



# MoralBERT: A Fine-Tuned Language Model for Capturing Moral Values in Social Discussions

Vjosa Preniqi\*  
Queen Mary University of London  
London, UK  
v.preniqi@qmul.ac.uk

Iacopo Ghinassi  
Queen Mary University of London  
London, UK  
i.ghinassi@qmul.ac.uk

Julia Ive  
Queen Mary University of London  
London, UK  
j.ive@qmul.ac.uk

Charalampos Saitis  
Queen Mary University of London  
London, UK  
c.saitis@qmul.ac.uk

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

## ABSTRACT

Moral values play a fundamental role in how we evaluate information, make decisions, and form judgements around important social issues. Controversial topics, including vaccination, abortion, racism, and sexual orientation, often elicit opinions and attitudes that are not solely based on evidence but rather reflect moral worldviews. Recent advances in Natural Language Processing (NLP) show that moral values can be gauged in human-generated textual content. Building on the Moral Foundations Theory (MFT), this paper introduces MoralBERT, a range of language representation models fine-tuned to capture moral sentiment in social discourse. We describe a framework for both aggregated and domain-adversarial training on multiple heterogeneous MFT human-annotated datasets sourced from Twitter (now X), Reddit, and Facebook that broaden textual content diversity in terms of social media audience interests, content presentation and style, and spreading patterns. We show that the proposed framework achieves an average F1 score that is between 11% and 32% higher than lexicon-based approaches, Word2Vec embeddings, and zero-shot classification with large language models such as GPT-4 for in-domain inference. Domain-adversarial training yields better out-of domain predictions than aggregate training while achieving comparable performance to zero-shot learning. Our approach contributes to annotation-free and effective morality learning, and provides useful insights towards a more comprehensive understanding of moral narratives in controversial social debates using NLP.

## CCS CONCEPTS

• **Computing methodologies** → **Machine learning**; **Natural language processing**; • **Human-centered computing** → **Social media**;

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 the author(s) 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](mailto:permissions@acm.org).

*GoodIT '24, September 04–06, 2024, Bremen, Germany*

© 2024 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 979-8-4007-1094-0/24/09

<https://doi.org/10.1145/3677525.3678694>

## KEYWORDS

Moral values, Social Media, Language use

### ACM Reference Format:

Vjosa Preniqi, Iacopo Ghinassi, Julia Ive, Charalampos Saitis, and Kyriaki Kalimeri. 2024. MoralBERT: A Fine-Tuned Language Model for Capturing Moral Values in Social Discussions. In *International Conference on Information Technology for Social Good (GoodIT '24)*, September 04–06, 2024, Bremen, Germany. ACM, New York, NY, USA, 10 pages. <https://doi.org/10.1145/3677525.3678694>

## 1 INTRODUCTION

Language is not merely a tool for communication, but a reflection of a plethora of intricate psychological constructs. The words and phrases people use can reveal underlying emotions [56], personality traits [66], and even moral values [24]. The latter occupy a salient position, significantly influencing stance taking on contentious social issues such as vaccine hesitancy [43], civil unrest [55], but also personal taste such as the type of music we like to listen [61, 62]. Here, we aim to improve the automatic assessment of moral values in text. This is an important task considering that a comprehensive understanding of moral values at a broader scale could greatly contribute to timely insights into attitudes and judgments concerning social issues, mitigating social polarisation or even uprisings through enhancing the efficacy of communication campaigns [42].

We employ the Moral Foundations Theory (MFT) as the theoretical underpinning to operationalise morality in the following six psychological “foundations” of moral reasoning, divided into “virtue/vice” polarities [30–32]: *Care/Harm* involves concern for others’ suffering and includes virtues like empathy and compassion; *Fairness/Cheating* focuses on issues of unfair treatment, inequality, and justice; *Loyalty/Betrayal* pertains to group obligations such as loyalty and the vigilance against betrayal; *Authority/Subversion* centers on social order and hierarchical responsibilities, emphasising obedience and respect; *Purity/Degradation* relates to physical and spiritual sanctity, incorporating virtues like chastity and self-control; *Liberty/Oppression* addresses feelings of reactance and resentment towards oppressors.

Alongside MFT came a lexicon to guide morality detection in text, illustrating the importance of and need for studying human morality as it manifests in verbal expression, but also highlighting the challenges of the task [25]. As interest in language and morality has

grown, improved dictionaries and other Natural Language Processing (NLP) resources have been developed to study the role of moral values in human life, including ground truth datasets with moral annotations [5, 9, 29, 35, 36, 57, 69]. These works offer quantitative evidence that people project their moral worldviews in a variety of social topics, from the emergence of symbolism and aesthetics of the resistance movement [51] to public perceptions during the COVID-19 vaccination campaigns [9, 11], amongst others.

Also building on the MFT, here we introduce MoralBERT, a range of language representation models fine-tuned to capture moral sentiment in social discourse. We describe a framework for both aggregated and domain-adversarial training on multiple heterogeneous MFT human-annotated datasets sourced from Twitter (before it was rebranded as X), Reddit, and Facebook that broaden textual content diversity in terms of social media audience interests, content presentation and style, and spreading patterns. We show that the proposed framework achieves an average F1 score that is between 11% and 32% higher than lexicon-based approaches, Word2Vec embeddings, and zero-shot classification with large language models such as GPT-4 for in-domain inference. Domain-adversarial training yields better out-of-domain predictions than aggregate training while achieving comparable performance to zero-shot learning.

MoralBERT holds substantial implications for future research. It opens up new possibilities for a more nuanced and context-sensitive understanding of moral narratives surrounding contentious social issues using NLP techniques. These insights can be instrumental in policy-making, social discourse, and conflict resolution by shedding light on the moral dimensions that underpin social stances.

## 2 RELATED WORK

The Moral Foundations Dictionary (MFD) [26] was one of the first lexicons which was created to capture the moral rhetoric in text according to the initial five dimensions and their virtue/vice polarities defined by the MFT [27]. The theory was subsequently extended with Liberty/Oppression as a foundation that deals with the domination and coercion by the more powerful upon the less powerful [30]. Until recently, much work on inferring morality from textual content does not include this foundation because the available linguistic resources are yet limited. Given its importance to recent controversial social discussions such as the vaccination debate, Araque et al. [4] introduced the LibertyMFD lexicon generated based on aligned documents from online news sources with different worldviews.

Traditionally, classification of moral elements in text has been approached via moral lexicons, lists of words depicting moral elements. Lexicons are generated manually [26, 67], via semi-automated methods [5, 71], or expanding a seed list with natural language processing (NLP) techniques [6, 60]. The lexicons are then used to classify morality using text similarity [8, 58]. Moral elements have also been described as knowledge graphs to perform zero-shot classification [7]. More recent methods adopt instead supervised machine learning [2, 38, 45, 47, 64]. Beiro et al. [9] explored the role of Liberty/Oppression in pro- and anti-vaccination Facebook posts using recurrent neural network classifiers with a short memory and entity linking information.

