The morality machine: tracking moral values in tweets

Teernstra, L.; Putten, P.W.H. van der; Noordegraaf-Eelens, L.; Verbeek, F.J.; Boström, H.; Knobbe, A.; ... ; Papapetrou, P.

The Morality Machine: Tracking Moral Values in Tweets

Livia Teernstra[1,2(✉)], Peter van der Putten[1], Liesbeth Noordegraaf-Eelens[2], and Fons Verbeek[1]

[1] Media Technology, Leiden University, Leiden, The Netherlands
{p.w.h.van.der.putten,f.j.verbeek}@liacs.leidenuniv.nl
[2] Faculty of Social Science, Erasmus University Rotterdam, Rotterdam, The Netherlands
livia@helloliefje.com, eelens@euc.eur.nl

Abstract. This paper introduces The Morality Machine, a system that tracks ethical sentiment in Twitter discussions. Empirical approaches to ethics are rare, and to our knowledge this system is the first to take a machine learning approach. It is based on Moral Foundations Theory, a framework of moral values that are assumed to be universal. Carefully handcrafted keyword dictionaries for Moral Foundations Theory exist, but experiments demonstrate that models that do not leverage these have similar or superior performance, thus proving the value of a more pure machine learning approach.

Keywords: Text classification · Moral values · Social technologies

1 Introduction

There has been growing interest in social sciences research to leverage intelligent data analysis to automatically gather and analyze large amounts of data. A potentially interesting yet relatively unexplored area is ethics, which so far has been approached more theoretically rather than empirically, especially with machine learning methods. The instantaneous and opinionated nature of Internet media such as Twitter provides an immediate outlet for emotions, opinions, information and interactions, loaded with moral perspectives [14]. Accordingly, Twitter is a promising data source for interdisciplinary research on ethics. However, most social science research examines the diffusion of information rather than the content [1,7,17]. Even when content is analysed, this has mostly been focused on commercial or political motivations [2,18]. Likewise within intelligent data analysis, social media monitoring is a popular topic, but it is typically limited to sentiment or opinion mining for business applications, and is lacking theoretical social science foundations. Hence, there is room for an approach that combines morality research with social network content analysis.

The main purpose of this study is to provide an overview of The Morality Machine, a proof of concept system that detects and monitors moral sentiment

H. Boström et al. (Eds.): IDA 2016, LNCS 9897, pp. 26–37, 2016.
DOI: 10.1007/978-3-319-46349-0_3

in Twitter communications, using a text classification approach. It is based on an ethical framework from social psychology called Moral Foundations Theory (MFT), which assumes universal moral foundations exist that can be used to categorize and study ethical problems and discourse [8,11].

As an example debate, this study will explore public opinion on austerity measures in the Eurozone, and specifically the discussion of the Greek exit of the Euro (the 'Grexit'). Austerity is a good topic to explore because it is often discussed in the context of moral hazard [10]. Some state that by bailing out Southern European nations that have shown lack of fiscal discipline, it is encouraging such behaviour rather than criticising it, and that the irresponsible behaviour of these governments is the root cause of the European financial crisis [4]. Conversely, others point out that richer countries have been main beneficiaries of economical support to poorer countries in the past, and that all EU countries have a duty to look after each other and protect the integrity of the EU. Consequently, the Grexit discussion is framed in a moral light, where 'good' and 'bad' nations and policies are distinguished, and there is no shortage of opinions. This austerity dispute will be used to contextualise the methodology since it has the potential to engage all moral foundations.

Related work that classify text into moral foundations typically use dictionary based techniques, meaning that large word lists grounded in psychological theory must first be created and validated manually, as opposed to being discovered automatically by machine learning [2,21,22]. When using these dictionaries, frequencies of morally related words generate moral loadings for texts [5,6,18]. However, the relative importance of these frequencies for detecting certain moral foundations are not derived from evidence. Thus, machine learning algorithms are useful since they can automatically determine lexical indicators for each foundation, without the need to create a dictionary beforehand. Additionally, lexical indicators for each moral foundation can be gleaned from the algorithm, which can be used for further research into moral expressions.

