LibertyMFD: A Lexicon to Assess the Moral Foundation of Liberty.

Oscar Araque o.araque@upm.es Universidad Politécnica de Madrid, ETSI Telecomunicación 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 to operationalise morality in a five-dimensional scale system. Recent developments of the theory urged for the introduction of a new foundation, the Liberty Foundation. Being only recently added to the theory, there are no available linguistic resources to assess whether liberty is present in text corpora. Given its importance to current social issues such as the vaccination debate, we propose two data-driven approaches, deriving two candidate lexicons generated based on aligned documents from online news sources with different worldviews. After extensive experimentation, we contribute to the research community a novel lexicon that assesses the liberty moral foundation in the way individuals with contrasting viewpoints express themselves through written text. The LibertyMFD dictionary can be a valuable tool for policymakers to understand diverse viewpoints on controversial social issues such as vaccination, abortion, or even uprisings, as they happen and on a large scale.

CCS CONCEPTS

• Computing methodologies \rightarrow Lexical semantics.

KEYWORDS

moral foundations theory, natural language processing, word embeddings, liberty, lexicon, moral values

ACM Reference Format:

Oscar Araque, Lorenzo Gatti, and Kyriaki Kalimeri. 2022. LibertyMFD: A Lexicon to Assess the Moral Foundation of Liberty.. In *Conference on Information Technology for Social Good (GoodIT'22), September 7–9, 2022, Limassol, Cyprus.* ACM, New York, NY, USA, 7 pages. https://doi.org/10.1145/3524458. 3547264

GoodIT'22, September 7–9, 2022, Limassol, Cyprus

© 2022 Association for Computing Machinery.

ACM ISBN 978-1-4503-9284-6/22/09...\$15.00

https://doi.org/10.1145/3524458.3547264

1 INTRODUCTION

Moral values are fundamental to our decision-making process on everyday matters. When taking a stance on 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 [28]. The Moral Foundations Theory (MFT) was created to explain morality across cultures [16]. 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*.

The MFT theoretical framework presents libertarians with 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 [6] and expresses the extreme side of the liberty moral foundation. Analysing the psychological dispositions of libertarians, Iyer et al. [19], found that libertarians are consistently less concerned about individuallevel concerns such as harm, benevolence, and altruism. They are also much less concerned with group-level moral issues, such as conformity, loyalty, and tradition, typically associated with conservative morality. Libertarians' cognitive style depends less on emotion and more on reason than conservatives. The liberty/oppression foundation, as described in the MFT, "deals with the domination and coercion by the more powerful upon the less so" [15].

The MFT is broadly adopted in the computational social science field since it defines a clear taxonomy of values together with a lexicon, the Moral Foundations Dictionary (MFD) [14], an essential resource for natural language processing applications. The MFD creators highlight the difficulty of creating such a resource since linguistic, cultural, and historical contexts reflect on language usage. Among the most significant limitations of the original 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 numeric indication of weight or importance to them. (iv) the *liberty* foundation is not considered due to its very recent addition to the main theory. To address these shortcomings, other lexicons were developed; the most broadly used are the MoralStrength lexicon [1] and the extended Moral Foundations Dictionary (eMFD) [18]. The eMFD continues along with the original MFD lexicon, treating each word related to a moral as either one of two opposites, "vice" and "virtue".

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

This approach leads to contradictory findings when both a vice and virtue are present in the same text, making it harder to assign a clear polarity to it. Following different approaches, both expand the number of lemmas per foundation with more commonly used terms. MoralStrength on the other hand, inspired by the literature in sentiment analysis, provides a notion of "moral valence" that ranges along the entire spectrum of each foundation from vice to virtue, providing more fine-grained and nuanced results. Despite addressing several of the most critical shortcomings of the original lexicon, these resources do not include the *liberty* foundation.