In general, there is a growing interest among researchers in analysing morality and the way modern machines perceive and capture it. For instance, Jiang et al. [41] presented DELPHI, an experimental framework based on neural networks capable of predicting judgements often aligned with human expectations (e.g., right and wrong). This work involved multiple experiments towards inclusive, ethically informed, and socially aware AI systems. Further, Liu et al. [49] presented POLITICS, a model for ideology prediction and stance detection. This model underwent training using novel ideology-focused pre-training objectives, which involved assessing articles on the same topic authored by media outlets with varying ideological perspectives. Another approach was introduced by Mokherian et al. [54] utilising unsupervised techniques to identify moral framing biases in news media.

More closely related to our work, Trager and colleagues [69] presented baseline models for moral values prediction using a pre-trained BERT model fine-tuned in Moral Foundation Reddit Corpus. However, this study is limited to a in-domain work and might not generalise well in other domains. Other studies have introduced out-of-domain approaches moral MFT predictions and presented techniques to enhance model generalisability [29, 57]. Guo et al. [29] presented a multi-label model for predicting moral values while using the domain adversarial training framework proposed by Ganin and Lempitsky [22] to align multiple datasets. Commonly moral classification studies in textual content utilised BERT (Bidirectional Encoder Representations from Transformers) [16]. Due to BERT's widespread adoption, several versions and successor models, including RoBERTa, T5, and DistilBERT, have been developed to effectively tackle a variety of tasks across multiple domains [10].

Recent studies have also explored the capabilities LLMs in understanding moral judgments. Ganguli et al. [21] showed that LLMs trained with reinforcement learning from human feedback can morally self-correct to avoid harmful outputs. These models can follow instructions and learn complex normative concepts like stereotyping, bias, and discrimination. Zhou et al. [74] proposed a theory-guided framework for prompting GPT-4 to perform moral reasoning based on established moral theories, demonstrating its capability to understand and make judgments according to these theories while aligning with human-annotated morality datasets. Scherrer et al. [65] assessed how LLMs encode moral beliefs, finding that in clear-cut scenarios, LLMs align with common sense, but in ambiguous situations, they often express uncertainty. Although most recent LLMs have shown a great performance in understanding complex societal themes, researchers from different fields have shown that smaller but more specialised models like BERT can still reach better accuracy in supervised learning tasks [14, 33].

Building on previous studies, here we present in-domain and out-of domain moral foundations predictions from three major social media platforms. Unlike most of the studies, which only discuss 5 major MFT dimensions, we also analyse Liberty/Oppression foundation [30]. Moreover, we built an extensive set of experiments while comparing MoralBERT models with MoralStrength lexicon [5] Word2Vec model with Random Forrest and zero-shot GPT-4.

### 3 DATA

In this study, we employ three datasets sourced from major social media platforms, all manually annotated for their moral content according to the MFT.

First, we use the Moral Foundations Twitter Corpus (MFTC), a collection of seven distinct datasets totalling 35,108 tweets that have been hand-annotated by at least three trained annotators for five moral foundations (Liberty was not included, see below), each with vice/virtue polarities, resulting in a total of 10 labels [35]. Also a “non-moral” label was used for tweets that are neutral or do not reflect any moral trait. Each tweet can have one or multiple moral labels. Final labels were determined by considering 50% agreement among the annotators. Here we employ six of the seven MFTC datasets, in total 20,628 tweets, focusing on the most populous topic collections, namely, **Hurricane Sandy**, **Baltimore Protest**, **All Lives Matter**, **Davidson Hate Speech**, the **2016 US Presidential Election**, and **Black Lives Matter** (BLM).

For Liberty/Oppression we incorporate newly available annotations for the BLM and 2016 Election datasets, collected via the same procedure and annotation scheme as MFTC [4].

We also use 13,995 Reddit posts from the Moral Foundations Reddit Corpus (MFRC) [69]. MFRC is organised into three buckets: **US politics** with subreddits *conservative*, *antiwork*, and *politics*; **French politics** with subreddits *conservative*, *europa*, *geopolitics*, *neoliberal* and *worldviews*; **Everyday Morality** with subreddits like *IAmTheAsshole*, *confession*, *nostalgia* and *relationship\_advice*. Similarly to MFTC, at least three trained annotators were used and a 50% agreement threshold was maintained for final labels. MFRC includes annotations for Proportionality and Equality, which we combine and label as Fairness. Similarly to MFTC, MFRC does not include annotations for the moral foundation of Liberty. Unlike MFTC, MFRC does not account for the polarity of moral foundations. To address this, we used VADER sentiment scores as weights for vice/virtue per foundation [36].

Lastly, we use a dataset of 1,509 Facebook posts related to **pro- and anti-vaccination**, each hand-annotated by nine researchers familiar with MFT following a similar annotation scheme to MFTC [9]. Annotation labels cover virtue and vice polarity for each MFT category, including Liberty/Oppression, and “non-moral” labels were also used to denote either moral neutrality or lack of any moral trait. Cohen’s kappa between annotators was 0.32, indicating fair agreement, but also speaking to the difficulty of detecting morality in text also for human annotators.

We consider each dataset as a distinct domain for morality learning. Distinct social media platforms possess diverse linguistic and social structural environments, potentially leading to variations in moral language [15]. They also differ in audience interests, content presentation/style, and spreading patterns [40, 63]. Accordingly, we hypothesised that the expression of moral values in tweets versus Reddit comments versus Facebook posts may vary across training data sourced from the corresponding corpora described above.

Table 1 illustrates the variation in the MFT label distribution including neutral (i.e., non-moral) text across the three social media datasets. The left graph in Figure 1 illustrates using manifold

approximation and projection (UMAP [50]) how the three corpora domains differ in the feature embedding space. Feature distributions are generally distinct across the three datasets. There is some overlap between FB and MFRC, possibly because content in the respective platforms tends to be longer and more elaborate, while tweets are shorter due to character limits, resulting in more fragmented discussions and frequent updates. This overlap becomes less prominent when neutral (non-moral) text is excluded (right graph in Figure 1), indicating that morally nuanced text is more clearly separated based on social media platform.