To our knowledge, this is the first study that aims to detect and monitor Moral Foundations using a machine learning approach. It is also the first study to examine moral expressions of the public regarding the Grexit. It uses generally accepted machine learning techniques to explore moral expressions in a natural real world setting, through the use of the Twitter platform. Specifically, this study will determine if supervised machine learning models are able to classify Tweets into moral foundations at an acceptable accuracy, potentially without relying on handcrafted dictionaries.

The remainder of this paper is structured as follows; Sect. 2 provides background on Moral Foundation Theory and text classification. Our methodology is outlined in Sect. 3 and experiments and results are described in Sect. 4. The paper ends with a discussion (Sect. 5) and conclusion (Sect. 6).

2 Related Work

Early ideas in moral reasoning originate in Greek Philosophy. Contemporary moral research asserts that the backbone of our moral decisions lie in a combination of biological and environmental factors [18]. These inert, deep-seated motivations can serve different social functions. This section provides an overview of Moral Foundations Theory (MFT), a framework of assumed universal moral values, and applies it to the case study of the Grexit. It also examines previous research on content analysis using MFT.

2.1 Moral Foundations Theory

The assertion behind MFT is that intrinsic, cognitive responses in individuals can be used to explain the variation in human moral reasoning across cultures [12]. Hence, the theory posits that there is an innate and universal morality which transcends cultural boundaries. This universal morality can be categorised into different foundations, which can be thought of as 'moral building blocks'. Each foundation is fostered within cultures, which serves the purpose of constructing narratives, virtues and institutions. The fostering of foundations differs between groups, where some may emphasize one foundation over another [8]. The foundations can be held simultaneously by individuals and societies, and may conflict with one another.

In the context of this research, six foundations will be used to classify Twitter data. Although there are normally five foundations which form the basis of MFT, a sixth (Liberty - Oppression) has previously been included in the model for other politically driven studies, so we included it [8]. The foundations are briefly described in Table 1, along with example Tweets. The moral foundation which drives opinions can stem from society at large, smaller communities, or individual moral preferences. As such, this study asserts no preference for a specific moral standpoint, as its main focus is learning to classify Grexit Tweets into moral foundations, as a case study for empirical ethics. Also by definition a framework is framed by an underlying theory, which should not be seen as objective or value free. MFT provides a useful framework to distinguish ethical statements, but we do not want to imply it is the only valid one. See for example [20] for a critical review.

2.2 Moral Foundations Text Analysis

In social sciences, dictionary based approaches are predominantly used for text classification. A Moral Foundations Dictionary (MFD) is also available. This dictionary gives linguistic indications for the five basic moral foundations (hence 'liberty' is excluded). The MFD was created for use with the Linguistic Inquiry Word Count (LIWC) program [9]. LIWC is one of the most widely used social science tools for text analysis and is also commonly used for Tweet classification [6,22]. Yet, there is no current research which uses the MFD with LIWC to detect moral foundations in Tweets.

Instead, textual analytics using the MFD with LIWC has been applied in analysis of long texts such as news articles and web blogs, where rhetorical moral assessments were assigned to each text [6,18]. The analysed texts are authored by opinion leaders, such as news media or bloggers, rather than the general public.

Table 1. Descriptions of moral foundations

Foundation	Description	Tweets
Care, Harm	The desire to cherish protect others, identification of a victim and sympathy with him	European control of the IMF is helping Greece
		Greece runs out of funding options despite Euro zone reprieve
Fairness, Cheating	The notions of justice and rights, applied to shared rules in a community. Relates to reciprocal altruism	Greece forced to sell assets and cut spending to pay back debts to EU
		It's easy for the Dutch to go hard on 'Greece'
Loyalty, Betrayal	Relating to 'in-groups'; friends, family, community, as well as showing virtues of patriotism	If I had to choose between #Greece and #Germany, I know which way I'd go...
		Greece may stay in the Eurozone for the time being there are no guarantees it can become a responsible member
Authority, Subversion	Submission to and respect for legitimate authority and traditions	Greece says Euro zone approves reform plan
		German elites are willing to let the Euro crash to guarantee their own political survival
Sanctity, Degradation	Stems from feelings of disgust and contamination. Relating to the virtue that 'the body is a temple', and should not be defiled	There really is no space inside the Euro for a radical left government
		The four-month extension on the Greek debt lowers the risk of Greece leaving the Euro zone
Liberty, Oppression	The resentment of tyranny and desire for autonomy. This is often in tension with the foundation of 'Authority'	Greece needs a path out of the Euro
		Greece really might leave the Euro

