

Studying Moral-based Differences in the Framing of Political Tweets

Markus Reiter-Haas¹, Simone Kopeinik², Elisabeth Lex¹

¹Graz University of Technology

²Know-Center GmbH

reiter-haas@tugraz.at, skopeinik@know-center.at, elisabeth.lex@tugraz.at

Abstract

In this paper, we study the moral framing of political content on Twitter. Specifically, we examine differences in moral framing in two datasets: (i) tweets from US-based politicians annotated with political affiliation and (ii) COVID-19 related tweets in German from followers of the leaders of the five major Austrian political parties. Our research is based on recent work that introduces an unsupervised approach to extract framing bias and intensity in news using a dictionary of moral virtues and vices. In this paper, we use a more extensive dictionary and adapt it to German-language tweets. Overall, in both datasets, we observe a moral framing that is congruent with the public perception of the political parties. In the US dataset, democrats have a tendency to frame tweets in terms of care, while loyalty is a characteristic frame for republicans. In the Austrian dataset, we find that the followers of the governing conservative party emphasize care, which is a key message and moral frame in the party's COVID-19 campaign slogan. Our work complements existing studies on moral framing in social media. Also, our empirical findings provide novel insights into moral-based framing on COVID-19 in Austria.

Introduction

Politicians and political campaigns increasingly use social media to connect and communicate with potential voters (Graham et al. 2013). The effectiveness of such communication is influenced by how the message is *framed* (Kusmanoff et al. 2020). Framing corresponds to the act of changing the formulation of a problem to affect the choices of people (Tversky and Kahneman 1981).

Recently, several related works focus on the characterization of frames: Walter and Ophir (2019) use topic modeling and network analysis to identify frames in news. Shurafa, Darwish, and Zaghouani (2020) categorize political discussions related to COVID-19 in Twitter into either *blame frames* or *support frames*. Wicke and Bolognesi (2020) find that the discourse around COVID-19 on Twitter is framed using war-related terminology.


In our work, we aim to study differences in moral-based framing in content created by members and followers of opposing political parties on Twitter. We base our

approach on the work of Makhberian et al. (2020), who have recently introduced an unsupervised, embedding-based method to characterize *moral frames* in text. Moral frames are frames that emphasize specific moral virtues and vices, such as care or harm. The approach of Makhberian et al. is grounded in the Moral Foundation Theory from the social sciences (Haidt and Joseph 2004), which defines five basic moral foundations and their associated virtues and vices (Haidt and Joseph 2007). Based on the theory, several moral foundation dictionaries (Graham, Haidt, and Nosek 2009; Frimer et al. 2017) have been developed that contain prototypical words for each moral foundation.

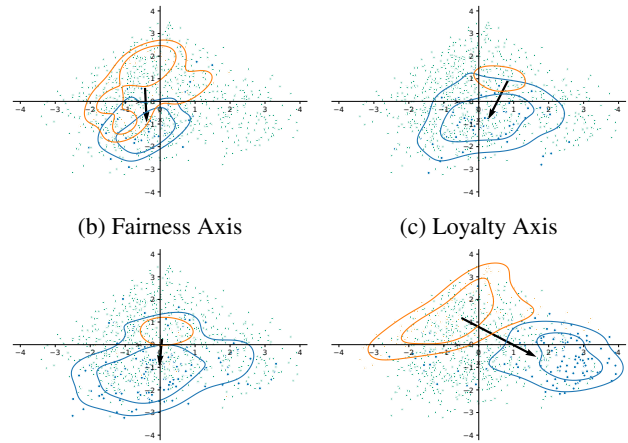
In this paper, we employ a similar approach to Makhberian et al. However, while they utilize the moral foundation dictionary by Graham, Haidt, and Nosek (2009), for our experiments, we use the more recent and more extensive dictionary by Frimer et al. (2017). Besides, we translate the content of that moral foundation dictionary to German using a list of sample translations of positive and negative valence words (Weichselbaum, Leder, and Ansoorge 2018) and two sets of word embeddings, i.e., one for English and one for German (Grave et al. 2018).