**Table 1: Distribution of human-annotated moral values in the three social media corpora used in this study. †Annotations for Liberty/Oppression are only available in FB (full corpus) and MFTC (BLM and 2016 US Election datasets).**


	MFTC	MFRC	FB	Total
Care	1658	737	357	2752
Harm	2027	1014	132	3173
Fairness	1575	623	173	2369
Cheating	2037	841	123	3001
Loyalty	1027	241	40	1308
Betrayal	1338	188	38	1564
Authority	824	330	110	1264
Subversion	565	357	204	1126
Purity	535	100	80	715
Degradation	746	187	112	1045
Non-Moral	7739	9842	367	17948
Liberty†	2136	2284	140	4560
Oppression†	1059	1028	65	2152
Non-Moral†	692	735	367	1794

### 4 MORALBERT

To capture moral expressions in social media discourse, we propose MoralBERT,<sup>1</sup> a series of transformer-based language models fine-tuned with the corpora presented in the previous section. These models use the BERT-base-uncased pretrained sequence classifier [17] with 768 hidden layers, 12 transformer layers and 110M parameters. We use the Adam optimiser with a learning rate of 5e-5 [18]. Due to data sparsity (see Table 1), we opted for a single-label classification approach, whereby each model predicts the presence or absence of a moral virtue or vice.

We further compare MoralBERT with MoralBERT<sub>adv</sub>, an extension of the former with domain adversarial training to account for heterogeneous training data. Following Guo et al. [29], we first obtain a domain invariant representation  $h = W_{inv}e$  from the BERT embedding  $e$ , where  $W_{inv} \in \mathbb{R}^{768 \times 768}$  is a learnable matrix, which means that every element in this matrix is a parameter that can be adjusted during the training process. We then obtain predictions of moral values  $y_m = \text{Softmax}(W_1(\text{ReLU}(W_2h)))$  where  $W_1 \in \mathbb{R}^{768 \times 768}$  and  $W_2 \in \mathbb{R}^{768 \times c}$  are learnable matrices with parameters that the training process adjusts to optimize the model’s

<sup>1</sup><https://github.com/vjosapreniqi/MoralBERT>



**Figure 1: UMAP visualisation of feature distributions across datasets (domains). The left graph includes both moral and non-moral (neutral) labelled text, the right graph includes only moral labelled text. Dots represent mean-pooled BERT embeddings for each text example.**

performance.  $c$  is the number of classes (0 represents neutral class, 1 represents the moral class),  $ReLU$  is the rectified linear unit activation function and  $Softmax$  is the normalised exponential function that gives the probability distribution over predicted classes. We include a domain classification head similar to the moral values classification head ( $\hat{y}_m$ ) to obtain domain predictions ( $\hat{y}_d$ ). The adversarial network connects the domain classifier head to the model via a gradient reversal layer, maximising domain classification loss while minimising the moral values prediction objective. We use the cross-entropy loss function for moral loss and domain loss representation.

We also added two regularization terms [29], an L2 norm and a reconstruction loss, so that the original BERT does not get driven too far away from the original output embeddings by the fine-tuning process:  $L_{norm} = \|W_{inv} - I\|^2$  and  $L_{rec} = \|W_{rech} - e\|^2$  where  $I$  is the identity matrix and  $W_{rech}$  a learnable matrix to reconstruct the original from the transformed embeddings. We calculate the total loss by adding  $L_{rec}$  and  $L_{norm}$  to the moral values classification loss and the domain classification loss.

For both MoralBERT and MoralBERT<sub>adv</sub> we assign class weights to address the class imbalance problem evident in Table 1. We do so by employing the approach of King and Zeng [44] so that for each class  $c$  we compute  $weight_c = N/N_c$  with  $N$  being the total number of samples in the training data and  $N_c$  the total number of samples in the training data belonging to class  $c$ . We trained models with a batch size of 16, and input sequences are tokenised to a maximum length of 150 tokens determined by the maximum sentence size across the three combined datasets. Each model was trained for five epochs, and the model checkpoints from the best epoch were saved for testing.

## 5 EXPERIMENTS

The performance of MoralBERT and MoralBERT<sub>adv</sub> was evaluated for in-domain and out-of-domain classification. To infer the 10

moral virtue/vice labels annotated across all three corpora (MFTC, MFRC, FB), in-domain models were trained using 80% of the combined data from all datasets and tested in the remaining 20%. For out-of-domain models, we first train on two of the three corpora (e.g., MFTC and MFRC) and then test on the left-out dataset (e.g., FB). Due to the partial annotation of Liberty/Oppression in our datasets (MFTC BLM, MFTC Election, FB), we carried out separate experiments to infer this moral foundation following the same in-domain train-test split (80%-20%) and out-of-domain setup (training on FB and testing for MFTC, and vice versa). For all experiments we report the F1 Binary score, which focuses solely on moral labels and measures each model’s accuracy in predicting true positives, and the F1 Macro score, which includes non-moral or neutral labels and measures each model’s accuracy in predicting both true positives and true negatives. In all data we use for fine-tuning and testing models, we cleaned the text by removing URLs, substituting mentions with “@user”, removing hashtags, substituting emojis with their textual descriptions, and removing any non-ASCII characters using the *re* Python library [20].

We employ two traditional baselines from previous works in this field. First, we use the MoralStrength lexicon [5] as a foundational estimate for each moral category. MoralStrength is an extension of MFD and offers a significantly larger set of morally annotated lemmas. It not only provides the moral valence score but also indicates the intensity of the lemma. Here, we use the MFT scores from MoralStrength to categorise values discretely to align with the output of MoralBERT. Second, we quantise the textual data using Word2Vec, a widely used model in Natural Language Processing known for its word embedding capabilities [53] and utilise a machine learning technique such as the Random Forest (RF) classifier with default parameters from the scikit-learn Python library [59] to predict each moral category. Word2Vec is a good method for handling large datasets and learning distributional properties of words, as well as syntactic and semantic word relationship [53]. It demonstrates high performance in terms of both accuracy and

computational efficiency [23]. However, Word2Vec embeddings do not model context, which makes them unsuitable for analyzing sentences.

We further compare MoralBERT and MoralBERT<sub>adv</sub> with a powerful LLM such as GPT-4 [1] deployed as a zero-shot classifier. LLMs like GPT-4 are trained on diverse text sources such as Wikipedia, GitHub, chat logs, books, and articles [13]. This enables them to generalise and understand language across various domains [19]. The earlier model, GPT-3, contains 175 billion parameters, a figure vastly greater than BERT-base and BERT-large models (110M and 340M parameters) [12]. Given the size, cost, and significant energy consumption of these models, we used them for only 20% of the data. The data selection was partially controlled. We selected 3,384 tweets from MFTC data, 2,793 Reddit posts from MFRC and 1509 posts from FB. We then prompted the classification task as follows:

*You will be provided with social media posts from Twitter, Reddit and Facebook, regarding different social topics. The social media posts will be delimited with #### characters. Classify each social media post into 12 Possible Moral Foundations as defined in Moral Foundation Theory. The available Moral Foundations are: {Moral Foundations Tags}. The explanation of the moral foundations is as follows: {Description tags}. This is a multi-label classification problem: where it's possible to assign one or multiple categories simultaneously. Report the results in JSON format such that the keys of the correct moral values are reported in a list.*

The {Moral Foundations Tags} represent the 12 moral virtues and vices, while the {Description Tags} provide a one-sentence description for each category.

