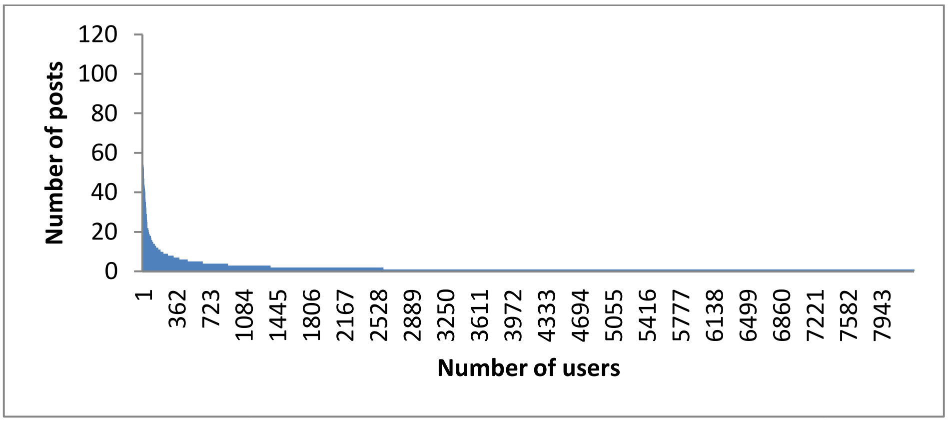

Figure 6: Distribution of users per number of posts

The average top contributor has 4,279 followers, 1,028 friends and has posted 33,134 times during his life in Twitter. Figure 7 displays the tag cloud of the top contributors’ names. Among these top contributors we identify multiple organisations and news agencies such as Demokratie Report, Anonymous Germany, DTN Germany, Svejk News, Netz4ktivisten, TimesDailyNews, Voice Dialogue and others. We have manually assessed the user accounts that belong to the group of top-contributors and 73.4% of the top contributors do not represent individual citizens but news agencies and other organisations. We can therefore conclude that ***policy discussions are mainly contributed by a small subset of active Twitter users that do not represent individual citizens but news agencies and other organisations.***

Figure 7:Names of the top contributors

The long-tail of the distribution (the remaining users) presents an average of 1,365 followers, 630 friends and 9,578 posts during their Twitter life. These numbers are still higher than the ones reported for the average Twitter user [Beevolve 2012], which follows 102 users, is followed by 50 users and posts 307 times during her Twitter life. These results provide an

indication that *the users contributing to policy topics are more active, popular and engaged than the average Twitter user*.

(a) Distribution of eParticipation projects	(b) Distribution of Twitter users

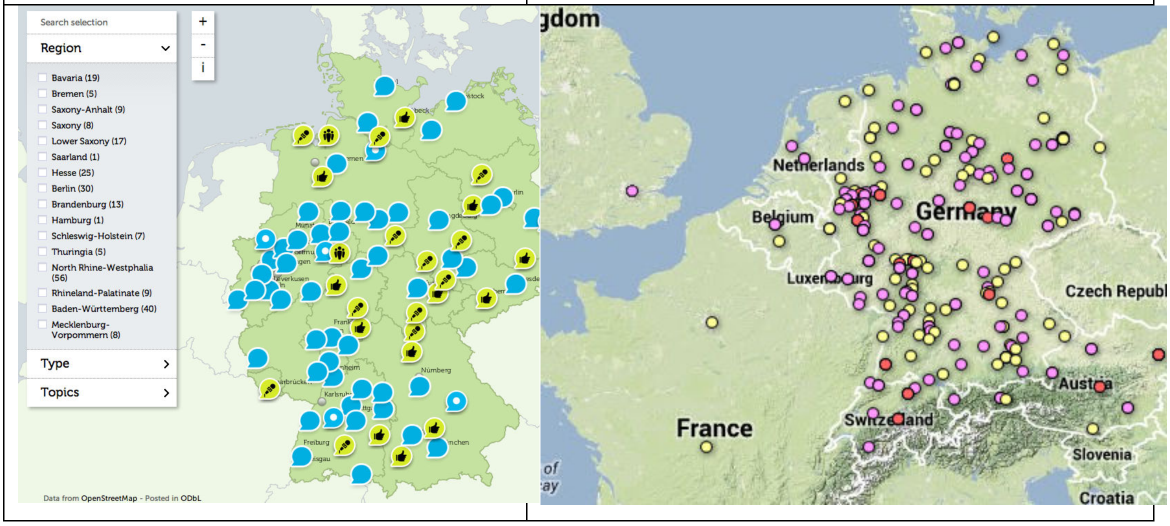

Figure 8: (a) Distribution of eParticipation projects in Germany (http://www.politik.de/politik-de/projekte_entdecken/beteiligungskarte) (b) Distribution of Twiter users: yellow are locations with less than 10 users, pink are locations with 10 to 50 users, red are locations with more than 50 users