For our study, we create two Twitter datasets. The first dataset contains tweets from US-politicians annotated with political affiliation (democrats vs. republicans). The second dataset contains COVID-19-related tweets from the followers of the five major Austrian political parties' leaders in the German language. From the tweets, we extract moral frames corresponding to the five moral foundations, their frame bias, i.e., the emphasis towards either virtue or vice, and frame intensity, i.e., the extent to which a frame is used. To study the prevalence of moral frames, we train a logistic regression classifier to predict party affiliation and investigate its coefficients.

In both datasets, we observe a moral framing congruent with the public perception of the political parties. In the US dataset, high frame intensity on care and fairness are predictors for democrats, while high frame intensity on loyalty and sanctity characterise republicans. In the Austrian dataset, we find a frame bias toward care in the COVID-19-related tweets of the conservative political party leader's followers. We attribute this to the followers' adoption of the conservative COVID-19 slogan's moral framing that stresses caring.



(a) Care Axis



(b) Fairness Axis

(c) Loyalty Axis

(d) Authority Axis

(e) Sanctity Axis

Figure 1: Axis of the five moral foundations. Each axis is created by the centroid of words assigned to virtues and the centroid of words for vices and surrounded by moral words associated with the other axes. The black arrow goes from the vices’ centroid to the virtues’ centroid and describes the axes. The high-dimensional space is reduced with Principal Component Analysis (PCA). All the axis point approximately in the same direction, which indicates that virtues are more similar to other virtues than to their corresponding vices, and vice versa. A kernel density estimation of the underlying point cloud is used for the colored contours.

Characterization of Moral-based Framing

In the following, we describe our approach to investigate moral-based differences in the framing of tweets.

Capturing Moral Frames

In our work, similar to Mokherian et al. (2020), we capture moral frames by combining the FrameAxis approach introduced in Kwak, An, and Ahn (2020) with a dictionary of moral values. FrameAxis enables the quantification of framing of a particular text using *semantic axes* (Kwak, An, and Ahn 2020). It is built upon the SemAxis approach (An, Kwak, and Ahn 2018), which defines semantic axes by the difference of opposing word pairs using their word embeddings in the vector space. FrameAxis learns in an unsupervised way by estimating the contribution of each word towards the target axis. The contribution per word is defined as the cosine similarity between its word embedding and the target axis in the vector space. For all contributions of every word in a given document, we calculate the *frame bias* and *frame intensity* of a moral frame. The frame bias corresponds to the mean of the contributions and the frame intensity to the variance of the contributions in relation to the baseline frame bias of the corpus. The latter denotes the mean of frame biases over the whole corpus.

As a dictionary of moral values, we use the Moral Foundation Dictionary version 2 (MFD-2) (Frimer et al. 2017). It is an extension of a moral values dictionary developed by Graham, Haidt, and Nosek (2009) and consists of prototypical words to *moral foundations*. Moral foundations are described in the moral foundation theory (MFT) as factors that guide emotional and ethical reactions to various social situations. MFT describes five foundations in the form of virtues and vices: (i) care/harm, i.e., the dislike for others’ suffering, (ii) fairness/cheating, i.e., dislike of cheating, (iii)

loyalty/betrayal, i.e., loyalty, (iv) authority/subversion, i.e., respect for authority, and (v) sanctity/degradation, i.e., concerns with purity (Mokherian et al. 2020).

