

The Language of Liberty: A preliminary study

Oscar Araque
o.araque@upm.es
Universidad Politécnica de Madrid
Madrid, Spain

Lorenzo Gatti
l.gatti@utwente.nl
University of Twente
Enschede, The Netherlands

Kyriaki Kalimeri
kyriaki.kalimeri@isi.it
ISI Foundation
Turin, Italy

ABSTRACT

Quantifying the moral narratives expressed in the user-generated text, news, or public discourses is fundamental for understanding individuals' concerns and viewpoints and preventing violent protests and social polarisation. The Moral Foundation Theory (MFT) was developed precisely to operationalise morality in a five-dimensional scale system. Recent developments of the theory urged for the introduction of a new foundation, liberty. Being only recently added to the theory, there are no available linguistic resources to assess liberty from text corpora. Given its importance to current social issues such as the vaccination debate, we propose a data-driven approach to derive a liberty lexicon based on aligned documents from online encyclopedias with different worldviews. Despite the preliminary nature of our study, we show proof of the concept that large encyclopedia corpora can point out differences in the way people with contrasting viewpoints express themselves. Such differences can be used to derive a novel lexicon, identifying linguistic markers of the liberty foundation.

KEYWORDS

moral foundations theory, natural language processing, word embeddings, Wikipedia

ACM Reference Format:

Oscar Araque, Lorenzo Gatti, and Kyriaki Kalimeri. 2021. The Language of Liberty: A preliminary study. In *Companion Proceedings of the Web Conference 2021 (WWW '21 Companion)*, April 19–23, 2021, Ljubljana, Slovenia. ACM, New York, NY, USA, 4 pages. <https://doi.org/10.1145/3442442.3452351>

1 INTRODUCTION

Moral values are fundamental to our decision-making process on everyday matters. When taking a stance towards a social issue, for instance, global warming or vaccine adherence, we consult - consciously or unconsciously - our moral system of values. Extracting and analysing moral content from user-generated text or public discourse, in general, is critical to understanding the decision-making process of individuals while getting a large scale perspective of evolving narratives [14]. The Moral Foundations Theory (MFT) was created to explain morality across cultures [7]. The theory initially proposed five foundations, namely *care*, *fairness*, *loyalty*, *authority*, and *sanctity*, while more recently, the theory was enhanced with a new sixth dimension: *liberty*.

This paper is published under the Creative Commons Attribution 4.0 International (CC-BY 4.0) license. Authors reserve their rights to disseminate the work on their personal and corporate Web sites with the appropriate attribution.

WWW '21 Companion, April 19–23, 2021, Ljubljana, Slovenia

© 2021 IW3C2 (International World Wide Web Conference Committee), published under Creative Commons CC-BY 4.0 License.

ACM ISBN 978-1-4503-8313-4/21/04.

<https://doi.org/10.1145/3442442.3452351>

MFT theoretical framework presents libertarians to have a unique moral-psychological profile, endorsing the principle of liberty as an end and devaluing many of the moral concerns typically endorsed by conservatives. Libertarianism is a political philosophy and movement that upholds liberty as a core principle [3] and express the extreme side of the liberty moral foundation. Analysing the psychological dispositions of libertarians, Iyer et al. [9], found that libertarians are consistently less concerned about individual-level concerns such as harm, benevolence, and altruism. They are also much less concerned with group-level moral issues, for instance, conformity, loyalty, and tradition, that are typically associated with conservative morality. Libertarians' cognitive style is presented to depend less on emotion and more on reason than conservatives. In a study carried out in more than ten countries, [10] found that one of the most reliable differences between liberals and conservatives is that individuals susceptible to threat and resistant to change typically find greater comfort in conservative rather than liberal ideologies.

MFT is broadly adopted in the computational social science field since it defines a clear taxonomy of values together with a term dictionary, the Moral Foundations Dictionary (MFD) [6], an essential resource for natural language processing applications. The MFD creators highlight the difficulty of creating such a resource since linguistic, cultural, and historical context reflect on language usage. Among the most significant limitations of the MFD, we have: (i) a limited amount of lemmas and stem of words; (ii) "radical" lemmas rarely used in everyday language, for instance, "homologous". and "apostasy"; and (iii) an association with a moral bi-polar scale, so-called vice and virtue, but without any indication of polarity or "strength". (iv) the "liberty" foundation is not considered due to its very recent addition to the main theory. The MoralStrength lexicon [1] addresses many of the shortcomings of MFD, expanding the number of lemmas per foundation with more commonly used terms introducing the notion of "moral polarity". Despite addressing the most critical shortcomings of the MFD and its efficiency in generic moral prediction tasks [1], the MoralStrength lexicon does not include the liberty foundation.