In addition to this analysis we have also investigated how users are geographically distributed. For this purpose we extracted the locations specified within the profiles of the collected Twitter users and geocoded them (extracted the latitude and longitude coordinates) by making use of the Google Maps API[10]. Figure 8(b) displays the geographic distribution of users: in yellow those locations with less than 10 users, in pink those locations with 10 to 50 users, and in red, those locations with more than 50 users. As expected, the **higher concentration of users occurs in constituencies of high population density** such as Berlin, Hamburg, Munich and Koln in Germany, Vienna in Austria, and Zurich in Switzerland. To investigate whether these locations are similar to the ones from which citizens engage in eParticipation platforms we compared this map with the distribution of eParticipation projects in Germany[11], Figure 8(a). Since we could not find concrete statistics about the geographical distribution of users engaged in eParticipation platforms in Germany, we made the assumption that the regions with higher number of eParticipation initiatives are also those ones where more users are engaged. Note that, while user statistics of eParticipation are available for Germany at country level[12] [E-Gov Survey, 2012] we haven't found any document reporting similar statistics at regional level. To map the locations of Twitter users within the 16 regions specified by politik.de[13] (Babaria, North Rhine Westphalia, Baden-Wurttemberg, etc.) we have made use of the region bounds provided by the Google API. The Pearson correlation coefficient between the number of Twitter users and the number of eParticipation projects in each region is 0.817, where 1 indicates total positive correlation and 0 is no correlation. This indicates that *users engaged in social media conversations around*

[10] https://developers.google.com/maps/

[11] https://www.politik.de

[12] http://www.bertelsmann-stiftung.de/cps/rde/xbcr/bst/xcms_bst_dms_31401__2.pdf

[13] https://www.politik.de

policy topics tend to be geographically concentrated in the same regions than users engaged in eParticipation platforms.

2.3.2 Characterising topics and sentiment of policy discussions

The purpose of this particular analysis is to understand how are policy topics represented in the policy discussions happening in social media and what is the overall sentiment about these topics. To extract the representativeness of policy topics (in terms of quantity of posts and users involved in the discussions) we first obtain the subsets of posts collected for each of the 42 topics and then identify the creators of these posts. To obtain the representativeness of positive and negative sentiment for each particular topic we compute the sentiment for each individual post associated to the topic and then extract the authors of those posts. This can give us an overview of how the opinions in favour and against policy topics are represented in the discussions. The results of this analysis are shown in Table 9.

Table 9: Representativeness of topics. For each topic the table includes: (1) its English translation, (2) the total number of posts about the topic, (3) the total number of users contributing to the topic, (4) the number of positive posts about the topic, (5) the number users contributing positively to the topic, (6) the number of negative posts about the topic, (7) the number of users contributing negatively to the topic, (8) the number of neutral posts about the topic and (9) the number of users contributing neutrally to the topic.

Topic	*Posts*	*Users*	*+Posts*	*+Users*	*-Post*	*-Users*	*NPost*	*NUsers*
privacy	3439	2130	491	404	404	361	2544	1629
network policy	3250	1615	515	392	323	262	2412	1271
minimum wage	2598	1558	683	578	285	240	1630	979
copyright	1297	954	221	183	68	62	1008	788
fracking	1079	688	236	191	194	174	649	431
domestic policy	910	478	175	146	108	79	627	323
genetic Engineering	808	454	72	57	**82**	51	654	410
Harz4	632	351	100	78	34	30	498	281
migrants	601	494	139	130	**143**	127	319	270
equality	536	421	164	145	33	30	339	280
female ratio	416	370	217	203	23	22	176	156
right wing	306	221	85	84	28	24	193	127
referendum	300	223	30	26	**108**	98	162	129
left wing extremism	245	199	50	49	26	25	169	142
education and training policy	235	213	94	88	39	38	102	98
energy policy	185	146	35	30	33	28	117	98
european policy	139	128	22	22	**25**	24	92	84
party donate	110	100	4	4	**7**	7	99	92
social policy	107	77	25	21	13	11	69	50
speed limit	75	66	5	5	**13**	10	57	52
financial policy	74	68	20	20	4	3	50	47
no smoking	70	66	16	16	4	4	50	48