The Moral Foundation Dictionary MFD-2 assigns words to virtues and vices. As virtues and vices are opposing moral values, we use them as poles to create *moral frame axes*. Then, for each pole, we associate its words with word embeddings, i.e., the 300-dimensional GloVe representation (Pennington, Socher, and Manning 2014) trained on 840 billion tokens and calculate their centroids for virtues and vices. Each pair of virtue and vice centroid forms a semantic axis, i.e., *moral frame axis*, that we use for FrameAxis instead of individual words. For each axis, we extract the frame biases and intensities per tweet by aggregating its words’ contributions (i.e., the cosine similarity with the axis) towards the corresponding moral value. Please note that we name axes using the name of the morals’ virtues in the remainder of this paper, e.g., the care axis.

Validation of Moral Frame Axes

We define four properties of the word embedding space to investigate the validity of the moral frame axes. *P1: All axes should be close to the zero point.* Note that each axis is dividing a moral space into a positive and a negative part. *P1* prohibits the dominance of one pole (i.e., the pole closer to the zero point) that could be caused by an association of an overwhelming majority of words. *P2: The words associated with a pole should be semantically closer to each other than to words of the opposite pole.* If words are added to or removed from an axis, then *P2* ensures its stability. *P3: The orientation of axes should not oppose.* Adherence to *P3* allows the axes to be combinable and form a meta-axis for virtues and vices, e.g., care virtues are closer to fairness virtues than fairness vices. *P4: The orientation of axes should differ in the*

	Precision	Recall	F1-score	Accuracy
Reproduced with MFD-2				
Care	0.828	0.827	0.827	0.827
Fairness	0.729	0.728	0.728	0.728
Authority	0.754	0.754	0.754	0.754
Loyalty	0.895	0.889	0.891	0.889
Sanctity	0.881	0.880	0.880	0.880
Original with MFD-1				
Care	0.746	0.768	0.734	0.768
Fairness	0.662	0.774	0.681	0.774
Authority	0.808	0.875	0.817	0.875
Loyalty	0.802	0.873	0.816	0.873
Sanctity	0.910	0.935	0.908	0.935

Table 1: Results of classification of moral frames on the annotated Twitter corpus. The performance of MFD-2 is compared with the results from Mokherberian et al. (2020).

hyperspace. We expect the axes to be orthogonal to a certain degree. A violation of $P4$, i.e., two axes are pointing directly in the same direction, suggests that these axes likely relate to the same concept and could be combined.

A visual analysis of the moral frame axes (see Figure 1) shows the first two principal components of word embeddings using probabilistic Principal Component Analysis (PCA), moral frame axes, and up to three density regions for virtues and vices using a kernel density estimation, which has a lowest level threshold of 33%. Results indicate all the four properties hold, e.g., all the axes point in the same direction. Due to some ambiguous words, there is some overlap in the projected point clouds (e.g., unharmed). Furthermore, some words (e.g., wounds) belong to both poles, i.e., virtue and vice in the dictionary. In addition to the visual depiction, we also perform the validation numerically¹.

Validation on Annotated Tweets

To validate our approach, we perform classification of moral frames similar to Mokherberian et al. (2020) on the Twitter dataset provided by Hoover et al. (2020), which is annotated with virtues and vices. We conduct our experiments using a logistic regression classifier with the MFD-2 dictionary. Table 1 contains the results of this experiment, and a comparison of our results with the results of Mokherberian et al. (2020). While we observe similar results as Mokherberian et al., we find that the use of MFD-2 improves the F1-score on care, fairness, and loyalty, but performs worse on authority and sanctity. In terms of accuracy, we achieve a higher performance on care and loyalty using MFD-2, but a lower performance on fairness, authority, and sanctity. We conclude that the classifier accurately captures moral frames in tweets.

Experiments and Results

We perform experiments on two datasets: firstly, in tweets in the English language created by US-based politicians, which

¹We provide the code, plots and examples of this research at: <https://github.com/socialcomplab/icwsm21-framing>

we gathered based on the Twitter user list provided by Barberá et al. (2015), and secondly, in German-speaking tweets that contain COVID-19 related content created by followers of the leaders of the five major Austrian parties. Our selection of datasets is motivated by their differences in contextual attributes, concretely their language (i.e., English vs. German), topics (i.e., various topics vs. COVID-19-related topics), account type (i.e., politicians vs. followers of top politicians), and distribution of political parties (i.e., two-party system in the US vs. multi-party system in Austria).