## 6 RESULTS

Table 2 shows that MoralBERT<sub>adv</sub> had the highest performance for in-domain predictions. It achieved a 17% higher F1 binary score compared to GPT-4, a 22% higher score than MoralStrength, and a 32% higher score Word2Vec with Random Forest models. The improved performance is also reflected in the F1 Macro scores. On average, MoralBERT<sub>adv</sub> surpasses GPT-4 by 11%, MoralStrength by 12%, and Word2Vec with Random Forest by 16% in F1 macro score.

Figures 2a and 2b show that MoralBERT<sub>adv</sub> performs marginally better than standard MoralBERT in F1 Binary and Macro average scores for out-of-domain predictions. For certain moral foundations, MoralBERT<sub>adv</sub> shows significant improvements. For instance, Degradation predictions improve on MFRC and MFTC, and Loyalty and Authority predictions enhance on MFRC. As such, these moral foundations may be expressed differently across domains, and domain adaptation in MoralBERT<sub>adv</sub> enables the model to identify these patterns.

We wanted to compare MoralBERT and GPT-4 for out-of-domain moral predictions. For this we used the MoralBERT model trained on MFTC and MFRC, and tested on FB, the smallest of our social media datasets with 1,509 posts, which allowed us to apply the zero-shot GPT-4 classification model to the entire Facebook dataset. Inference on larger data could not be performed due to higher cost. Table 3 reveals that the prediction results of MoralBERT, MoralBERT<sub>adv</sub>, and GPT-4 are very similar, with GPT-4 achieving an average of 1% higher F1 binary Score and 2% higher F1 macro score.

For Liberty/Oppression in-domain predictions showed in Figure 5, the MoralBERT and MoralBERT<sub>adv</sub> performed better than GPT-4 with an average of 33% higher F1 Binary Score and 19% higher F1 Macro Score. In an out-of-domain setup for predicting this foundation in Facebook posts, zero-shot GPT-4 performed better than MoralBERT and MoralBERT<sub>adv</sub>, achieving an average F1 binary score 18% higher and an F1 macro score 15% higher. The low performance of MoralBERT when tested on Facebook Data may be attributed to the marginal inner-annotator agreement (0.38 Cohen's kappa coefficient) observed in the Facebook posts, indicating that these posts might be complex and ambiguous. In contrast, MoralBERT and MoralBERT<sub>adv</sub> performed significantly better when predicting Liberty/Oppression on MFTC tweets about BLM and the 2016 US Elections, with an average F1 binary score 25% higher and an F1 macro score 8% higher.


To quantitatively analyse the models' performance, we presented individual examples from the social media posts in the three datasets, annotated by human annotators, along with the results from the MoralBERT<sub>adv</sub> and the GPT-4 classification model in Table 4. From the examples it can be seen that the text in the posts contains informal language, grammar mistakes, and many abbreviations. Further, some of the posts are written in an argumentative tone, and some use more personal and emotional nuances. We can also see that Reddit comments and Facebook posts are typically much longer than tweets. Since GPT-4 model is used as zero-shot approach it predicts moral labels based on the prompting request and the general knowledge that this model has Moral Foundation Theory. Previous works have shown that LLMs like GPT-4 can indeed perform moral reasoning through the lens of moral theories [74] which is evident in our examples as well. On the other hand, MoralBERT<sub>adv</sub> learns to align more closely with the trends seen in the fine-tuned annotation examples. However, MoralBERT<sub>adv</sub> often predicts both Liberty and Oppression even if only one is mentioned in the text.

## 7 DISCUSSION

Moral values and judgments significantly influence our daily lives. Psychologists argue that moral judgment is not a rigorous reasoning process. Instead, it is influenced by more personal factors, including intuitions and emotions [28]. Making moral judgments is intrinsically challenging, even for humans, due to lack of a universal standard [74]. People from different beliefs and cultural backgrounds can have significantly different attitudes toward the same topic [37]. Furthermore, moral inferences are highly context-dependent [3] and different contexts can lead to distinct judgments [74]. Similarly, Guo et al. [29] showed that also the writing culture matters; in their study, using a model trained on MFTC to predict moral values in news articles (eMFD dataset [36]) was shown significantly more challenging than predicting moral values on another Twitter dataset with discourse around COVID-19 vaccination. In a more meticulous analysis, Lisco et al. [46] demonstrated that the predictability of moral values depends heavily on the distribution of moral rhetoric within a domain, and the further apart the domains are, the weaker the predictions of moral foundations become. This is evident in our study too, with MoralBERT models trained within the domain to be notably more successful at making moral inferences than those trained out-of-domain. To mitigate this issue, we implemented the

**Table 2: In-domain prediction results for 10 Moral Foundations, showing F1 Binary and Macro average scores. Standard deviations are based on 1,000 bootstrap samples. C. = Care, H. = Harm, Ch. = Cheating, F. = Fairness, L. = Loyalty, B. = Betrayal, A. = Authority, S. =Subversion, P. = Purity, D. = Degradation.**

	F1 Binary					F1 Macro				
	MS Lex.	W2V RF	GPT-4	MoralBERT	MoralBERT <sub>adv</sub>	MS Lex.	W2V RF	GPT-4	MoralBERT	MoralBERT <sub>adv</sub>
C.	.31 ± .02	.14 ± .02	.42 ± .01	.48 ± .02	<b>.50 ± .02</b>	.63 ± .01	.55 ± .01	.66 ± .01	.71 ± .01	<b>.73 ± .01</b>
H.	.38 ± .02	.07 ± .01	.41 ± .01	.55 ± .01	<b>.56 ± .02</b>	.65 ± .01	.51 ± .01	.64 ± .01	.75 ± .01	<b>.76 ± .01</b>
F.	.32 ± .02	.35 ± .03	.30 ± .01	.56 ± .02	<b>.57 ± .02</b>	.62 ± .01	.66 ± .01	.55 ± .01	.76 ± .01	<b>.77 ± .01</b>
Ch.	.19 ± .02	.15 ± .02	.34 ± .02	.60 ± .01	<b>.61 ± .01</b>	.57 ± .01	.55 ± .01	.64 ± .01	.78 ± .01	<b>.79 ± .01</b>
L.	.36 ± .02	.28 ± .03	.39 ± .02	.57 ± .02	<b>.64 ± .03</b>	.66 ± .01	.63 ± .02	.68 ± .01	.78 ± .01	<b>.81 ± .01</b>
B.	.14 ± .02	.13 ± .02	.18 ± .02	.32 ± .03	<b>.40 ± .04</b>	.55 ± .01	.55 ± .01	.57 ± .01	.65 ± .02	<b>.69 ± .02</b>
A.	.24 ± .02	.22 ± .03	.19 ± .01	.37 ± .02	<b>.39 ± .03</b>	.59 ± .01	.60 ± .02	.55 ± .01	.67 ± .01	<b>.69 ± .01</b>
S.	.25 ± .02	.10 ± .03	.22 ± .02	.36 ± .02	<b>.37 ± .02</b>	.60 ± .01	.54 ± .01	.58 ± .01	<b>.67 ± .01</b>	<b>.67 ± .01</b>
P.	.17 ± .02	.06 ± .03	.44 ± .02	.49 ± .03	<b>.49 ± .02</b>	.56 ± .01	.53 ± .01	.71 ± .01	<b>.74 ± .01</b>	.73 ± .01
D.	<b>.28 ± .02</b>	.12 ± .03	.21 ± .02	.23 ± .02	.25 ± .03	<b>.62 ± .01</b>	.55 ± .01	.59 ± .01	.60 ± .01	.61 ± .02
Avg.	.26 ± .02	.16 ± .03	.31 ± .01	.45 ± .02	<b>.48 ± .02</b>	.61 ± .01	.57 ± .01	.62 ± .01	.71 ± .01	<b>.73 ± .01</b>