care money	66	61	10	10	1	1	55	50
transport policy	61	54	15	15	3	3	43	37
generational justice	55	54	13	13	2	2	40	40
debt brake	51	45	6	6	6	6	39	34
environmental policy	43	39	8	8	3	3	32	30
npd ban	36	24	3	3	**19**	7	14	14
non smoking shelter	35	35	4	4	**7**	7	24	24
socially ticket	28	21	1	1	**4**	4	23	16

Our analysis shows that 30 out of the 42 collected topics were discussed during the monitored week. The posts distribution per topic is displayed in Figure 9. As we can see in this figure, ***few topics are extensively discussed*** during the analysed period, such as privacy, network policy, minimum wage, or copyright, ***while the majority of topics are underrepresented***.

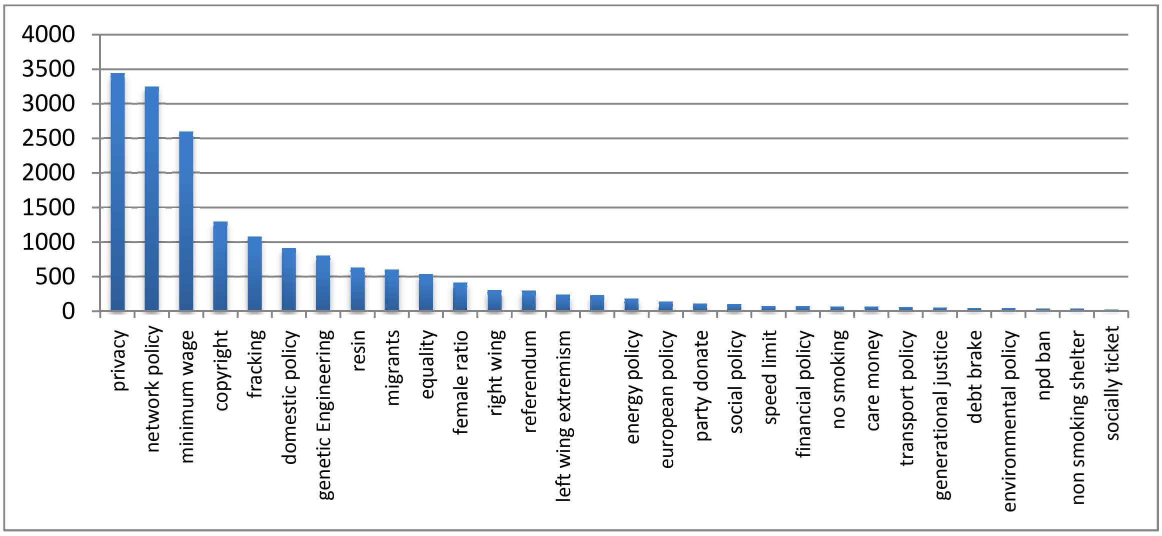

Figure 9: Post distribution per topic

This topic distribution is also reflected in the hashtags used within the Twitter conversations. Hashtags are metadata tags that Twitter users include in their posts to explicitly indicate the topics under discussion. As we can see in Figure 10 the most popular hashtags of our dataset include privacy (Datenschutz), minimum wage (Mindestlohn), copyright (Urheberrecht), fracking or genetic engineering (Gentechnik), which are among the most popular topics in terms of frequency of associated posts.