Here we lay the groundwork for a linguistic resource that can detect the presence of the *liberty* moral dimension in people's narratives. Since this foundation was a later addition to the MFT theory, there are no initial linguistic indicators that can be used as seed words for further expanding the lexicon, a pre-requisite for the eMFD and MoralStrength approaches. To overcome this issue, we gathered data from news sources on most political ideologies and leanings in the USA political scene. More specifically, we obtained data from AllSides¹, an American news aggregator that presents different versions of similar news stories from sources of the political right, left, and centre [23] to help readers break through their filter bubble, as well as the libertarian magazine Reason². Seen through the lenses of our theoretical framework, the Moral Foundations Theory, Reason prioritises the moral value of liberty with respect to the news outlets from other political views.

We employ two data-driven approaches, the first based on word embedding similarity and the second on compositional semantics. To decide on the most efficient approach, we compare the two models qualitatively and quantitatively via two separate classification tasks. Firstly, we employ the obtained lexicons to predict whether an unseen document originated from AllSides or Reason datasets. Secondly, in an unsupervised manner, we infer the presence of the moral value of liberty in Facebook comments on Pages regarding vaccination. We conclude that the lexicon generation based on the compositional semantics approach was overall the most efficient one, and hence the one we propose as the final *LibertyMFD* lexicon.

In the age of social media, social crisis, disasters, and epidemics, drive not only affliction in the physical world, but also prompt a deluge of opinions, (mis)information, and advice to billions of internet users [24]. Here, we contribute to the research and policymaking communities with a resource that facilitates the understanding of the liberty and moral value of user-generated communications on controversial phenomena timely and at large scale.

2 THEORETICAL BACKGROUND AND RELATED LITERATURE

The original Moral Foundations Dictionary (MFD) [14] consists of a collection of lemmas assembled by experts and was typically used together with the Linguistic Inquiry and Word Count (LIWC) software [30] to estimate moral traits and to investigate differences in moral concerns between different cultural groups. Garten et al. [12] proposed the Distributed Dictionary Representations (DDR) method based on psychological dictionaries and semantic similarity to quantify the presence of moral sentiment around a given topic. Later on, the authors extended the method, incorporating demographic embeddings into the language representations [13].

In an attempt to address several of the limitations of the MFD, Araque et al. [1] proposed a data-driven generated lexicon the MoralStrength, which expanded the original MFD employing the WordNet synsets and crowdsourced annotations. Different from the MFD, where each foundation is considered a bipolar of "virtue" and "vice", MoralStrength treats each foundation as a continuum, assigning a numeric value of moral valence to each lemma that indicates the weight with which the lemma is expressing the specific value. Hopp et al. [18] developed the extended Moral Foundations Dictionary (eMFD), a lexicon which expanded the MFD based on crowdsourced annotations. Each lemma in eMFD is assigned a continuously weighted vector that expresses the probability that the lemma belongs to any of the five moral foundations.

Notably, none of the above lexicons though included the liberty moral foundation. A first attempt to derive a lexicon from assessing the presence of liberty in the text was presented by Araque et al. [2]. They considered pairs of Wikipedia Pages³ and their Conservapedia⁴ counterparts as a natural experiment. They created a series of word embeddings which were then compared through cosine similarity to a set of seed words defined by experts to generate a lexicon. Although innovative, their design comes with the obvious conceptual limitations of considering the Wikipedia project as expressing a strongly libertarian position and initiating the embeddings with a list of manually selected seed words from expert annotators. Here, we overcome these limitations by considering news outlets with a clear political orientation provided by experts in the field⁵. At the same time, our initial seed words emerge in a data-driven way from the actual linguistic behavioural patterns employed by the two sides.

3 MATERIALS AND METHODS

Our initial assumption is that libertarians are highly sensitive to the moral value of liberty (and lack thereof). Hence, their communication will contain more words related to the liberty foundation than people self-identifying with a different part of the political spectrum. We propose two entirely data-driven models for assessing the liberty foundation aiming to obtain a final *LibertyMFD* lexicon. The first model (Sect. 3.2) is based on word embeddings, and the second one (Sect. 3.3) on computational semantics.

3.1 Data Collection

To assess the linguistic patterns of the entire political spectrum of the U.S., we obtained data from the AllSides and Reason platforms. Moreover, we included the only available resource with annotations regarding the liberty moral foundation [27].