Here, we lay the groundwork for a linguistic resource that assesses the liberty moral dimension in people's narratives. We consider the Wikipedia¹ pages and their Conservapedia counterparts as a natural experiment. While studies have shown that Wikipedia articles exhibit a quality comparable to conventional encyclopedias, it has been criticised by Christian conservatives to show strong liberal bias [8], especially in controversial issues such as abortion, homosexuality, and global warming. Conservapedia was created

¹Link to Wikipedia site: <https://www.wikipedia.org/>

precisely to express the views of several conflicting topics according to right-conservative ideas². Seen through the lenses of our theoretical framework, the Moral Foundations Theory, Conservapedia aims to defend the moral values that its readership believes not adequately expressed in the respective Wikipedia pages. As a starting step, we restrict our analysis to the categories that are directly related to the political domain³.

The scope of this study is to provide researchers with a resource to gauge the moral value of *liberty* from the user-generated text. Based on the well-established theoretical framework of MFT, we combine a natural experiment approach with unsupervised machine learning techniques to derive a liberty lexicon based on online encyclopedia documents. Such lexicon will contribute to the computational linguistic resources that tackle moral values in large, user-generated corpora, given their importance to many current social issues, among which the vaccination debate [5, 11].

2 DATA COLLECTION AND METHODS

We propose a machine learning approach where, without any a priori linguistic information, we will attempt to identify the linguistic markers that characterize the expression of liberal and conservative values. We are based on the assumption that editors of Conservapedia created a new entry on a topic they believed discussed on Wikipedia in a very liberal way [8].

Starting from the title of each Conservapedia page, we searched for the corresponding page in Wikipedia. We managed to align more than 37,000 articles across Wikipedia and Conservapedia; of these, about 28,000 pages share an identical title, while the remaining ones are aligned based on redirect pages. In total, the whole corpus contains 106 million tokens and 558,000 unique words. We performed additional filtering to address the political domain, using page categories that refer to political issues. Also, to improve the dataset’s quality, we have computed a length ratio that compares the document lengths of a Wikipedia/Conservapedia pair. We define the ratio as the number of words in a Wikipedia document over the number of terms of the Conservapedia pair. In this way, we drop document pairs with a ratio higher than 10, resulting in 2,026 documents, 1,013 from Conservapedia and 1,013 from Wikipedia. Basic preprocessing was performed on the original text, extracted using WikiExtractor⁴; in particular, we removed stopwords, normalized tokens (e.g., transforming numeric expressions to *num*), filtered punctuation and short words (i.e., terms shorter than three letters).

Inspired by Turney et al. [15], we define two sets of seed words representing the conservative and liberal orientation. Starting from the liberty questionnaire [9], we crafted a set of seed words, shown in Table 1, which are then used to expand the lexicon’s vocabulary. Taking into account word frequency, we obtain an annotated lexicon that models libertarian and conservative word usage. Such a resource can be employed to predict the presence and the polarity of the “liberty” moral foundation in a previously unseen text. As in Turney et al. [15], we compute each word’s polarity through a word embedding model. To this end, we used gensim’s doc2vec [13] to

obtain both document (not used in the experiments here reported) and word vectors after lemmatizing the corpus with Spacy. We used the default doc2vec options, but we increased the embeddings’ dimension to 300, a standard parameter setting.

Using the resulting word embedding model, we implemented a lexicon generation method based on the cosine similarity between the selected seed words and words from the available documents. In this way, let S_C be the set of seed words for the conservative orientation, and S_L the seed words for the liberal direction. We compute the moral polarity of a word w_i from the documents as

$$\sum_{w_j \in S_L} \text{sim}(w_i, w_j) - \sum_{w_k \in S_C} \text{sim}(w_i, w_k)$$

where *sim* is the cosine similarity, as computed by the word embedding model. The polarity is positive if w_i is related to the liberal seed words and negative if the relation occurs towards the conservative seed words. To obtain the polarities, we use the word embedding model we trained on our full dataset (before category filtering) to ensure that the word usage characterization of the embeddings is more accurate than that of a pre-trained model.