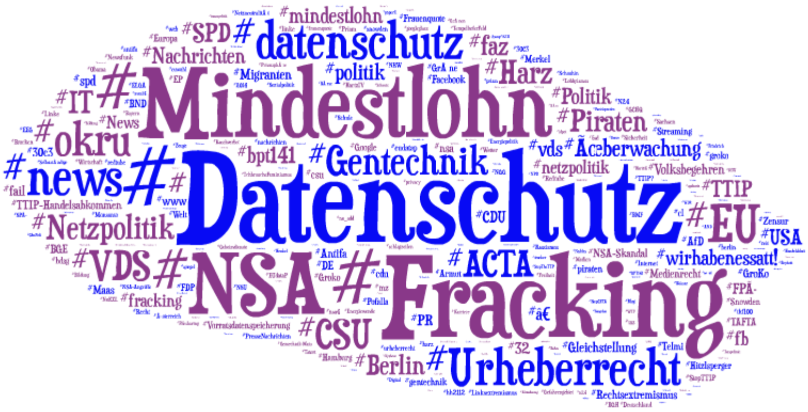

Figure 10: Hashtags extracted from the collected dataset

Regarding the sentiment around the collected topics, we have highlighted in bold (Table 9) those topics for which the number of negative contributions is higher than the number of positive contributions. Among these topics we find genetic engineering, immigration, the possibility of a referendum, donations to political parties or the speed limit. The rest of the monitored topics present a slightly higher number of positive than negative tweets. Note that some topics present a high volume of posts around positive and negative opinions, i.e., they are notably controversial. Among these topics we can highlight privacy, fracking, or domestic policy.

As a measure of user engagement in conversations around policy topics we have analysed the reply chain of the collected conversations. 45% of the collected posts in our dataset are replies to previously initiated discussions. Contrasting this result with earlier studies based on different collected Twitter datasets,[14] where a maximum of 23% of posts where replies, this percentage of engagements in discussions is comparatively high, i.e., *users that engage in policy discussions in Twitter more actively than in other topics.*

2.4 Conclusions of the performed analysis

Understanding who are the users discussing policy in social media and how policy topics are debated could help Policy Makers assessing how their views and opinions should be weighted and considered to inform policy making.

WP5 has focused on this first year of the project on developing methods that can monitor and analyse policy discussions. As an example of the application of these methods we have presented in this deliverable a research study in which Twitter discussions around 42 policy topics (including 8,296 users and 17,790 posts) have been analysed. These topics, provided by GESIS, were selected by sixteen Policy Makers from different political institutions in Germany.

We analysed the different types of user groups discussing policy topics as well as their geographical distribution. Our results show that a small percentage of users (top contributors) are responsible for most of the generated discussions (around 6% of users are responsible of 36% of the conversations). 73.4% of the top contributors are not individual citizens but news agencies and other organisations. Our results also show that the Twitter user discussing policy topics is more active, popular and engaged than the average Twitter user. Regarding the geographical distribution of these users we have observed that: (i) they tend to be concentrated in locations with high population density and, (ii) they tend to be concentrated in the same regions than users engaged in eParticipation platforms.

[14] http://www.sysomos.com/insidetwitter/engagement/

We have also analysed the popularity and sentiment of the different conversations around policy topics. Our results indicate that a small subset of topics is extensively discussed (privacy, network policy, minimum wage, copyright, etc.) while the volume of conversations is relatively low for the rest of the topics. Regarding the analysed sentiment, the topics accumulating a higher percentage of negative comments include: genetic engineering, immigrants or the possibility of a referendum. While most of the analysed topics present a higher number of positive than negative comments, some of these topics are particularly controversial.

Note that the dataset used during this research has been made available to the research community (ECSMDataset 2014)

2.5 Monitoring Policy Discussion Dynamics: Implementation Details

In this section we describe the main software components implemented to monitor and analyse policy discussions. As we can see in Figure 11 our architecture is based on three main components:

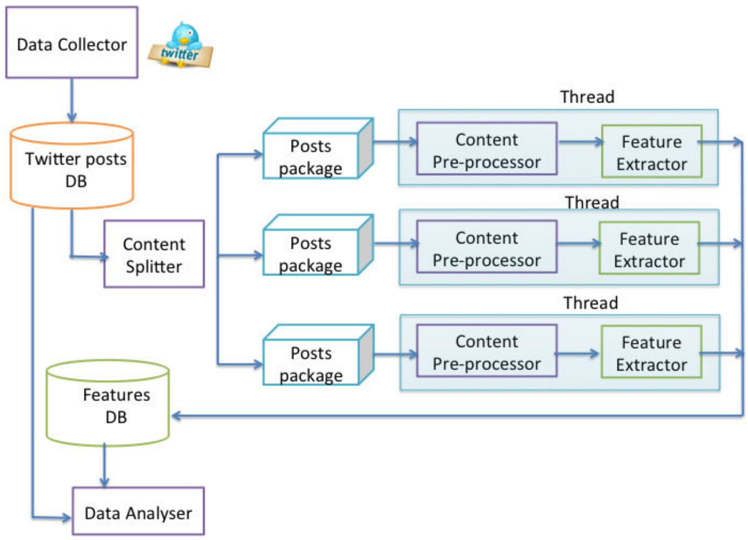

Figure 11: Software Architecture

The Data Collector: This component connects to the Twitter API and downloads data around specific keywords. In the context of Sense4us, these keywords can be manually input by the user, or automatically extracted from a policy document. This module is implemented in Java and supports connections to the Twitter search API as well as to the Twitter Stream API. Key functionalities of this module are described in the table below.