**Figure 2: Out-of-domain classification: for each test dataset, models are fine-tuned on the other two datasets. Bar heights represent F1 Binary and Macro average scores; error bars indicate standard deviation estimated via 1,000 bootstraps.**

domain adversarial module (MoralBERT<sub>adv</sub> models) which resulted in marginal improvements in out-of-domain prediction models,

demonstrating that out-of-domain moral inferences remain a challenging task.

**Table 3: Zero-shot (GPT-4) versus out-of-domain (MoralBERT) classification, showing F1 Binary and Macro average scores and standard deviation estimated via 1,000 bootstraps. Models are fine-tuned on MFTC and MFRC and tested on FB.**

	F1 Binary			F1 Macro		
	GPT-4	MoralBERT	MoralBERT <sub>adv</sub>	GPT-4	MoralBERT	MoralBERT <sub>adv</sub>
Care	<b>.51 ± .02</b>	.48 ± .02	.50 ± .02	.62 ± .01	.64 ± .01	<b>.65 ± .01</b>
Harm	.25 ± .02	.25 ± .03	<b>.28 ± .03</b>	.47 ± .01	<b>.57 ± .02</b>	.57 ± .01
Fairness	<b>.34 ± .02</b>	.26 ± .02	.29 ± .02	<b>.59 ± .01</b>	.43 ± .01	.55 ± .01
Cheating	.14 ± .03	<b>.16 ± .02</b>	.17 ± .02	<b>.53 ± .02</b>	.52 ± .01	.42 ± .01
Loyalty	.14 ± .06	<b>.18 ± .06</b>	.16 ± .05	.56 ± .03	<b>.58 ± .03</b>	.56 ± .03
Betrayal	.06 ± .03	.05 ± .01	<b>.08 ± .04</b>	.51 ± .02	.42 ± .01	<b>.53 ± .02</b>
Authority	<b>.21 ± .03</b>	.15 ± .02	.15 ± .03	<b>.56 ± .02</b>	.48 ± .01	.52 ± .01
Subversion	.23 ± .03	.28 ± .02	<b>.29 ± .02</b>	<b>.57 ± .02</b>	.45 ± .01	.52 ± .01
Purity	<b>.34 ± .04</b>	.21 ± .04	.25 ± .04	<b>.65 ± .02</b>	.58 ± .02	.60 ± .02
Degradation	<b>.25 ± .03</b>	.19 ± .02	.22 ± .03	<b>.59 ± .02</b>	.47 ± .01	.55 ± .02
Avg.	<b>.25 ± .03</b>	.22 ± .03	.24 ± .03	<b>.57 ± .02</b>	.51 ± .01	.55 ± .02

**Table 4: Examples of human-annotated and machine-learned moral values in social media discourse. GPT-4 is zero-shot classification; MoralBERT<sub>adv</sub> predictions are out-of-domain. R = Reddit; T = Twitter; F = Facebook.**

Text	Human	GPT-4	MoralBERT <sub>adv</sub>
"And yet, more and more space and laws protect these people in their host countries. When are people in power going to wake up? OH RIGHT!, Le pen was, and it backfired on her" [R]	Authority	Authority, Subversion	Care, Authority
"I'll be blunt. I don't care whether a government (Macron's or anyone else's) has gender parity in its cabinet, all I actually hope is that the best people are chosen, regardless of having a wiener or not. For me this is what equality should be like" [R]	Fairness	Fairness	Fairness
"Those who deceive young men by selling war as an adventure are cruel monsters." [T]	Harm, Cheating, Oppression	Care, Harm	Harm, Cheating, Betrayal
"My tribute today to Sardar Patel-a Congress stalwart,who strove for communal harmony; dedicated his life to the unity" [T]	Loyalty	Care, Fairness, Loyalty	Loyalty
"Viruses and bacteria have no respect for religious beliefs. They will attack regardless. VERY few faiths promote an anti-vaccine agenda. Most consider the body to be a sacred gift that must receive proper care" [F]	Care, Purity	Care, Purity	Care, Purity
"It's a travesty that kids are exposed to the insanity of Big Pharma. Parents must take the CO route and protect their kids. Meanwhile, get active in anti-vaccine groups since the PTB really do want mandatory vaccines or the kids will be given over to foster homes" [F]	Care, Subversion, Liberty	Liberty, Oppression, Authority	Care, Subversion, Liberty, Oppression

In our study, we face both challenges; diverse linguistic styles, since each dataset is sourced from a different social media platform, and a variety of social topics treated. Another important factor in predicting moral values with language models like BERT is the distribution of moral labels in fine-tuning data [29]. Unlike FB, the MFTC and MFRC corpora are highly imbalanced, with non-moral labelled text dominating over text labelled with moral values. The class weighting technique we employed played a significant role in addressing this issue in both the MoralBERT and MoralBERT<sub>adv</sub> models.

Regarding the experimental design, we opted for both a single-label approach, (predicting each moral value/vice separately), and a multi-label approach (predicting all moral values/vices at once) and as expected the results of the latter were significantly weaker. Thus we report the single-label experiments only. The drop in the

multi-label prediction approach relates to the intrinsic interdependence of moral dimensions [48], which in our case is particularly challenging because some dimensions, like Care and Harm, often overlap with other moral dimensions. Instead, in the single-label setup, the prediction task is simpler, with the models focusing on learning specifically the dimension of interest and distinguishing it from morally neutral text. Furthermore, the single-label design approach in general showed better performance for out of domain predictions, in line with the findings of Guo et al. [29].