Table 1: Words used as seed words for the lexicon generation method. Words in bold originate from the questionnaire proposed by Iyer et al. [9].

Libertarian seed words
liberty, society, free, freedom, choice, equal , reformist, libertarian, rational, broad-minded, high-minded, indulgent, intelligent, reasonable, unbiased, unbigoted, unconventional.
Conservative seed words
private, property, norm, tradition , conserve, nation, traditional, right, conventional, orthodox, preserve, national, army, family, bank, capital, republican, country

3 PRELIMINARY RESULTS AND DISCUSSION

Even if at an early stage, our approach shows promising results; Figure 1 depicts the exploratory view of the word frequency distribution in the political category pages of the two encyclopedias: Conservapedia and Wikipedia. Both axes correspond to the rank frequency with which a term occurs in the respective category of documents in Wikipedia (horizontal axis) and Conservapedia (vertical axis). The rank increases from left to right and low to high. Hence, at the top right of Figure 1, we find the most commonly used words by both communities. Respectively, the most common words for Conservapedia and Wikipedia are represented at the top left and bottom right.

We can notice marked differences in word usage in the two resources: Wikipedia authors tend to use more objective/neutral words (*affordable care, american politician*), in addition to many non-political terms. In Conservapedia prevail derogatory terms such as *rino*, which stands for “Republican In Name Only”, and

²Link to Conservapedia site: <https://conservapedia.com/>

³The categories included in this study are ‘Republicans’, ‘Conservatives’, ‘Republican Party’, ‘Liberalism’, ‘Democrats’, ‘Liberals’, ‘Democratic Party’.