Datasets

For the US Twitter dataset, we collect the most recent tweets of democrats and republicans using the party-associated Twitter handles (Barberá et al. 2015). The resulting dataset consists of 1,388,198 tweets, i.e., 704,392 tweets from 243 democratic (D) and 683,806 from 252 republican (R) accounts. We label the tweets according to the account owner’s party affiliation.

For the Austrian Twitter dataset, we manually extract the Twitter handles of the five major Austrian parties’ lead politicians, i.e., @BMeinl for the liberal party (NEOS), @WKogler for the green party (Greens), @norbertghofer for the national-focused freedom party (FPÖ), @rendiwagner for the social-democratic party (SPÖ), and @sebastiankurz for the conservative people’s party (ÖVP). Then, we collect the most recent tweets of followers and labeled each tweet of the follower with the politician they follow. To avoid mutual labels, we restrict our collection to users that follow only one of the five accounts. Besides, we only consider tweets that contain COVID-19 related hashtags (e.g., #Corona). This results in a collection of 22,205 tweets, i.e., 17,230 tweets labeled with @sebastiankurz, 2,090 labeled with @WKogler, 1,164 labeled with @rendiwagner, 901 labeled with @BMeinl, and 820 labeled with @norbertghofer.

We normalize the tweets in both datasets and (i.e., lowercase, removing URLs, punctuation), remove stopwords, and apply tokenization before extracting the frame biases and intensities for training a logistic regression classifier¹.

Moral-based Framing in US-based Tweets

We group the tweets by parties and report the coefficients of the logistic regression classifier in Table 2a. The frame biases do not deviate considerably and, in general, share the same direction on all moral frames but on authority, which is positive for republicans and negative for democrats. We observe that democrats score higher in fairness and lower in sanctity, whereas republicans score higher in the frame bias for care and exhibit a high negative score in the frame bias for loyalty. Concerning the frame intensities, we observe opposing and more distinct results. The frame intensity for care is much higher for democrats, and conversely, the frame intensity for loyalty is higher for republicans. The frame intensities on fairness and sanctity agree with their corresponding frame biases, i.e., fairness has a higher frame intensity for democrats, while sanctity has a higher frame intensity for republicans. We find that our observations are congruent with Graham, Haidt, and Nosek (2009), i.e., liberals are predominantly associated with care and fairness.

Moral Frames		D	R	@BMeinl	@WKogler	@norbertghofer	@rendiwagner	@sebastiankurz
Bias	Care	2.505	3.791	-0.788	-0.463	-4.682	2.141	6.931
	Fairness	2.130	1.115	-1.375	12.494	-9.165	2.408	-13.408
	Authority	-0.385	2.343	-0.035	-1.078	-0.366	2.561	-0.291
	Loyalty	-1.419	-5.269	-0.078	-0.998	2.627	-7.987	0.649
	Sanctity	0.476	2.102	-3.673	-0.457	15.645	-0.145	0.458
Intensity	Care	9.701	-0.634	0.077	-0.039	0.001	-0.010	-0.011
	Fairness	4.376	-8.154	0.072	-0.046	0.038	-0.066	0.018
	Authority	-3.453	-6.329	0.008	0.002	-0.009	0.007	-0.003
	Loyalty	0.166	9.956	0.003	0.003	-0.003	0.001	-0.009
	Sanctity	-2.967	3.261	0.104	-0.008	-0.023	-0.020	-0.034

(a) Moral frames in US politics. Democrats (D) and republicans (R) differ most in terms of frame intensities (in bold).