3.1.1 News Dataset. The Reason is an American libertarian monthly magazine published by the Reason Foundation. To obtain a representative sample of data expressing the libertarian point of view on multiple societal issues, we downloaded 14,319 articles from the

¹https://www.allsides.com/unbiased-balanced-news ²https://reason.com

³Link to the Wikipedia Project.

⁴Link to the Conservapedia Project.

⁵Media Bias Chart provided by AllSides

GoodIT'22, September 7-9, 2022, Limassol, Cyprus

website, ranging from June 1968 to June 2021⁶. About 50 documents were discarded, as they contained no text. Each article (averaging 1235 words in length) consists of a title, publication date, and the magazine section where it appears. The articles published after 1994 usually contain a summary, one or more keywords and the author's name.

AllSides is an American company that assesses the political bias of prominent media outlets, aligning similar news stories from sources of the political right, left, and centre in a mission to show readers news outside their filter bubble. We collected 5,584 "news roundups"⁷ published between June 2013 and April 2021. These consist of a short description of the news item, the headline and the first paragraphs (71 words on average, including headline and description, calculated on all 11,379 individual articles downloaded) of the original publication and its source. The AllSides media bias chart⁸ classifies Reason as the only libertarian-leaning source; however, they opted to assign articles originating from Reason magazine with the "right-leaning" political side. Here, we excluded all 67 articles from Reason included in the AllSides dataset.

To ensure the stability and generalisability of the approach, we split this dataset in two (Table 1). The training set, consisting of 80% of the whole *News* dataset, is used for the generation of the *liberty* lexicons. The left out set, i.e. the remaining 20% articles, will only be used for testing the quality of the generated lexicons.

3.1.2 Vaccine Dataset. To evaluate the obtained lexicon, we also employ data from Facebook Pages with a clear stance about vaccination [27]. The data include anonymised posts and comments from approximately 200 Facebook Pages collected via the Facebook API from January 2012-June 2019. Overall, there are 607, 105 comments and posts from vaccines supporting and opposing sides. Approximately 1,500 comments were randomly selected and annotated manually as to whether the liberty moral foundation is present or not in the text snippet and whether it is associated with the "virtue" (liberty) or "vice" (oppression) polarity of the dimension. It is worth mentioning that this dataset, although limited, is the only available resource containing annotations about the moral foundation of liberty. In contrast, the most commonly used benchmark dataset for analysing the moral foundations [17] does not include the liberty/oppression foundation.

3.2 Word Embedding Similarity

Inspired by the work of Turney et al. [31], the first approach to generate the *LibertyMFD* is based on word embedding similarity among vectors of both AllSides and Reason. This approach requires an initial set of seed words that represent the domains we aim at distinguishing, in our case, the liberty/oppression moral dimension from the *others*, namely all remaining moral foundations (such as authority, purity) and non-moral text. Due to the lack of any apriori linguistic information, in previous work [2] experts manually defined the list of seed words; however, slight variations in the seed word list may generate very different final lexicons.

Here, to avoid the shortcomings of an arbitrary selection, we obtain the set of seed words in a data-driven way, estimating the *frequency shifts* [11] of the lemmas between the AllSides and Reason datasets. Let the relative frequency of a word *w* in a set of documents *D* be:

$$p_{w}^{(D)} = \frac{f_{w}^{(D)}}{\sum_{w' \in W^{(D)}} f_{w'}^{(D)}}$$

where $w' \in W^{(d)}$ are the words in vocabulary set $W^{(D)}$ except for w. We compute the frequency shift in the relative frequency per word w between two different sets of documents:

$$\delta p_w = p_w^{(2)} - p_w^{(1)}$$

The corresponding seed word lists emerge as salient differences in word frequency shifts. We apply a minimum frequency threshold at 100, under which a lemma is filtered out of the seed word list. Then, we compute each word's respective vector employing the word2vec algorithm [21] with a standard parameter setting and a vector dimension equal to 100. The lexicon is then generated by estimating the cosine similarity between the word vectors obtained with the emerged seed words. If S_C is the set of seed words for the "others" orientation, and S_L the seed words for the "liberty" direction, the moral polarity of a word w_i is computed from the documents as

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

where *sim* is the cosine similarity estimated by the word embedding model. The polarity is a positive number if w_i is related to the "liberty/oppression" seed words and a negative one if the relation occurs towards the "other" seed words. For simplicity, we will refer to this model as the *WE model*.