⁴<https://github.com/attardi/wikiextractor>

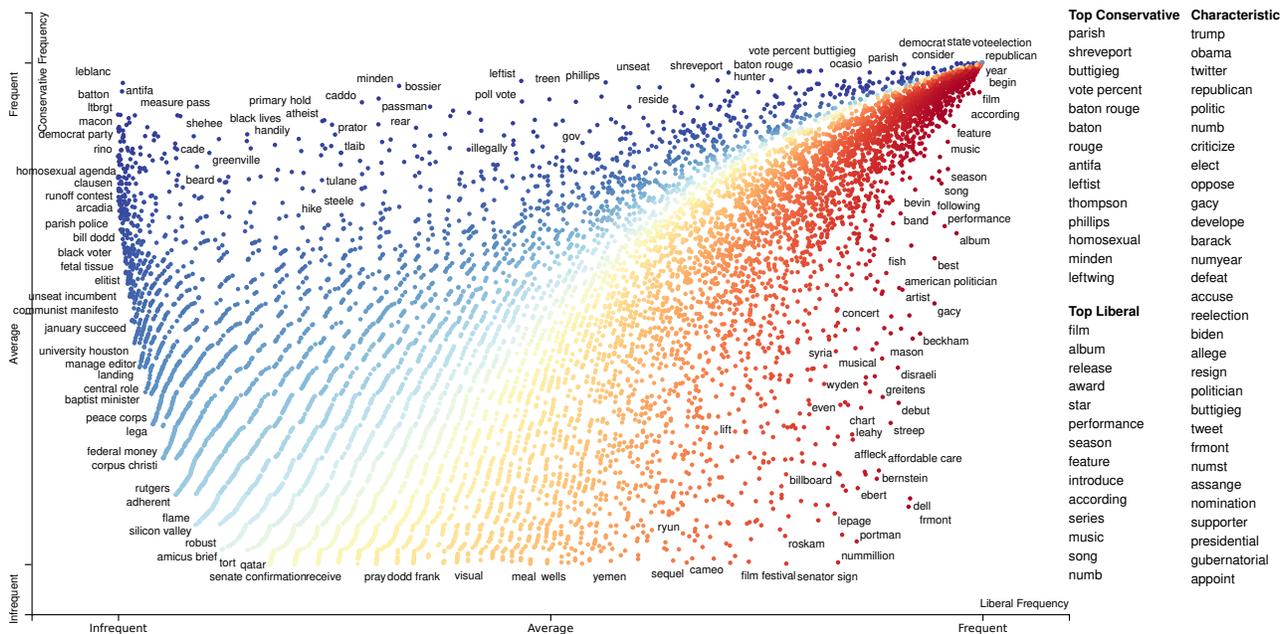


Figure 1: Visualization of words and phrase distributions in Wikipedia and Conservapedia sample of the corpus [12]. Points are colored red or blue based on the association of their corresponding terms with Conservapedia or Wikipedia. The most associated terms are listed under “Top Conservapedia” and “Top Wikipedia” headings.

Democrat Party, but also topics of high concern to the conservative community such as the *homosexual agenda*, *communist manifesto*, and *fetal tissue*.

Table 2: Top ten lemmas for the liberal and the conservative side, ordered by polarity value, and three selected, indicative high-ranking terms from both sides.

Top 10 Liberal terms		Top 10 Conservative terms	
absurd	2.41	cemetery	-2.49
honest	2.35	diocese	-2.35
irrational	2.30	territory	-2.29
energetic	2.19	province	-2.24
misunderstand	1.96	treasury	-2.09
economics	1.90	mansion	-2.04
fairness	1.89	monastery	-2.02
inappropriate	1.87	principality	-2.01
crazy	1.86	church	-2.00
innate	1.81	kingdom	-1.99
agnostic	1.69	heritage	-1.55
scientifically	1.53	directorate	-1.52
atheism	1.46	force	-1.34

Following the embedding-based approach with our initial seed lemmas, we derive a lexicon that encodes the linguistic range of the “liberty” dimension in this corpus. Table 2 shows the top ten emerging lemmas per dimension ranked by absolute moral polarity, while the last three elements per dimension are manually

selected. Despite the brevity of the excerpt, we can draw some initial remarks. Liberal terms are more related to economy (economics), emotional and cognitive states (absurd, honest, energetic, irrational), and moral reasoning (irrational, fairness). On the Conservative side, the lemmas are in general about property-owning (territory, mansion), religious views (cemetery, diocese), and authority (principality, kingdom). Both sides exhibit terms that are in line with the psychological profiles depicted in the moral and political psychology literature [4, 9, 10]. Going through the emerged elements, we continue to encounter words that are constant with each side’s moral profiles. The full generated lexicon is available at <https://github.com/oaraque/moral-foundations>.

Importantly, we notice the extent to which the selection of seed words impacts the resulting lexicon. For instance, since many adjectives are included in the libertarian seed words (e.g., *reformist*, *rational*, *broad-minded*), the lexicon has a prevalence of adjectives with higher polarity for both sides. Additionally, the inclusion of some words loosely related to religion in the conservative seed set, such as *tradition* and *orthodox*, causes some high-polarity conservative terms are associated with this topic. The term “orthodox” was initially included in the list of seed words as a synonym of conservative, in the general sense of “conforming to the canon of a philosophical current”. This calls for careful refining of the seed words employed for the lexicon creation and better embeddings that take different word senses into account.

The proposed model is based on word embeddings known to express the hidden distributional relations of words; in real-life applications with high possible societal impact, we want to make

sure that we adequately express concrete concepts. In this case, “irrational” is not an adjective that characterises the libertarian profile according to the literature. The inclusion of additional data sources and improved strategies for expanding the lexicon should help filter these outliers.

This study’s future steps include methodological improvements to address the current lexicon shortcomings and a more sophisticated approach to seed word generation. Enhancing the lexicon with sentiment scores [2] will improve the interpretability of the resource when employed to analyse complex societal discourses. Moreover, we aim to extend this approach to other categories of pages to capture different liberty dimension’s nuances, such as economic and societal. Finally, we acknowledge the conceptual limitation of our approach; Wikipedia strives for a neutral point of view, avoiding the claim of an objective and immutable truth. With its vast author population, Wikipedia is likely to express the viewpoints of the entire conservative-libertarian spectrum. The proposed setup is one of the few large-scale, open-source text corpora, where the same concepts are presented in two distinct ways, fine-tuned to express a different set of moral values.