(b) Moral frames in Austrian politics. Frame biases are distinct between the followers of the party leaders, whereas intensities are very small in comparison. Minimum and maximum of frame biases per moral are in bold. Frame bias in fairness exhibits the greatest difference.

Table 2: Reported results correspond to the coefficients of the logistic regression classifier.

Moral-based Framing in Austrian-based Tweets

To investigate differences in moral-based framing in the Austrian Twitter dataset, we first translate the content of the MFD-2 dictionary. To that end, we use a list of sample translations of positive and negative valence words (Weichselbaum, Leder, and Ansoerge 2018) and two sets of word embeddings, i.e., one for English and one for German (Grave et al. 2018). Using a translation matrix estimated from the valance word translations, we translate the words of the MFD-2 to similar words in German in terms of their word embeddings. We see that the top words seem to be congruent with the moral values, e.g., top translation of authority being *Befehl* – *command*, but also observe words of opposite moral values in their vicinity, e.g., harm having *Schadenfreude* – *malicious joy* as the second, and *Freude* – *joy* and third nearest neighbor. Such inconsistencies are expected since we previously established that some words are neither clearly associated with virtues nor vices.

Then, we group the tweets by followers of Austrian party leaders and report the coefficients of the logistic regression classifier in Table 2b. We find substantial differences in frame biases between the tweets of the groups, but not in their frame intensities. The reported frame biases reaffirm the parties’ public perception, with fairness having a stronger association with left parties (with @WKogler followers being the highest), while sanctity is predominantly associating with right parties (i.e., the highest for @norbertghofer followers). Noteworthy, the followers of @sebastiankurz have the lowest association with fairness, which might indicate a contention point between the viewpoints of the governing coalition, i.e., the ÖVP (@sebastiankurz) and Greens (@WKogler). Moreover, the results show that @sebastiankurz followers are mostly associated with care, a moral frame that is prevalent in the government’s COVID-19 information campaign through the slogan “*Schau auf dich - schau auf mich*”, which translates to “*take care of you - take care of me*”. Followers of @rendiwagner, who is also a scientist and epidemiologist, are associated with authority. We suspect that is the result of her emphasizing to listen to doctors and experts. For followers of @BMeinl, all frame biases

are negative, which we relate to the party being an opposition party arguing against government COVID-19 policies. In summary, we find differences in the moral framing of the tweets on COVID-19 of the followers of the party leaders that reflect the ideology and messages of the corresponding political parties.

Conclusion

In conclusion, our experimental results show that the moral framing in the tweets of US-based politicians and the tweets of the followers of Austrian politicians is congruent with the public perception of the political parties. In the tweets from US-based politicians, we find that democrats are associated with high frame intensity in care and fairness, whereas high frame intensity in loyalty and sanctity is associated with republicans. In the tweets from followers of the five major Austrian parties’ leaders, we find that high frame bias in fairness is mostly associated with followers of the green party’s leader, while high frame bias in sanctity predominantly indicates followers of the freedom party’s leader. Besides, we find that followers of the ruling conservative party’s leader have a notable frame bias towards care in the case of COVID-19-related tweets. We attribute this to the followers’ adoption of the framing of the conservative COVID-19 slogan that stresses caring. From a methodological perspective, our experiments show that the use of the extended moral foundations dictionary MFD-2 increases the accuracy of moral frame characterization.

We recognize several limitations of our work: our analysis is restricted to two specific political Twitter datasets. We chose these datasets, as the interpretation of results requires the researchers’ domain understanding and language skills. Through making a validity analysis of the approach, we aimed to mitigate the potential impact of constraints. Also, since we did not filter out retweets, 63 tweets in the Austrian dataset are from the political party leaders.

For future work, we aim to research the interplay of frame bias and intensity in more detail. We will also study how followers engage with moral frames shared by politicians and if they are more prevalent in retweets or comments.