In this work we introduce classification models for the Liberty/Oppression foundation, which were not previously considered in transformer-based approaches to morality inference from social media discourse [29, 57, 69]. Given that Liberty/Oppression values are rooted in reason more than emotion [39], they remain crucial for understanding decision-making across various contexts. They have been shown to be particularly relevant to current social issues

**Table 5: In-domain and out-of-domain predictions of Liberty/Oppression, showing F1 Binary and Macro average scores; standard deviation is calculated with 1000 bootstraps.**

	F1 Binary			F1 Macro		
	<i>in-domain experiments</i>					
	GPT-4	MoralBERT	MoralBERT <sub>adv</sub>	GPT-4	MoralBERT	MoralBERT <sub>adv</sub>
Liberty	.24 ± .02	.63 ± .01	<b>.66 ± .01</b>	.48 ± .01	.70 ± .01	<b>.71 ± .01</b>
Oppression	.17 ± .02	<b>.45 ± .02</b>	.40 ± .02	.51 ± .01	<b>.68 ± .01</b>	.55 ± .01
	<i>out-of-domain experiments, test dataset is FB (vaccination)</i>					
Liberty	<b>.39 ± .03</b>	.19 ± .02	.19 ± .02	<b>.62 ± .01</b>	.39 ± .01	.39 ± .01
Oppression	<b>.20 ± .03</b>	.05 ± .01	.09 ± .01	<b>.56 ± .02</b>	.49 ± .01	.30 ± .01
	<i>out-of-domain experiments, test dataset is MFTC (BLM and 2016 US Elections)</i>					
Liberty	.17 ± .01	<b>.59 ± .01</b>	.57 ± .01	.39 ± .01	.52 ± .01	<b>.53 ± .01</b>
Oppression	.17 ± .02	.25 ± .02	<b>.27 ± .02</b>	.48 ± .01	.49 ± .01	<b>.50 ± .01</b>

such as the vaccination debate, poverty, and radicalisation [6, 51]. In the present experiments, despite having significantly more limited training data resources, inference for Liberty/Oppression was comparable to that for the other 10 moral foundations.

We benchmarked our MoralBERT and MoralBERT<sub>adv</sub> with both lexicon-based approaches, namely MoralStrength lexicon, Word2Vec, on top of Random Forrest as more traditional baselines, and large language models (LLM zero-shot GPT-4 classification). We showed that our models on average outperformed all other approaches for in-domain set-up with an approximate 11% to 32% improvement in F1 score, while for out-of domain the performances drops slightly as expected, but remain comparable with the GPT-4 classification results. Overall, both MoralBERT and MoralBERT<sub>adv</sub> were better for predicting Liberty/Oppression in Twitter data. Noteworthy is the fact that LLM models like GPT-4 are trained on billion of parameters, are extremely large, expensive, and consume significant amounts of energy, leading to various environmental implications and issues [72]. This shows that BERT-based models can be just as effective as larger LLMs once fine-tuned with considerably less resources. Moreover, the training is based on human annotators, ensuring that the models learn from human moral reasoning. This is crucial especially for moral values assessment, misinterpretation of which can lead to social polarisation amplification. Another important point is that BERT-based approaches still provide interpretable results fundamental in assessing and examining how the model makes decisions on controversial social issues.

There is an ever-increasing interest in understanding moral values via natural language processing even in more artistic fields and beyond social media contexts. Recently, researchers have explored moral values in movie synopses [23] and lyrics [61, 62] using different dictionary and lexicon approaches. Our approach, utilising fine-tuned models for predicting moral values, will provide a valuable starting point for exploring morality in various contexts. Understanding moral values from written content can greatly enhance communication and support social campaigns, but it also carries risks if used for malicious or manipulative purposes. Automatic annotation of morality in text can misrepresent individuals' moral positions or unfairly categorise them, leading to social stigmatisation and discrimination [70].

Our research has certain limitations. First, although we gathered a substantial amount of textual content from three different social media platforms, a large portion of the data was labeled as non-moral or neutral. This leads to data imbalance issues which we tried to handle using a standard class weighting technique. Second, we could only partially use the data for the Liberty/Oppression foundation classification because this foundation was not present in all the datasets. Third, our study focused exclusively on English-language posts, which limits our understanding of how moral rhetoric is shaped across different cultures.

In the future, we aim to expand our investigations into multilingual models for understanding cross-cultural moral narratives. Also, we will explore moral expressions in other domains, such as music lyrics, which often contain more complex linguistic structures and figurative expressions. Additionally, we plan on employing techniques for distilling the knowledge [34] from LLMs like GPT-4 and Llama 2 [68] in creating synthetic data which can be then used for fine-tuning language models like BERT models which are comparative for narrower tasks. This will help on further improving the current results in capturing moral values while reducing the need for manual annotation through the use of synthetic data generation [52, 73]. By doing so, we can leverage the strengths of both types of models and improve our models' understanding of moral expressions in text across various domains and situations. In general, we believe that this work is particularly timely, considering the current surge in research dedicated to identifying moral narratives in textual data. Even though there is room for improvement, our approach still holds significant value for the research community and beyond.

## ACKNOWLEDGMENTS

VP and IG are supported by PhD studentships from Queen Mary University of London's Centre for Doctoral Training in Data-informed Audience-centric Media Engineering. KK acknowledges support from the Lagrange Project of the Institute for Scientific Interchange Foundation (ISI Foundation) which is funded by Fondazione Cassa di Risparmio di Torino (Fondazione CRT).