Table 10: Key functionalities for the Data Collector Module

Function Name	*Input*	*Output*	*Description*
GetDataForTermList	List of terms DB Name and connection details (properties file) Limit Date	Relational DB filled with the JSON files returned by the Twitter API	This module receives a list of terms, collects Twitter data for those terms and stores information in a DB.

GetDataForUserList	List of Twitter users DB Name and connection details (properties file) Limit Date	Relational DB filled with the JSON files returned by the Twitter API	This module receives a list of Twitter users, collects posts from their stream, and stores information in a DB.

These two functionalities allow us to collect information for policy discussions, and users over time. If no LimitDate is specified in the Data Collector, it keeps gathering information until manually stopped.

The Feature Extractor: This component extracts, for each post and user a set of features that can be later on used for analysing data. In particular, for each user and post these modules extract the following set of features:

User Features

- *In-degree*: This feature measures the number of incoming connections to the user.
- *Out-degree*: This feature measures the number of outgoing connections from the user.
- *Post Count*: Measures the number of posts that the user has made over her life in the system.
- *Policy Post Count*: Measures the number of posts that the user has made within a particular dataset (normally representing a policy topic)
- *User Age*: Measures the length of time that the user has been a member of the Twitter.
- *Post Rate*: Measures the number of posts made by the user per day.
- *Geolocation*: uses the Google Maps API[15] to extract geocoordinates from the location specified by the user (generally using a String), if available.

Content Features

- *Post length*: Number of terms in the post.
- Complexity: Cumulative entropy of terms within the posts to gauge the concentration of language and its dispersion across different terms. Let n be the number of unique terms within the post p and fi the frequency of the term t within p. Therefore, complexity is given by:

$$complexity(p) = \frac{1}{n}\sum_{i=1}^{n} f_i(\log_n - \log f_i)$$

- *Readability*: This feature gauges how hard the post is to parse by humans. To measure readability we use the Gunning Fox Index[16] using average sentence length (ASL) and the percentage of complex words (PCW).

$$0.4 * (ASL + PCW)$$

- *Referral Count*: number of hyperlinks (URLs) present in the posts.

[15] https://developers.google.com/maps/

[16] http://en.wikipedia.org/wiki/Gunning_fog_index

- *Mentions*: number of mentions to other users within the posts.
- *Informativeness*: The novelty of the post's terms with respect to the other posts. We derive this measure using the Term Frequency-Inverse Document Frequency (TF-IDF) measure, which is commonly used in Information Retrieval:

$$\sum_{t \in p} tf_{t,p} \times idf_t$$

- *Polarity*: Average polarity (sentiment) of the post.
- *Time of the day*: Time when the tweet was posted (e.g., 20:00)

For optimization purposes the Feature Extractor component has been multithreaded. The component divides the set of posts into packages, pre-process them and extracts the corresponding features. Once the features are extracted they are stored in database for further analysis. The component is also implemented in Java and depends on the modules described in Section 1 for the extraction of the sentiment feature. The key functionalities of this module are described in the following table.

Table 11: Key functionalities for the Feature Extractor Module

Function Name	*Input*	*Output*	*Description*
Extract User Features	DB of Twitter data	DB filled with user features	This function process the Twitter data collected and extracts for each user a set of features
Extract Content Features	DB of Twitter data	DB filled with content features	This function process the Twitter data collected and extracts for each post a set of features

Note that this component has been designed to facilitate the incorporation of new content and user features, if needed.

The Data Analyser: This component is the one in charge of providing the different analysis as the ones presented in this section. It is implemented using Java, SQL and R.