For example, research on the Ground Zero Mosque showed that blog authors showed more lexical similarity among virtuous terms for the foundations care, fairness and authority [6]. One can then gather that expression of the other foundations may be constructed differently amongst cultural groups. Due to the differences in textual expressions of moral opinion, dictionary based approaches can be problematic when drawing conclusions about moral reasoning. And given that the dictionary is hand built rather than learned, it is very dependent on it being correct and complete. All in all, the use of MFT and MFD in text analysis is in its infancy, and there is notable room for improvement.

3 Experiments and Results

In this section we provide an overview of the experiments and results, and then review each step of the process and the accompanying results in more detail.

3.1 Overall Procedure

Tweets were collected, and for a random sample frequent keywords were generated as well as bigrams, and bigrams and Tweets were labelled. This gave us 3 data sets: just raw data, raw data with bigrams and raw data with the moral foundations dictionary (MFD). We created two variants of each, one with and one without stop words removed. Skipping stop word removal worked best, so on this data we then carried out learning curve experiments to assess the impact of training set size. For the best performing variants we ran an additional five fold cross validation test. The best model was then deployed to the full data set minus the labeled Tweets to illustrate how the model can be used to track moral sentiment on new data.

3.2 Data Collection

In order to gather initial public reactions to Eurozone meetings, English language Tweets with keywords 'Euro' and 'Greece' were collected from three specific times in 2015, using a custom built streaming Twitter data collector. The search term 'Grexit' was omitted, as it is more prominent in the financial sector, so it excludes Tweets from those who are not familiar with the term. Moreover, 'grexit' tends to carry a certain connotation, focusing only on Greece leaving the Eurozone, rather than economic issues as a whole. The exact dates, number of Tweets and events are outlined in Table 2. Each week of data collection yielded between 4000 and 7000 Tweets, resulting in a total of 18,986 Tweets. The duplicate entries were then removed (including re-Tweets), leaving only unique Tweets (N = 8,292). Note all our coding was done with in Python including the Python Natural Language ToolKit [3].

Table 2. Data collection time periods

Data set	Date Range	N	Event
1	24/02/2015 to 03/03/2015	7,037	Eurozone Finance ministers agreed to extend the Greek bailout for another 4 months
2	28/04/2015 to 04/05/2015	4,856	Eurozone Finance ministers meet to discuss reform packages from Athens
3	11/05/2015 to 23/05/2015	7,066	Athens announces repayments to International Monetary Fund to avoid default

3.3 Data Preparation

Tweets in our system are primarily represented as distributions across sets of keywords (bag of words) and these distributions are then fed into the classifiers. A baseline set of keywords to be used is the Moral Foundations Dictionary (MFD). In our machine learning approach we can already improve over basic MFD label counting because the relationship between MFD keywords and moral foundations classes is learned. In addition we generate keyword sets from the data.

First, data was changed to lower case and hexadecimal codes for emojis were removed, leaving plain text for coding and analysis. Also URLs were replaced with the code 'URL' in order to determine the frequency of link sharing, rather than the most popularly shared links. Next, generated frequency counts for the 100 most common words and 100 most common bi-grams (pairs of consecutive words) were produced for efficiency. Optionally, once the relevance of the data was confirmed by Ethics scholars, a list of common stop words was applied. Stop words contain the most common words in a language and corpus. Removal of these words often yields more accurate predictions in linguistic processing and classification [19]. The most common words were examined without removing stop words, then the most frequently Tweeted words in the data set were added to a standard stop word list, including 'URL', 'greece' and 'euro'. We kept the raw version of the data and keyword sets as well.

The next step was to manually label a random selection of 2000 Tweets with the correct moral foundation. The codes were initially based on the MFD, where related words and synonyms were used to guide classification. Beginning with a dictionary-based approach was useful in order to obtain a more tangible picture of lexical indicators for each of the foundations. However, since the MFD didn't include liberty, a list of synonyms for this foundation was created. Then, detailed descriptions of each of the foundations were used to better understand the nuances in each foundation, as outlined in the work of Graham et al. [8] and Haidt [11]. So the combination of specific, related words as well as detailed descriptions of the foundations were used to code the Tweets.