3.3 Compositional Semantics

The second approach is based on the *Compositional Semantics* (CS) method, which has been extensively used to generate emotion lexicons [3, 29]. We adapt the domain of application to the liberty moral values; in this case, we consider the documents of Reason as discussing liberty/oppression more frequently than the documents from AllSides. The CS method projects the moral values from a document to its words, assuming that each word is associated with the moral value expressed in the documents where the word appears more frequently.

Let us consider a document-by-moral matrix M_{DM} which contains the distribution of the liberty foundation across the Reason and AllSides datasets. We compute a word-by-document matrix M_{WD} which contains the number of occurrences for each word in the vocabulary in a given document, normalised with the number of words per document. To obtain the word-by-moral matrix, we perform the following multiplication:

$M_{WM} = M_{WD} \cdot M_{DM}$

This allows us to merge words with moral values by summing the products of the weight of a word with the weight of the moral values in each document. Finally, we normalise the lexicon scores

⁶Reason.com. Retrieved May 27th, 2022. Data downloaded according to the Terms of Use from https://reason.com/archives/.

⁷https://www.allsides.com/story/admin

⁸Classification of Reason magazine.

Table 1: Overview of the datasets employed in this study. The *News* dataset was used for the generation of the lexicons (Training Set), while for the evaluation of the performance, only the Left-out Set from the *News* dataset was employed. The *Vaccine* dataset was used for the evaluation of the lexicons. For each set, the total number of instances is provided.

Dataset	Label	No. of instances (%)
News (Train Set - Used to generate the lexicons)	Liberty/Oppression (Reason) Other (AllSides)	11,470 (56%) 9,062 (44%)
News (Left-out Set - Used to evaluate the lexicons)	Liberty/Oppression (Reason) Other (AllSides)	2,849 (56%) 2,284 (44%)
Vaccine (Used to evaluate the lexicons)	Liberty/Oppression (annotations) Other	178 (11%) 1398 (89%)

(column-wise), remove over-representation issues, and scale each lemma (row-wise) to sum up to one. Previous lexicon validations showed that this is an adequate normalisation approach [3]. For simplicity, we will refer to this model as the *CS model*.

3.4 Evaluation Design

3.4.1 Lexicon Overlap Score and Coverage. We employed a series of qualitative and quantitative methods to gauge the differences between the two generated lexicons, namely the WE and the CS lexicons. For each lexicon, we estimate the average coverage per dataset as a percentage of the total number of lemmas in the dataset that also exist in the respective lexicon. Then, we assess the statistical properties of the liberty score distributions obtained from two lexicons and provide excerpts covering both polarities of each lexicon. We also provide the lexicon overlap score (LOS) as defined in [32]. The simple version of the lexicons overlaps score considers the size of both lexicons and the words they have in common. Since the simple LOS does not consider the general direction of the polarity, we also report the binary LOS.

3.4.2 Unsupervised Document Categorisation. We perform an unsupervised evaluation of the WE and CS lexicons on the Vaccine dataset in line with the evaluation approach proposed for the comparison of polarity lexicons in sentiment analysis tasks [32]. For each document, we average the liberty scores assigned to the lemmas in the lexicon under evaluation. In our case, a document is a comment on a Facebook Page post. Hence, we obtain two numeric scores for each document, describing how much the document expresses the value of liberty according to the WE and the CS lexicons, respectively. We then binarise the scores according to the median value of the respective lexicon, deciding for each document whether it expresses "liberty/oppression" moral values or not. Negation correction was not applied, as foundation polarities do not directly translate as opposites (e.g., "not libertarian" is not the same as "oppressive"). Here, in a very straightforward way, we assess the extent to which the two lexicons assign a common or a diverse label to an unseen document.