## REFERENCES

- [1] Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altenschmidt, Sam Altman, Shyamal Anadkat, et al. 2023. Gpt-4 technical report. *arXiv preprint arXiv:2303.08774* (2023).
- [2] Milad Alshomary, Roxanne El Baff, Timon Gurcke, and Henning Wachsmuth. 2022. The Moral Debater: A Study on the Computational Generation of Morally Framed Arguments. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (ACL '22)*. Association for Computational Linguistics, Dublin, Ireland, 8782–8797. <https://aclanthology.org/2022.acl-long.601.pdf>
- [3] Prithviraj Ammanabrolu, Liwei Jiang, Maarten Sap, Hannaneh Hajishirzi, and Yejin Choi. 2022. Aligning to Social Norms and Values in Interactive Narratives. In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*. 5994–6017.
- [4] Oscar Araque, Lorenzo Gatti, Sergio Consoli, and Kyriaki Kalimeri. 2024. A Novel Lexicon for the Moral Foundation of Liberty. *arXiv:2407.11862 [cs.CL]* <https://arxiv.org/abs/2407.11862>
- [5] 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), 1–11. <https://linkinghub.elsevier.com/retrieve/pii/S095070511930526X>
- [6] Oscar Araque, Lorenzo Gatti, and Kyriaki Kalimeri. 2022. LibertyMFD: A Lexicon to Assess the Moral Foundation of Liberty. In *Proceedings of the 2022 ACM Conference on Information Technology for Social Good*. 154–160.
- [7] Luigi Asprino, Luana Bulla, Stefano De Giorgis, Aldo Gangemi, Ludovica Marinucci, and Misael Mongiovi. 2022. Uncovering Values: Detecting Latent Moral Content from Natural Language with Explainable and Non-Trained Methods. In *Proceedings of Deep Learning Inside Out: The 3rd Workshop on Knowledge Extraction and Integration for Deep Learning Architectures (DeeLIO '22)*. Association for Computational Linguistics, Dublin, Ireland and Online, 33–41. <https://aclanthology.org/2022.deeLIO-1.4>
- [8] Mohamed Bahgat, Steven R. Wilson, and Walid Magdy. 2020. Towards Using Word Embedding Vector Space for Better Cohort Analysis. In *Proceedings of the International AAAI Conference on Web and Social Media (ICWSM '20)*. AAAI Press, Atlanta, Georgia, 919–923. <https://ojs.aaai.org/index.php/ICWSM/article/view/7358>
- [9] Mariano Gastón Beiró, Jacopo D'Ignazi, Victoria Perez Bustos, María Florencia Prado, and Kyriaki Kalimeri. 2023. Moral narratives around the vaccination debate on facebook. In *Proceedings of the ACM Web Conference 2023*. 4134–4141.
- [10] Jordan J Bird, Anikó Ekárt, and Diego R Faria. 2023. Chatbot Interaction with Artificial Intelligence: human data augmentation with T5 and language transformer ensemble for text classification. *Journal of Ambient Intelligence and Humanized Computing* 14, 4 (2023), 3129–3144.
- [11] Judith Borghouts, Yicong Huang, Sydney Gibbs, Suellen Hopfer, Chen Li, and Gloria Mark. 2023. Understanding underlying moral values and language use of COVID-19 vaccine attitudes on twitter. *PNAS nexus* 2, 3 (2023), pgad013.
- [12] Mitchell Bosley, Musashi Jacobs-Harukawa, Hauke Licht, and Alexander Hoyle. 2023. Do we still need BERT in the age of GPT? Comparing the benefits of domain-adaptation and in-context-learning approaches to using LLMs for Political Science Research. (2023).
- [13] Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. 2020. Language models are few-shot learners. *Advances in neural information processing systems* 33 (2020), 1877–1901.
- [14] Shan Chen, Yingya Li, Sheng Lu, Hoang Van, Hugo JW Aerts, Guergana K Savova, and Danielle S Bitterman. 2024. Evaluating the ChatGPT family of models for biomedical reasoning and classification. *Journal of the American Medical Informatics Association* 31, 4 (2024), 940–948.
- [15] Stephan A Curiskis, Barry Drake, Thomas R Osborn, and Paul J Kennedy. 2020. An evaluation of document clustering and topic modelling in two online social networks: Twitter and Reddit. *Information Processing & Management* 57, 2 (2020), 102034.
- [16] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2018. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805* (2018).
- [17] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2018. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. <https://huggingface.co/google/bert-base-uncased>. Accessed: 2024-06-06.
- [18] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics (NAACL '19)*. 4171–4186. <https://aclanthology.org/N19-1423>
- [19] Seunghoon Doh, Keunwoo Choi, Jongpil Lee, and Juhan Nam. 2023. LP-MusicCaps: LLM-Based Pseudo Music Captioning. In *Ismir 2023 Hybrid Conference*.
- [20] Python Software Foundation. 2023. *Python Regular Expression (re) Library*. <https://docs.python.org/3/library/re.html>. Accessed: 2024-06-06.
- [21] Deep Ganguli, Amanda Askell, Nicholas Schiefer, Thomas I Liao, Kamile Lukošiušė, Anna Chen, Anna Goldie, Azalia Mirhoseini, Catherine Olsson, Danny Hernandez, et al. 2023. The capacity for moral self-correction in large language models. *arXiv preprint arXiv:2302.07459* (2023).
- [22] Yaroslav Ganin and Victor Lempitsky. 2015. Unsupervised domain adaptation by backpropagation. In *International conference on machine learning*. PMLR, 1180–1189.
- [23] Carlos González-Santos, Miguel A Vega-Rodríguez, Carlos J Pérez, Joaquín M López-Muñoz, and Iñaki Martínez-Sarrigui. 2023. Automatic assignment of moral foundations to movies by word embedding. *Knowledge-Based Systems* 270 (2023), 110539.
- [24] Jesse Graham, Jonathan Haidt, Sena Koleva, Matt Motyl, Ravi Iyer, Sean P. Wojcik, and Peter H. Ditto. 2013. Moral Foundations Theory: The Pragmatic Validity of Moral Pluralism. In *Advances in Experimental Social Psychology*. Vol. 47. Elsevier, Amsterdam, the Netherlands, 55–130. <https://doi.org/10.1016/B978-0-12-407236-7.00002-4>
- [25] 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.
- [26] 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), 1029–1046. <https://doi.org/10.1037/a0015141>
- [27] Jesse Graham, Brian A Nosek, Jonathan Haidt, Ravi Iyer, Spassena Koleva, and Peter H Ditto. 2011. Mapping the moral domain. *Journal of personality and social psychology* 101, 2 (Aug. 2011), 366–85.
- [28] Joshua Greene and Jonathan Haidt. 2002. How (and where) does moral judgment work? *Trends in cognitive sciences* 6, 12 (2002), 517–523.
- [29] Siyi Guo, Negar Mokherian, and Kristina Lerman. 2023. A Data Fusion Framework for Multi-Domain Morality Learning. In *Proceedings of the International AAAI Conference on Web and Social Media*, Vol. 17. 281–291.
- [30] Jonathan Haidt. 2012. *The righteous mind: Why good people are divided by politics and religion*. Vintage.
- [31] Jonathan Haidt and Jesse Graham. 2007. When morality opposes justice: Conservatives have moral intuitions that liberals may not recognize. *Social justice research* 20, 1 (2007), 98–116.
- [32] Jonathan Haidt and Craig Joseph. 2004. Intuitive ethics: How innately prepared intuitions generate culturally variable virtues. *Daedalus* 133, 4 (2004), 55–66.