Manual labeling of Tweets is a challenging task, given the inherent ambiguity of some Tweets, the short length of Tweets, the potential of multiple moral foundations being covered in Tweets and use of writing styles such as irony, sarcasm, satire or mere trolling. We choose not to filter out hard to label Tweets as this could bias the sample, nor did we want to include an 'unknown' category, as it would limit the usefulness of the model for monitoring. We considered approaching it as a multi-label problem, however in our view there were far more cases where the Tweet was simply hard to label due to ambiguity than that there was sufficient evidence to conclude that multiple foundations were being addressed, also given the short lengths of Tweets in contrast to the longer texts (blog posts, articles) studied in related work. For similar reasons we discounted an approach where we would have scored the Tweets on the various dimensions to a particular degree. So we kept it simple by manually labeling each Tweet with a single label. These other approaches are indeed interesting areas for future research, but we decided to generate baseline results first.

The most frequent class occurred in 21 % of cases, thus a majority vote baseline model has an accuracy of 21 %. Two coders also labeled a set of bigrams (N = 112) to determine the degree which coders could agree on moral classes. The coders agreed on 66 % of the classifications. It is acknowledged that coding bigrams more difficult than coding Tweets, yet it gives an indication of intercoder agreement in classifying moral foundations with little contextual information. Therefore, any accuracy higher than 21 % is an improvement of the classifier over selection of the most frequently occurring class, and any accuracy around 66 % would show that the classifier is matching human classification of bigrams.

4 Modeling and Evaluation

Previous research using the MFD was conducted on long texts, examining moral loadings and linguistic relations between these texts [6,18]. Since Tweets are short, single-label output (one classification per Tweet) was chosen over multiple labels. We used Multinomial Naive Bayes (NB) and Maximum Entropy (ME) as classification algorithms [15]. The most relevant key difference for this study relates to the independence of features, where NB assumes conditional independence and ME can exploit contextual information (such relationships between words) for classification. Despite the fact that the independence assumption is typically violated, Naive Bayes has shown in general to be a robust classification method, especially for noisy, high variance problems [16].

The data was split into a training (N = 1,300) and a test set (N = 700). Classifiers were built on the raw data (no stop word removal) and the clean data (stop word removal). NB showed higher overall accuracy (raw = 65 %, clean = 64 %) than ME (raw = 57 %, clean = 55 %). Removing stop words did not seem to increase classifier accuracy for either algorithm. To study the impact of training set size, we trained classifiers on raw data training sets of increasing size, with increments of 100 Tweets, up until a maximum of 1300, whilst keeping the test set constant. We also varied the feature set between the raw features, raw features with the MFD and raw features with the bigrams. The results in Fig. 1 show that NB performance is not significantly improved by adding the dictionary, and performance drops if bigrams are added. Detailed results for ME are omitted for brevity, but ME performs best with the addition of bi-grams, achieving 57 % accuracy, and for training set sizes of 300 instances or more, NB outperforms ME. Under almost all conditions, the NB classifier outperformed ME, shown in Fig. 2 (best feature set for each). Over time, the learning curves of both classifiers flattens. It is therefore expected that additional training data will not improve classifier accuracy.

These results were confirmed by a 5-fold cross-validation comparison, where the mean accuracy for NB was 64.7 % (SD = 0.03, p = .000) compared with the ME mean accuracy of 54.2 % (SD = 0.02, p = .000). The difference in classifier accuracy is significant (T = 13.9, p = .000). Overall, the NB classifier is 10 % points more accurate than ME in classifying Tweets into moral foundations.

Confusion matrices, precision, recall and F measures are provided in Tables 3 and 4 for this model (ME results are omitted for brevity). These tables refer to a subsample of Tweets used for training and testing the model. The most frequently correctly classified foundation was care (N = 108), followed by authority (N = 87). Liberty was the least often correctly classified foundation (N = 38). Despite care being most frequently classified correctly, the precision, recall and