Table 12: Key functionalities for the Data Analyser Module

Function Name	*Input*	*Output*	*Description*
AnalyseActivityEvolution	Date initDate Date finalDate Int period	Volume of posts grouped by time period	Creates a histogram that describes the evolution of posts over time

			considering a time period between initDate and FinalDate
AnalyseDemographics	DB of Twitter data DB of Features Boolean (posts/users)	Amount of users/or posts per location	This functionality groups users or posts considering their demographic locations.
AnalyseSentimentOverTime	DB of Twitter data DB of Features **Selected topic/ Selected user** Date initDate Date finalDate Int period	Time evolution of sentiment for the selected policy topic	This functionality computes sentiment scores over time for the selected topic or user. For this purpose it considers slots of size period between initDate and finalDate
AnalyseTopicsOverTime	DB of Twitter data DB of Features **Topic list** Date initDate Date finalDate Int period	Time evolution of policy topics	This functionality computes topic frequency per time period between the specified dates

Note that these functionalities allow us to track the volume of discussions over time, as well as their evolution in terms of topics and sentiment. The output of these functionalities may have different visualisations. We are still collaborating with the technical and use case partners of the Sense4us project to decide on the optimal visualisations for these functionalities. For example, Figure 12 shows a proposal to visualise evolution of policy discussions around three different policy topics in Twitter (climate change, nuclear energy and solar energy). The graph measures the evolution of each topic in terms of the amount of associated posts. The time period selected is “month”.

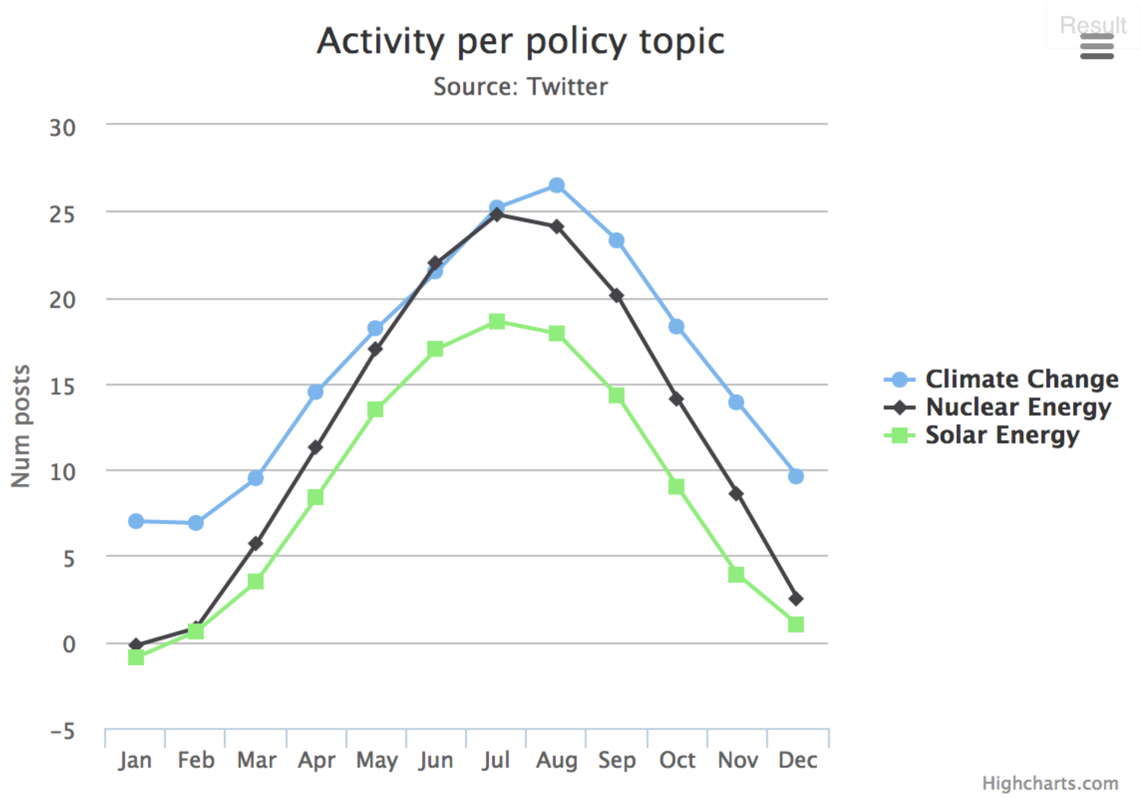

Figure 12: Evolution of policy conversations around three different policy topics: Climate change, Nuclear Energy and Solar Energy

3 Conclusions

This deliverable summarises the work of WP5 conducted during the first year of the project. The goal of this WP has been to develop tools to monitor and analyse policy discussions in social media as well as the citizens' sentiment towards policies based on these discussions.