References

- An, J.; Kwak, H.; and Ahn, Y.-Y. 2018. SemAxis: A Lightweight Framework to Characterize Domain-Specific Word Semantics Beyond Sentiment. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, 2450–2461.
- Barberá, P.; Jost, J. T.; Nagler, J.; Tucker, J. A.; and Bonneau, R. 2015. Tweeting from left to right: Is online political communication more than an echo chamber? *Psychological science* 26(10): 1531–1542.
- Frimer, J.; Haidt, J.; Graham, J.; Dehghani, M.; and Boghrati, R. 2017. Moral foundations dictionaries for linguistic analyses, 2.0. *Unpublished Manuscript*. Retrieved from: www.jeremyfrimer.com/uploads/2/1/2/7/21278832/summary.pdf.
- Graham, J.; Haidt, J.; and Nosek, B. A. 2009. Liberals and conservatives rely on different sets of moral foundations. *Journal of personality and social psychology* 96(5): 1029.
- Graham, T.; Broersma, M.; Hazelhoff, K.; and Van't Haar, G. 2013. Between broadcasting political messages and interacting with voters: The use of Twitter during the 2010 UK general election campaign. *Information, communication & society* 16(5): 692–716.
- Grave, E.; Bojanowski, P.; Gupta, P.; Joulin, A.; and Mikolov, T. 2018. Learning Word Vectors for 157 Languages. In *Proceedings of the International Conference on Language Resources and Evaluation (LREC 2018)*.
- Haidt, J.; and Joseph, C. 2004. Intuitive ethics: How innately prepared intuitions generate culturally variable virtues. *Daedalus* 133(4): 55–66.
- Haidt, J.; and Joseph, C. 2007. The moral mind: How five sets of innate intuitions guide the development of many culture-specific virtues, and perhaps even modules. *The innate mind* 3: 367–391.
- Hoover, J.; Portillo-Wightman, G.; Yeh, L.; Havaldar, S.; Davani, A. M.; Lin, Y.; Kennedy, B.; Atari, M.; Kamel, Z.; Mendlen, M.; et al. 2020. Moral Foundations Twitter Corpus: A collection of 35k tweets annotated for moral sentiment. *Social Psychological and Personality Science* 11(8): 1057–1071.
- Kusmanoff, A. M.; Fidler, F.; Gordon, A.; Garrard, G. E.; and Bekessy, S. A. 2020. Five lessons to guide more effective biodiversity conservation message framing. *Conservation Biology* 34(5): 1131–1141.
- Kwak, H.; An, J.; and Ahn, Y.-Y. 2020. FrameAxis: Characterizing Framing Bias and Intensity with Word Embedding. *arXiv preprint arXiv:2002.08608*.
- Mokherian, N.; Abeliuk, A.; Cummings, P.; and Lerman, K. 2020. Moral Framing and Ideological Bias of News. In *International Conference on Social Informatics*, 206–219. Springer.
- Pennington, J.; Socher, R.; and Manning, C. D. 2014. Glove: Global vectors for word representation. In *Proceedings of the 2014 conference on empirical methods in natural language processing (EMNLP)*, 1532–1543.
- Shurafa, C.; Darwish, K.; and Zaghouni, W. 2020. Political Framing: US COVID19 Blame Game. In *International Conference on Social Informatics*, 333–351. Springer.
- Tversky, A.; and Kahneman, D. 1981. The framing of decisions and the psychology of choice. *Science* 211(4481): 453–458.
- Walter, D.; and Ophir, Y. 2019. News frame analysis: an inductive mixed-method computational approach. *Communication Methods and Measures* 13(4): 248–266.
- Weichselbaum, H.; Leder, H.; and Ansorge, U. 2018. Implicit and explicit evaluation of visual symmetry as a function of art expertise. *i-Perception* 9(2): 2041669518761464.
- Wicke, P.; and Bolognesi, M. M. 2020. Framing COVID-19: How we conceptualize and discuss the pandemic on Twitter. *PLoS ONE* 15(9): e0240010.