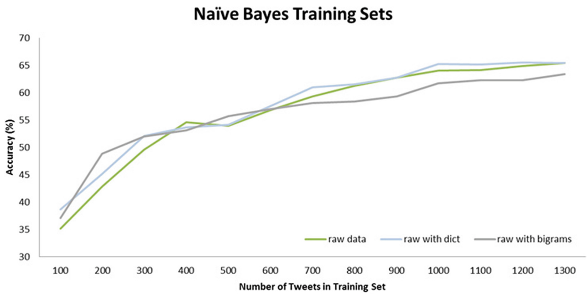

Fig. 1. NB classifier test set accuracy for different training set sizes

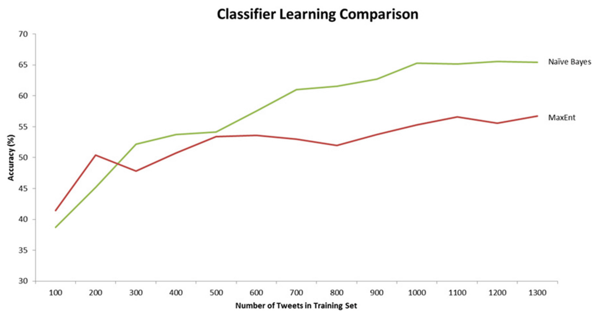

Fig. 2. NB and ME classifiers test set accuracy for different training set sizes

F-measures in Table 4 show otherwise. Taking the relative accuracy into account, authority was the most accurately classified (F = 0.73), followed by sanctity (F = 0.66) and care (F = 0.63). Fairness was the least accurate (F = 0.58). Therefore, this model overall works best in identifying Tweets stemming from the foundation of authority.

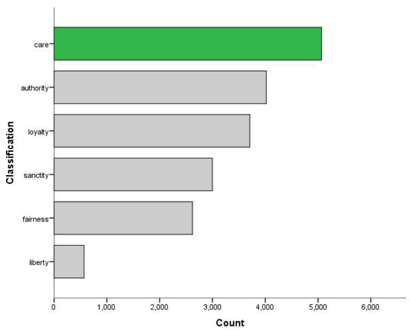

Fig. 3. Classification of all Tweets

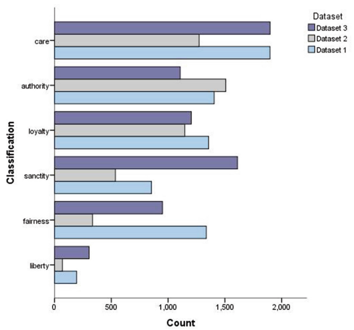

Fig. 4. Classification of Tweets per time period

Table 3. Naive Bayes Confusion Matrix, comparing actual frequencies (rows) and predicted frequencies (columns)

	Authority	Care	Fairness	Liberty	Loyalty	Sanctity
Authority	<87>	13	2	.	9	4
Care	6	<108>	7	5	13	7
Fairness	8	25	<61>	7	23	6
Liberty	4	12	4	<38>	10	2
Loyalty	10	15	5	1	<77>	8
Sanctity	10	23	3	.	14	<73>

Table 4. Naive Bayes accuracy for each class

	TP	FN	FP	Precision	Recall	F-Measure
Authority	87	28	38	0.696	0.757	0.725
Care	108	38	88	0.551	0.74	0.632
Fairness	61	69	21	0.744	0.469	0.575
Liberty	38	32	13	0.745	0.543	0.628
Loyalty	77	39	69	0.527	0.664	0.589
Sanctity	73	27	27	0.73	0.593	0.655
Total	444	256	256			

4.1 Deployment

The most accurate algorithm, with the least training time required (NB, raw data, no MFD) was trained with all labeled data (N = 2000). Following learning, the model was used to classify the remaining Tweets (N = 16,986). Deployment of the model enabled analysis of changes in moral concerns following key meetings regarding the Greek exit of the Eurozone. There were 3 different time frames where Tweets were collected. Figure 3 demonstrates that Tweets were classified most frequently in the care category (N = 5068). Hence, over the first half of 2015, individuals on Twitter showed care as the primary moral concern in the Grexit debate, authority as the second, and loyalty as the third. However, over time, the predominant moral underpinning of the rhetoric can change. Indeed, Fig. 4 shows that in the first and third time periods, care was the most common concern, whereas in the second time period, authority dominated the discussion overall. In all time periods, liberty was the least discussed foundation, especially in data set 2, where the foundation barely emerged. Thus, the hypothesis that liberty is a necessary foundation for this research is disconfirmed. Application of the classifier shows that people on Twitter are not primarily concerned with liberty or oppression of any party in this debate.