3.4.3 Supervised Classification of Liberty. To evaluate the performance of the two lexicons, we designed two supervised classification tasks. In the first one, a classifier predicts whether an unseen document from the *News* dataset was published by Reason (i.e., expresses a libertarian position) or from AllSides. The classifier is trained (and evaluated, using cross-validation) on the left-out set of articles that were not considered for generating the lexicons. Similarly, for the *Vaccine* dataset, we aim to accurately infer the labels assigned by the expert annotators on whether a post or a comment contains the liberty foundation (regardless of whether it is expressed as virtue or vice) or not. This second experiment aims at determining to what extent the lexicons help identify the presence of liberty in new datasets and domains, going beyond the limits of our naïve unsupervised approach.

We trained an SVM classifier [10] with a linear kernel for each one of the above tasks and each lexicon. The feature vectors are constructed as follows; each document is represented by a vector of equal size to the lexicon. For those lemmas in the document present in the lexicon, the vector contains the respective polarity score, otherwise zero. Since this type of representation dramatically simplifies the linguistic information present in the document, we enhance the classification design with two more experiments. We extend each vector representation with the "statistical summary" functions, namely the average, maximum, median, variance, and a peak-to-peak score of the lexicon values of that document. This offers the learning models a more complete view of the text.

As a baseline model, we train an SVM classifier using a unigram representation with a standard vocabulary size of 5,000 tokens. We report the macro-averaged F-score with 10-fold cross-validation for all experiments. We also test and report the performance of the "statistical summary" used in isolation.

4 RESULTS AND DISCUSSION

This study provides a thorough comparison of the two lexicons employing both supervised and unsupervised approaches.

4.1 Lexicon Overlap Score and Coverage.

Figure 1 depicts the distribution of moral scores for the two generated resources. The score distributions from the two lexicons are statistically different (p < 0.001), evident from both the spread and shape. The WE model has a unimodal distribution with a skewed average towards the negative values. The liberty values of the CS model instead have a right-skewed bimodal distribution with an imposing first mode and a long tail.

Another interesting characteristic is that the average value of the WE lexicon distribution is skewed to the negative values, while the first mode of the CS lexicon is 0. Intuitively, this is reasonable, as it indicates that most of the lemmas are not indicative of the *liberty* foundation. The obtained distributions follow patterns found in previous lexicons; in fact, the normal distribution of the WE lexicon is similar to that of SentiWordNet [5], while the CS lexicon has a similar distribution to the NRC Emotion lexicon [4, 25].



Figure 1: Distributions of the generated lexical resources, the word embedding similarity, WE model, on the left hand side and the Compositional Semantics, CS model on the right hand side. Dashed lines indicate the quartiles of the distributions.

Checking the coverage of the two models, we observe that the coverage of the CS lexicon is much higher across all datasets, as reported in Table 2. The simple LOS states to what extent the lexicons share a common vocabulary 77.2%, while the binary LOS is 39.5%. At the same time, if we consider the lemmas with a frequency of appearance above the average for both lexicons, we obtain 93.5% shared lemmas, with a binary LOS of 44.4%. The above finding shows that the two lexicons are more in accordance with polarity scores for the most popular lemmas.

Moreover, the two methods capture different patterns of language usage related to *liberty*. As can be seen in Table 3, "expense" and "distribution" are both highly related to liberty/oppression. Given that the lexicons are generated starting from news, these terms likely refer to economic spending and usage of taxes, but their relative order is different in the two resources.

4.2 Unsupervised Document Categorisation.

We further evaluate the WE and CS lexicon differences, respectively, in an unsupervised classification task. Here, we aim to infer whether a document, from the *Vaccine* dataset, expresses a concept relevant to the moral foundation of liberty. We evaluate the outcome comparing against the manual annotations provided by experts. Given the limited linguistic coverage observed already in Table 2, the performance is expected to be relatively low. Nonetheless, this evaluation offers the first indication of the performance of the two models. An average macro F-score of 45.32 for the CS lexicon, against 23.46 for the WE one, gives a clear benefit to the CS model.

4.3 Supervised Classification of Liberty.

We evaluated the obtained lexicons as per their predictive power in two independent text classification tasks on both the *News* and the *Vaccine* datasets. As described in Sect. 3.4, we compare the classification performance of the two proposed lexicon generation methods, including a baseline for reference.