The developed tools include methods to: (i) monitor and collect social media data, (ii) pre-process this data to extract relevant content and user features, (iii) analyse this data to extract insights from policy discussions (iv) compute the sentiment of the monitored discussions, (v) track the evolution of discussion dynamics and, (vi) predict sentiment evolution.

The design and development of these tools have been driven by the conducted research. During the first year of the project, WP5 investigated the use of contextual and conceptual semantics for calculating sentiment as well as the use of user and content features to track discussion dynamics around a variety of topics.

Our results show an increase in accuracy when using semantic information to compute sentiment. Our results also show relevant insights about policy discussions in social media. After tracking and analysing 42 policy related topics supplied by GESIS our analyses revealed that a small percentage of users (top contributors) are responsible for most of the generated and that the top contributors are not individual citizens but mostly news agencies and other organisations. Our analysis also revealed that a small subset of topics is extensively discussed while the majority go relatively unnoticed.

4 References

[1] Adamic, L., and Glance, N. (2005) The political blogosphere and the 2004 U.S. election: Divided they blog. 3rd Intl. Workshop on Link Discovery (LinkKDD)

[2] Aday, S. (2010) Blogs and bullets: New media in contentious politics. United States Institute of Peace. http://www.usip.org/files/resources/pw65.pdf

[3] Beevolve (2012) An Exhaustive Study of Twitter Users Across the World. http://www.beevolve.com/twitter-statistics

[4] Bhuiyan, S. (2011) Social media and its effectiveness in the political reform movement in Egypt. 2011. Middle east media educator 1, 1, 14–20

[5] Conover, M., Ratkiewicz, J., Francisco, M., Gonçalves, B., Menczer, F., & Flammini, A. (2011). Political polarization on twitter. In ICWSM

[6] E-Gov Survey (2012) http://www.unpan.org/dpadm/egovernment/unegovernmentsurveys/tabid/600/language/en-us/default.aspx

[7] Honigman, B. (2012) 100 Fascinating Social Media Statistics and Figures From 2012. http://www.huffingtonpost.com/brian-honigman/100-fascinating-social-me_b_2185281.html

[8] IBM (2012) http://www.businessofgovernment.org/blog/business-government/next-four-years-citizen-participation

[9] Madden, M. (2010) Older Adults and Social Media, Pew Internet http://pewinternet.org/Reports/2010/Older-Adults-and-Social-Media.aspx

[10] Magro, Michael J. (2012) A review of social media use in e-government. Administrative Sciences 2.2: 148-161

[11] Poblete, B., Garcia, R., Mendoza, M., & Jaimes, A. 2011. Do all birds tweet the same?: characterizing twitter around the world. In Proceedings of the 20th ACM international conference on Information and knowledge management (pp. 1025-1030). ACM

[12] Tumasjan, A., Sprenger, T. O., Sandner, P. G., & Welpe, I. M. (2010) Predicting elections with twitter: What 140 characters reveal about political sentiment. In Proceedings of the 4th International AAAI Conference on Weblogs and Social media.

[13] ECSMDataset (2014) http://kmi.open.ac.uk/people/miriam/ECSM2014/dataset.zip

[14] Bifet, A., Frank, E.: Sentiment knowledge discovery in twitter streaming data. In: Discovery Science. Canberra, Australia (2010)

[15] Pak, A., Paroubek, P.: Twitter as a corpus for sentiment analysis and opinion mining. In: Proceedings of LREC 2010. Valletta, Malta (2010)

[16] Agarwal, A., Xie, B., Vovsha, I., Rambow, O., Passonneau, R.: Sentiment analysis of twitter data. In: Proc. ACL 2011 Workshop on Languages in Social Media. Portland, Oregon (2011)

[17] Barbosa, L., Feng, J.: Robust sentiment detection on twitter from biased and noisy data. In: Proceedings of COLING. Beijing, China (2010)

[18] Kouloumpis, E., Wilson, T., Moore, J.: Twitter sentiment analysis: The good the bad and the omg! In: Proceedings of the ICWSM. Barcelona, Spain (2011)

[19] Saif, H., He, Y., Alani, H.: Alleviating data sparsity for twitter sentiment analysis. In:Proc. Workshop on Making Sense of Microposts (#MSM2012) in WWW 2012. Lyon, France (2012)

[20] Saif, H., He, Y., Alani, H.: Semantic sentiment analysis of twitter. In: Proc. 11th Int. Semantic Web Conf. (ISWC). Boston, MA (2012)

[21] H. Saif, M. Fernandez, Y. He, and H. Alani. Evaluation datasets for twitter sentiment analysis a survey and a new dataset, the sts-gold. In Proceedings, 1st Workshop on Emotion and Sentiment in Social and Expressive Media (ESSEM) in conjunction with AI*IA Conference, Turin, Italy, 2013.