Libertarian values are essential in understanding decision making in a variety of contexts. Quantifying the expressed moral in user-generated text from social media, news journals, or even online fora will help unveil the drivers of moral judgments towards critical social issues, such as poverty, radicalisation, and vaccine adherence.

ACKNOWLEDGMENTS

This work was partly supported by the European Union’s Horizon 2020 Research and Innovation Programme under project Participation (grant agreement no. 962547). KK acknowledges support from the Lagrange Project of the Institute for Scientific Interchange Foundation (ISI Foundation) funded by Fondazione Cassa di Risparmio di Torino (Fondazione CRT).

REFERENCES

- [1] Oscar Araque, Lorenzo Gatti, and Kyriaki Kalimeri. 2020. MoralStrength: Exploiting a moral lexicon and embedding similarity for moral foundations prediction. *Knowledge-based systems* 191 (2020).
- [2] O. Araque, L. Gatti, J. Staiano, and M. Guerini. 2019. DepecheMood++: a Bilingual Emotion Lexicon Built Through Simple Yet Powerful Techniques. *IEEE Transactions on Affective Computing* (2019). <https://doi.org/10.1109/TAFFC.2019.2934444>
- [3] Randy E Barnett. 2014. *The structure of liberty: Justice and the rule of law*. OUP Oxford.
- [4] Dana R Carney, John T Jost, Samuel D Gosling, and Jeff Potter. 2008. The secret lives of liberals and conservatives: Personality profiles, interaction styles, and the things they leave behind. *Political Psychology* 29, 6 (2008), 807–840.
- [5] Alessandro Cossard, Gianmarco De Francisci Morales, Kyriaki Kalimeri, Yelena Mejova, Daniela Paolotti, and Michele Starnini. 2020. Falling into the echo chamber: the Italian vaccination debate on Twitter. In *Proceedings of the International AAAI Conference on Web and Social Media*, Vol. 14. 130–140.
- [6] Jesse Graham, Jonathan Haidt, and Brian A Nosek. 2009. Liberals and conservatives rely on different sets of moral foundations. *Journal of personality and social psychology* 96, 5 (2009), 1029.
- [7] Jonathan Haidt and Craig Joseph. 2004. Intuitive ethics: how innately prepared intuitions generate culturally variable virtues. *Daedalus* 133, 4 (2004), 55–66. <https://doi.org/10.1162/0011526042365555>
- [8] Christoph Hube. 2017. Bias in Wikipedia. In *Proceedings of the 26th International Conference on World Wide Web Companion*. 717–721.
- [9] Ravi Iyer, Spassena Koleva, Jesse Graham, Peter Ditto, and Jonathan Haidt. 2012. Understanding libertarian morality: The psychological dispositions of self-identified libertarians. *PloS one* 7, 8 (2012), e42366.
- [10] John T Jost, Jack Glaser, Arie W Kruglanski, and Frank J Sulloway. 2003. Political conservatism as motivated social cognition. *Psychological bulletin* 129, 3 (2003), 339.
- [11] Kyriaki Kalimeri, Mariano G. Beiró, Alessandra Urbinati, Andrea Bonanomi, Alessandro Rosina, and Ciro Cattuto. 2019. Human values and attitudes towards vaccination in social media. In *Companion Proceedings of The 2019 World Wide Web Conference*. 248–254.
- [12] Jason S. Kessler. 2017. Scattertext: a Browser-Based Tool for Visualizing how Corpora Differ. In *Proceedings of ACL-2017 System Demonstrations*. Association for Computational Linguistics, Vancouver, Canada.
- [13] Quoc Le and Tomas Mikolov. 2014. Distributed representations of sentences and documents. In *International Conference on Machine Learning*. 1188–1196.
- [14] Ashwin Rao, Fred Morstatter, Minda Hu, Emily Chen, Keith Burghardt, Emilio Ferrara, and Kristina Lerman. 2020. Political Partisanship and Anti-Science Attitudes in Online Discussions about Covid-19. *arXiv preprint arXiv:2011.08498* (2020).
- [15] Peter D Turney and Michael L Littman. 2003. Measuring praise and criticism: Inference of semantic orientation from association. *ACM Transactions on Information Systems (TOIS)* 21, 4 (2003), 315–346.