Table 4 reports the average macro F-score and the standard deviation obtained on a 10-fold cross-validation setup. Worth mentioning is the big difference in performance between the two datasets; this should not surprise, as the *Vaccine* dataset is significantly smaller, is high class imbalanced, and is focused on a domain that is very different from that of news. The baseline model reaches a high performance in both datasets; in the case of the *News* dataset, both WE and CS generated lexicons that improve the performance over the baseline. For the *Vaccine* dataset instead, only the CS models provides a slightly better average. The difficulty of the WE model to adapt to a new dataset like the *Vaccine* may be related to the limited and fixed amount of information that can be encoded [22].

The ability of the CS model to generalise better in the *Vaccine* dataset prediction, outperforming the baseline and the WE model, demonstrates the importance of having extended coverage. Moreover, the long tail distribution of the moral scores of the CS model help distinguish more efficiently documents of a lesser-known domain not just in terms of vocabulary but also of writing style. Such ability can be parallelized with the human semantic ability to compose lexical meanings to form a potentially unlimited number of complex linguistic expressions.

Interestingly, the introduction of basic statistical summary features alone is not informative of the liberty foundation; however, when added to the main feature vectors, it gives the classifier a complete view of the presence of the *liberty* foundation in text. To conclude, both the CS and WE model, with or without the statistical summary, can easily outperform a unigram baseline, showing the value of both lexicons in the domain of news. When considering a novel, unseen domain, our findings suggest that the compositional semantics approach (CS model) provides with a more efficient, domain-independent lexicon; hereafter we will refer to the output of the CS model as the *LibertyMFD lexicon*.

5 CONCLUSIONS

The foundation of Liberty centres on individualism and independence and can have multiple facets ranging from social justice advocacy for immigrants to resistance to oppressive regulations with prosocial character [8]. Literature on computational social science highlighted the importance of the liberty foundation in understanding peoples' fundamental rationale towards critical social issues such as attitudes towards vaccination [20], charitable giving [26], familial coercion and human rights [9].

Although the impact of the liberty foundation goes well beyond these topics, and despite the increasing body of research around the main five dimensions of the moral foundation's theory, there is still no linguistic resource assessing the liberty foundation. Here, we address precisely this shortcoming. Having as a cornerstone assumption that libertarian journals will express more intensively the value of liberty than the others, we considered the online news sources as a natural experiment to train and test our models.

Words in the *LibertyMFD* were selected after extensive experimentation on two different data-driven lexicon generation approaches. Further, we assessed the performance of the lexicon on a manually annotated dataset of user-generated text regarding the Table 2: Average coverage of the two lexicons over all datasets. Coverage is measured with the number of words from each document that appear in each lexicon.

Dataset	WE lexicon coverage	CS lexicon coverage
News Train Set	82.0	96.8
News Left-out Set	82.0	96.6
Vaccine	77.4	90.9

Table 3: Example lemmas from the WE and CS lexicons. We indicatively present some of the lemmas with the highest and lowest polarity score for the two lexicons.

Model	Liberty Polarity	Lemmas
WE	High	$distribution, knowledge, we alth, control, function, productive, expense, values, producing, \ skills$
WE	Low	Georgia, senator, Harris, Iowa, governor, congress man, Ohio, impeachment, scandal and the senator of the sen
CS	High	corporations, coercion, graduate, develop, incomes, budgets, expense, scientists, advice
CS	Low	financial, Americans, enemy, responsible, suspect, citizens, racial, bias, suburban, refuse and the suburban and the suburb

Table 4: Macro averaged f-score in the *News* left-out set and *Vaccine* datasets. WE and CS models are the lexicons generated with the embedding similarity and compositional semantics approaches. In **bold**, we indicate the best performing model for each dataset.

Model	News Dataset	Vaccine Dataset
Unigram (baseline)	94.29 (±0.91)	76.07 (± 5.13)
WE model	97.54 (± 0.82)	73.85 (± 5.29)
Stat. summary (WE)	82.35 (± 3.52)	37.62 (± 3.26)
WE model + Stat. summary	98.42 (± 0.44)	74.21 (± 5.03)
CS model	97.40 (± 0.73)	76.70 (± 4.80)
Stat. summary (CS)	82.60 (± 1.76)	44.41 (± 3.32)
CS model + Stat. summary	98.27 (±0.57)	77.51 (± 5.08)

issue of vaccine hesitancy, one of the topics where liberty is known to be a critical factor [7]. Our findings showed that the lexicon generated based on the compositional semantics approach had a broader coverage even in an out of domain dataset. At the same time, the polarity scores assigned to each lemma appeared to be more predictive of the liberty moral foundation in unseen documents. As a future improvement, we aim to refine the lexicon polarity scores to achieve a more robust association with the two polarities of the liberty foundation.