[22] H. Saif, M. Fernandez, Y. He, and H. Alani. 2014a. On Stopwords, Filtering and Data Sparsity for Sentiment Analysis of Twitter. In Proc. 9th Language Resources and Evaluation Conference (LREC), Reykjavik, Iceland, 2014.

[23] H. Saif, M. Fernandez, Y. He, and H. Alani. 2014b. SentiCircles for Contextual and Conceptual Semantic Sentiment analysis of Twitter. Extended Semantic Web Conference (ESWC), Crete, 2014.

[24] H. Saif, M. Fernandez, Y. He, and H. Alani. 2014c. Adapting Sentiment Lexicons using Contextual Semantics for Twitter Sentiment Analysis. In Proceeding of the first semantic sentiment analysis workshop: conjunction with the eleventh Extended Semantic Web conference (ESWC). Crete, Greece.

[25] H. Saif, Y. He, M. Fernandez and H. Alani. 2014d Semantic Patterns for Sentiment Analysis of Twitter, The 13th International Semantic Web Conference (ISWC), Riva del Garda - Trentino Italy

[26] Liu, B.: Sentiment analysis and subjectivity. Taylor and Francis Group(2010)

[27] Aue, A., Gamon, M.: Customizing sentiment classifiers to new domains: A case study. In: Proc. recent advances in natural language processing (RANLP). Borovets, Bulgaria (2005)

[28] O'Connor, B., Balasubramanyan, R., Routledge, B.R., Smith, N.A.: From tweets to polls: Linking text sentiment to public opinion time series. ICWSM 11, 122–129 (2010)

[29] Bollen, J., Mao, H., Zeng, X.: Twitter mood predicts the stock market. Journal of Computational Science 2(1), 1–8 (2011)

[30] Taboada, M., Brooke, J., Tofiloski, M., Voll, K., Stede, M.: Lexicon-based methods for sentiment analysis. Computational linguistics 37(2), 267–307 (2011)

[31] Baccianella, S., Esuli,A., Sebastiani, F.: Sentiwordnet 3.0: An enhanced lexical resource for sentiment analysis and opinion mining. In: Seventh conference on International Language Resources and Evaluation, Malta. Retrieved May. Valletta, Malta (2010)

[32] Wilson, T., Wiebe, J., Hoffmann, P.: Recognizing contextual polarity in phrase-level sentiment analysis. In: Proceedings of the conference on Human Language Technology and Empirical Methods in Natural Language Processing. Vancouver, British Columbia, Canada (2005)

[33] Thelwall, M., Buckley, K., Paltoglou, G.: Sentiment strength detection for the social web. J. American Society for Information Science and Technology 63(1), 163–173 (2012)

[34] Thelwall, M., Buckley, K., Paltoglou, G., Cai, D., Kappas, A.: Sentiment strength detection in short informal text. J. American Society for Info. Science and Technology 61(12) (2010)

[35] Cambria, E.: An introduction to concept-level sentiment analysis. In: Advances in Soft Computing and Its Applications, pp. 478–483. Springer (2013)

[36] Turney, P.D., Pantel, P.,et al.: From frequency to meaning: Vector space models of semantics. Journal of artificial intelligence research 37(1), 141–188 (2010)

[37] Wittgenstein, L.: Philosophical Investigations. Blackwell, London, UK (1953, 2001)

[38] Speriosu, M., Sudan, N., Upadhyay, S. and Baldridge, J. Twitter polarity classification with label propagation over lexical links and the follower graph. In Proceedings of the EMNLP First workshop on Unsupervised Learning in NLP, Edinburgh, Scotland, 2011.

[39] Fernandez, M., Wandhoefer, T., Allen, B., Elisabeth Cano, A., Alani, H.(2014). Using Social Media To Inform Policy Making: To whom are we listening?. In Proceedings of the European Conference on Social Media (ECSM). UK.