The running means of the Tweets made through the data sets shown in Figs. 5, 6 and 7. These means show the discourse over the number of Tweets,

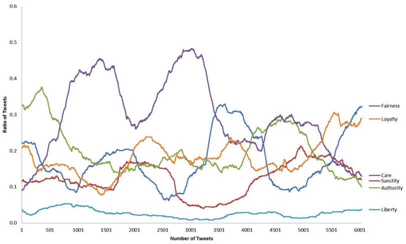

Fig. 5. Running Mean for data set 1

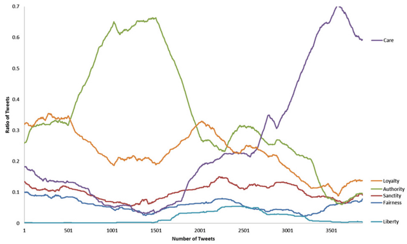

Fig. 6. Running Mean for data set 2

disregarding the time and day they were sent. This compensates for different time zones and allows time for news to disseminate. In Fig. 5, care has two dominant peaks, despite initial discussion referring to authority. Towards the end of the week, loyalty and fairness was behind the discussion. Figure 6 shows that in the first Tweets of data set 2, authority is a key concern, but is replaced with care in the later Tweets. In the final data set (see Fig. 7) there are multiple points of interest. The first peak shows that authority drove the early Tweets, followed by loyalty. At the end of the discussion, care became the dominant foundation. Liberty was not a relative point of concern in any data set.

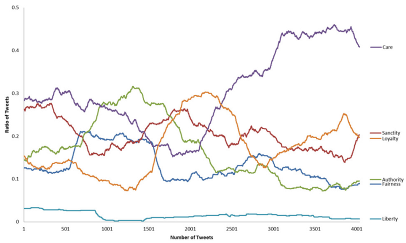

Fig. 7. Running Mean for data set 3

One key finding is that the data shows that public discussion is not in line with analysts moral view of the situation, as shown especially in Figs. 6 and 7. Economic analysts tend to approach Grexit discussion from angles of fairness and loyalty, such as the potential loyalty of Greece shifting beyond the European Union if they were to leave [13]. The public discussion frequently centers on the foundation of care, referring to helping Greece with extensions or bailouts. Loyalty was indeed more present than fairness, but clearly care and authority were salient moral foundations especially in data sets 2 and 3.

5 Discussion

The results show that the NB classifier is a good starting point for attributing moral foundations to Tweets. The three most accurately classified foundations (care, authority and sanctity) agree with previous research [6]. The learning curves show that coding more than 2,000 Tweets for training a classifier will not improve accuracy, at least for the feature sets and ground truth used.

For time periods monitored care is the primary moral concern of the public, which is somewhat in contrast to the dominant economic views that are concerned more with loyalty and fairness.

Perhaps most surprisingly, results also showed that addition of words from the MFD did not improve model accuracy. Therefore, the usefulness of the MFD in a frequency based classification approach is called into question. If using this dictionary is desired in future research, improvements to the dictionary should be made by including words identified as the most informative features following training the NB algorithm. However, the efforts in improvement of the MFD may only have marginal implications for model accuracy. It may be prudent to discontinue the MFD, since these dictionaries are costly to build and maintain, and a pure machine learning approach has similar accuracy and uses less assumptions.

6 Conclusion

This study presents several experiments to determine if machine learning methods can be used to accurately detect moral foundations in Tweets regarding the Grexit. A Naive Bayes (NB) model trained on raw data was 10 % points more accurate than the a Maximum Entropy (ME) model, with best results achieved on raw data without bigram or Moral Foundations Dictionary (MFD) attributes. Specifically, the fact that the NB model doesn't require the handcrafted MFD is an interesting result.

At this point, it is difficult to compare with other moral foundation classification research, as thus far none have used a machine learning approach. However, the accuracy of the NB model is comparable to the agreement of moral classification between humans for bigrams (64.7 % compared with 66 %, respectively). Moreover, the model is roughly 3 times more accurate than the ZeroR measure of 21.4 %. Hence, using a NB classifier is a good starting point for categorization of Tweets into their dominant moral foundations.