The proposed resource is generated for the English language; although the framework is directly expandable to other languages it would require retraining of the model on local language datasets to efficiently grasp the linguistic particularities of each language. We recommend practitioners who aim to utilise the resource in languages different than English simply via automatic text translation, to interpret the results with caution, since the moral expression of the words may differ across languages.

We contribute to the research and policymakers communities with an open-source lexicon that can be particularly helpful to understanding better moral judgments, dispositions, and attitudes formation from spontaneous digital data. Given the penetration of social media in our communication and information ecosystem, such resources can be employed to analyse large scale usergenerated texts nowcasting peoples' attitudes towards fast-evolving social phenomena with direct implications for the humanitarian sector.

ACKNOWLEDGMENTS

The work of Oscar Araque was 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

- 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] Oscar Araque, Lorenzo Gatti, and Kyriaki Kalimeri. 2021. The Language of Liberty: A Preliminary Study. In *Companion Proceedings of the Web Conference* 2021 (Ljubljana, Slovenia) (WWW '21). Association for Computing Machinery, New York, NY, USA, 623?626. https://doi.org/10.1145/3442442.3452351
- [3] O. Araque, L. Gatti, J. Staiano, and M. Guerini. 2019. DepecheMood++: a Bilingual Emotion Lexicon Built Through Simple Yet Powerful Techniques. *IEEE Transac*tions on Affective Computing (2019), 1–1. https://doi.org/10.1109/TAFFC.2019. 2934444
- [4] Oscar Araque and Carlos A. Iglesias. 2020. An Approach for Radicalization Detection Based on Emotion Signals and Semantic Similarity. *IEEE Access* 8 (2020), 17877–17891. https://doi.org/10.1109/ACCESS.2020.2967219
- [5] Stefano Baccianella, Andrea Esuli, and Fabrizio Sebastiani. 2010. SentiWordNet 3.0: An Enhanced Lexical Resource for Sentiment Analysis and Opinion Mining. In Proceedings of the Seventh International Conference on Language Resources and

LibertyMFD: A Lexicon to Assess the Moral Foundation of Liberty.

GoodIT'22, September 7-9, 2022, Limassol, Cyprus

Evaluation (LREC'10). European Language Resources Association (ELRA), Valletta, Malta. http://www.lrec-conf.org/proceedings/lrec2010/pdf/769_Paper.pdf

- [6] Randy E Barnett. 2014. The structure of liberty: Justice and the rule of law. OUP Oxford.
- [7] Lorenzo Betti, Gianmarco De Francisci Morales, Laetitia Gauvin, Kyriaki Kalimeri, Yelena Mejova, Daniela Paolotti, and Michele Starnini. 2021. Detecting adherence to the recommended childhood vaccination schedule from user-generated content in a US parenting forum. *PLoS computational biology* 17, 4 (2021), e1008919.
- [8] Jason Brennan. 2012. Libertarianism: What everyone needs to know. Oxford University Press.
- [9] Scott Clifford, Vijeth Iyengar, Roberto Cabeza, and Walter Sinnott-Armstrong. 2015. Moral foundations vignettes: A standardized stimulus database of scenarios based on moral foundations theory. *Behavior research methods* 47, 4 (2015), 1178–1198.
- [10] Rong-En Fan, Kai-Wei Chang, Cho-Jui Hsieh, Xiang-Rui Wang, and Chih-Jen Lin. 2008. LIBLINEAR: A library for large linear classification. *the Journal of machine Learning research* 9 (2008), 1871–1874.
- [11] Ryan J Gallagher, Morgan R Frank, Lewis Mitchell, Aaron J Schwartz, Andrew J Reagan, Christopher M Danforth, and Peter Sheridan Dodds. 2021. Generalized word shift graphs: a method for visualizing and explaining pairwise comparisons between texts. *EPJ Data Science* 10, 1 (2021), 4. https://doi.org/10.1140/epjds/ s13688-021-00260-3
- [12] Justin Garten, Joe Hoover, Kate M Johnson, Reihane Boghrati, Carol Iskiwitch, and Morteza Dehghani. 2018. Dictionaries and distributions: Combining expert knowledge and large scale textual data content analysis. *Behavior Research Methods* 50, 1 (2018), 344–361.
- [13] Justin Garten, Brendan Kennedy, Joe Hoover, Kenji Sagae, and Morteza Dehghani. 2019. Incorporating Demographic Embeddings Into Language Understanding. *Cognitive science* 43, 1 (2019).
- [14] 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. https://doi.org/10.1037/a0015141
- [15] Jonathan Haidt. 2012. The righteous mind: Why good people are divided by politics and religion. Vintage.
- [16] Jonathan Haidt and Craig Joseph. 2004. Intuitive ethics: How innately prepared intuitions generate culturally variable virtues. *Daedalus* 133, 4 (2004), 55–66.
- [17] Joe Hoover, Gwenyth Portillo-Wightman, Leigh Yeh, Shreya Havaldar, Aida Mostafazadeh Davani, Ying Lin, Brendan Kennedy, Mohammad Atari, Zahra Kamel, Madelyn Mendlen, et al. 2020. Moral Foundations Twitter Corpus: A collection of 35k tweets annotated for moral sentiment. *Social Psychological and Personality Science* 11, 8 (2020), 1057–1071.
- [18] Frederic R Hopp, Jacob T Fisher, Devin Cornell, Richard Huskey, and René Weber. 2021. The extended Moral Foundations Dictionary (eMFD): Development and applications of a crowd-sourced approach to extracting moral intuitions from

text. Behavior research methods 53, 1 (2021), 232-246.

- [19] 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.
- [20] 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.
- [21] Quoc Le and Tomas Mikolov. 2014. Distributed representations of sentences and documents. In International Conference on Machine Learning. 1188-1196.
- [22] Alessandro Lenci. 2018. Distributional models of word meaning. Annual review of Linguistics 4 (2018), 151–171.
- [23] AllSides LLC. 2022. AllSides Media Bias Ratings? https://www.allsides.com/ media-bias/media-bias-ratings
- [24] Yelena Mejova and Kyriaki Kalimeri. 2020. COVID-19 on Facebook ads: competing agendas around a public health crisis. In Proceedings of the 3rd ACM SIGCAS Conference on Computing and Sustainable Societies. 22–31.
- [25] Saif M. Mohammad and Švetlana Kiritchenko. 2015. Using Hashtags to Capture Fine Emotion Categories from Tweets. *Computational Intelligence* 31, 2 (2015), 301–326. https://doi.org/10.1111/coin.12024
- [26] Artur Nilsson, Arvid Erlandsson, and Daniel Västfjäll. 2020. Moral foundations theory and the psychology of charitable giving. *European Journal of Personality* 34, 3 (2020), 431–447.
- [27] Maria Florencia Prado, Victoria Perez Bustos, Kyriaki Kalimeri, and Mariano G. Beiro. 2022. Narratives Around the Vaccination Discourse on the Facebook Platform. arXiv preprint (2022).
- [28] 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).
- [29] Jacopo Staiano and Marco Guerini. 2014. Depeche Mood: a Lexicon for Emotion Analysis from Crowd Annotated News. In Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers). Association for Computational Linguistics, Baltimore, Maryland, 427–433. https://doi.org/10.3115/v1/P14-2070
- [30] Yla R Tausczik and James W Pennebaker. 2010. The psychological meaning of words: LIWC and computerized text analysis methods. *Journal of language and* social psychology 29, 1 (2010), 24–54.
- [31] 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.
- [32] Felix JM Welter. 2020. Quantitative comparison of polarity lexicons in sentiment analysis tasks: Using a lexicon overlap score for similarity measurement between lexicons. SKILL 2020-Studierendenkonferenz Informatik (2020).