References

1. Adamic, L.A., Glance, N.: The political blogosphere and the 2004 US election: divided they blog. In: Proceedings of the 3rd International Workshop on Link Discovery, pp. 36–43. ACM (2005)
2. Agarwal, A., Xie, B., Vovsha, I., Rambow, O., Passonneau, R.: Sentiment analysis of Twitter data. In: Proceedings of the Workshop on Languages in Social Media, pp. 30–38. Association for Computational Linguistics (2011)
3. Bird, S., Klein, E., Loper, E.: Natural Language Processing with Python. O'Reilly Media, Sebastopol (2009)
4. Bond, J.: It aint over till the fat lady sings. Sens-Public (2012). http://www.sens-public.org/article979.html

5. Clifford, S., Jerit, J.: How words do the work of politics: moral foundations theory and the debate over stem cell research. J. Politics **75**(03), 659–671 (2013)
6. Dehghani, M., Sagae, K., Sachdeva, S., Gratch, J.: Linguistic analysis of the debate over the construction of the Ground Zero Mosque. J. Inform. Technol. Politics **11**, 1–14 (2014)
7. Freelon, D.: On the interpretation of digital trace data in communication and social computing research. J. Broadcast. Electron. Media **58**(1), 59–75 (2014)
8. Graham, J., Haidt, J., Koleva, S., Motyl, M., Iyer, R., Wojcik, S.P., Ditto, P.H.: Moral foundations theory: the pragmatic validity of moral pluralism. Adv. Exp. Soc. Psychol. **47**, 55–130 (2013)
9. Graham, J., Haidt, J., Nosek, B.A.: Liberals and conservatives rely on different sets of moral foundations. J. Pers. Soc. Psychol. **96**(5), 1029 (2009)
10. Grauwe, P.: The eurozone as a morality play. Intereconomics Rev. Eur. Econ. Policy **46**(5), 230–231 (2011)
11. Haidt, J.: The righteous mind: why good people are divided by politics and religion. Vintage, New York (2012)
12. Haidt, J., Joseph, C.: Intuitive ethics: how innately prepared intuitions generate culturally variable virtues. Daedalus **133**(4), 55–66 (2004)
13. Lazarou, A.: Greece: The many faces of Yanis Varoufakis. Green Left Weekly (104) (2015)
14. Lazer, D., Pentland, A.S., Adamic, L., Aral, S., Barabasi, A.L., Brewer, D., Christakis, N., Contractor, N., Fowler, J., Gutmann, M., et al.: Life in the network: the coming age of computational social science. Science **323**(5915), 721 (2009). (New York, NY)
15. Manning, C.D., Schütze, H.: Foundations of Statistical Natural Language Processing. MIT Press, Cambridge (1999)
16. van der Putten, P., van Someren, M.: A bias-variance analysis of a real world learning problem: the CoIL Challenge 2000. Mach. Learn. **57**(1), 177–195 (2004)
17. Ratkiewicz, J., Conover, M., Meiss, M., Gonçalves, B., Patil, S., Flammini, A., Menczer, F.: Truthy: mapping the spread of astroturf in microblog streams. In: Proceedings of the 20th International Conference Companion on World Wide Web, pp. 249–252. ACM (2011)
18. Sagi, E., Dehghani, M.: Measuring moral rhetoric in text. Soc. Sci. Comput. Rev. **32**(2), 132–144 (2014)
19. Saif, H., Fernández, M., Alani, H.: Automatic stopword generation using contextual semantics for sentiment analysis of Twitter. In: CEUR Workshop Proceedings, vol. 1272 (2014)
20. Suhler, C.L., Churchland, P.: Can innate, modular foundations explain morality? Challenges for Haidt's moral foundations theory. J. Cogn. Neurosci. **9**, 2103–2116 (2011)
21. Taboada, M., Brooke, J., Tofiloski, M., Voll, K., Stede, M.: Lexicon-based methods for sentiment analysis. Comput. Linguist. **37**(2), 267–307 (2011)
22. Tumasjan, A., Sprenger, T.O., Sandner, P.G., Welpe, I.M.: Predicting elections with Twitter: what 140 characters reveal about political sentiment. ICWSM **10**, 178–185 (2010)