- [33] Evan Hernandez, Diwakar Mahajan, Jonas Wulff, Micah J Smith, Zachary Ziegler, Daniel Nadler, Peter Szolovits, Alistair Johnson, Emily Alsentzer, et al. 2023. Do we still need clinical language models?. In *Conference on Health, Inference, and Learning*. PMLR, 578–597.
- [34] Geoffrey Hinton, Oriol Vinyals, and Jeff Dean. 2015. Distilling the knowledge in a neural network. *arXiv preprint arXiv:1503.02531* (2015).
- [35] Joe Hoover, Gwenyth Portillo-Wightman, Leigh Yeh, Shreya Havaladar, 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.
- [36] 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 (2021), 232–246.
- [37] Minda Hu, Ashwin Rao, Mayank Kejriwal, and Kristina Lerman. 2021. Socioeconomic correlates of anti-science attitudes in the US. *Future Internet* 13, 6 (2021), 160.
- [38] Xiaolei Huang, Alexandra Wormley, and Adam Cohen. 2022. Learning to Adapt Domain Shifts of Moral Values via Instance Weighting. In *Proceedings of the 33rd ACM Conference on Hypertext and Social Media (HT '22)*. Association for Computing Machinery, 121–131. <https://doi.org/10.1145/3511095.3531269>
- [39] Ravi Iyer, Spassena Koleva, Jesse Graham, Peter Ditto, and Jonathan Haidt. 2012. Understanding libertarian morality: The psychological dispositions of self-identified libertarians. (2012).
- [40] Kokil Jaidka, Sharath Guntuku, and Lyle Ungar. 2018. Facebook versus Twitter: Differences in self-disclosure and trait prediction. In *Proceedings of the International AAAI Conference on Web and Social Media*, Vol. 12.
- [41] Liwei Jiang, Jena D Hwang, Chandra Bhagavatula, Ronan Le Bras, Jenny Liang, Jesse Dodge, Keisuke Sakaguchi, Maxwell Forbes, Jon Borchardt, Saadia Gabriel, et al. 2021. Can machines learn morality? the delphi experiment. *arXiv preprint arXiv:2110.07574* (2021).
- [42] Kyriaki Kalimeri, Mariano G. Beiró, Matteo Delfino, Robert Raleigh, and Ciro Cattuto. 2019. Predicting demographics, moral foundations, and human values from digital behaviours. *Computers in Human Behavior* 92 (2019), 428–445. <https://doi.org/10.1016/j.chb.2018.11.024>
- [43] 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 (WWW '19), 248–254. <https://doi.org/10.1145/3308560.3316489>
- [44] Gary King and Langche Zeng. 2001. Logistic regression in rare events data. *Political analysis* 9, 2 (2001), 137–163.
- [45] Alex Gwo Jen Lan and Ivandrè Paraboni. 2022. Text- and author-dependent moral foundations classification. *New Review of Hypermedia and Multimedia* 0, 0 (2022), 1–21. <https://doi.org/10.1080/13614568.2022.2092655>
- [46] Enrico Liscio, Oscar Araque, Lorenzo Gatti, Ionut Constantinescu, Catholijn Jonker, Kyriaki Kalimeri, and Pradeep Kumar Murukannaiah. 2023. What does a text classifier learn about morality? An explainable method for cross-domain comparison of moral rhetoric. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*. 14113–14132.
- [47] Enrico Liscio, Alin E. Dondera, Andrei Geadau, Catholijn M. Jonker, and Pradeep K. Murukannaiah. 2022. Cross-Domain Classification of Moral Values. In *Findings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics (NAACL '22)*. Association for Computational Linguistics, Seattle, USA, 2727–2745. <https://aclanthology.org/2022.findings-naacl.209.pdf>
- [48] Enrico Liscio, Roger Lera-Leri, Filippo Bistaffa, Roel IJ. Dobbe, Catholijn M. Jonker, Maite Lopez-Sanchez, Juan A. Rodriguez-Aguilar, and Pradeep K. Murukannaiah. 2023. Value Inference in Sociotechnical Systems: Blue Sky Ideas Track. In *Proceedings of the 22nd International Conference on Autonomous Agents and Multiagent Systems (AAMAS '23)*. IFAAMAS, London, United Kingdom, 1–7.
- [49] Yujian Liu, Xinliang Frederick Zhang, David Wegsman, Nick Beauchamp, and Lu Wang. 2022. POLITICS: pretraining with same-story article comparison for ideology prediction and stance detection. *arXiv preprint arXiv:2205.00619* (2022).
- [50] Leland McInnes, John Healy, and James Melville. 2018. UMAP: Uniform Manifold Approximation and Projection for Dimension Reduction. *arXiv preprint arXiv:1802.03426* (2018).
- [51] Yelena Mejova, Kyriaki Kalimeri, and Gianmarco De Francisci Morales. 2023. Authority without Care: Moral Values behind the Mask Mandate Response. In *Proceedings of the International AAAI Conference on Web and Social Media*, Vol. 17. 614–625.
- [52] Yu Meng, Jiaxin Huang, Yu Zhang, and Jiawei Han. 2022. Generating training data with language models: Towards zero-shot language understanding. *Advances in Neural Information Processing Systems* 35 (2022), 462–477.
- [53] Tomas Mikolov, Ilya Sutskever, Kai Chen, Greg S Corrado, and Jeff Dean. 2013. Distributed representations of words and phrases and their compositionality. *Advances in neural information processing systems* 26 (2013).
- [54] Negar Mokherian, Andrés Abeliuk, Patrick Cummings, and Kristina Lerman. 2020. Moral framing and ideological bias of news. In *Social Informatics: 12th International Conference, SocInfo 2020, Pisa, Italy, October 6–9, 2020, Proceedings 12*. Springer, 206–219.
- [55] Marlon Mooijman, Joe Hoover, Ying Lin, Heng Ji, and Morteza Dehghani. 2018. Moralization in social networks and the emergence of violence during protests. *Nature Human Behaviour* 2, 6 (2018), 389–396. <https://doi.org/10.1038/s41562-018-0353-0>
- [56] Pansy Nandwani and Rupali Verma. 2021. A review on sentiment analysis and emotion detection from text. *Social Network Analysis and Mining* 11, 1 (2021), 81.
- [57] Tuan Dung Nguyen, Ziyu Chen, Nicholas George Carroll, Alasdair Tran, Colin Klein, and Lexing Xie. 2024. Measuring Moral Dimensions in Social Media with Mformer. In *Proceedings of the International AAAI Conference on Web and Social Media*, Vol. 18. 1134–1147.
- [58] Matheus C. Pavan, Vitor G. Santos, Alex G. J. Lan, Joao Martins, Wesley Ramos Santos, Caio Deutsch, Pablo B. Costa, Fernando C. Hsieh, and Ivandrè Paraboni. 2020. Morality Classification in Natural Language Text. *IEEE Transactions on Affective Computing* 3045, c (2020), 1–8. <https://doi.org/10.1109/taffc.2020.3034050>
- [59] Fabian Pedregosa, Gaël Varoquaux, Alexandre Gramfort, Vincent Michel, Bertrand Thirion, Olivier Grisel, Mathieu Blondel, Peter Prettenhofer, Ron Weiss, Vincent Dubourg, et al. 2011. Scikit-learn: Machine learning in Python. *the Journal of machine Learning research* 12 (2011), 2825–2830.
- [60] Vladimir Ponzovskiy, Murat Ardag, Lusine Grigoryan, Ryan Boyd, Henrik Dobe-wall, and Peter Holtz. 2020. Development and Validation of the Personal Values Dictionary: A Theory-Driven Tool for Investigating References to Basic Human Values in Text. *European Journal of Personality* 34, 5 (2020), 885–902. <https://doi.org/10.1002/per.2294>
- [61] Vjosa Prentiqi, Kyriaki Kalimeri, and Charalampos Saitis. 2022. "More Than Words": Linking Music Preferences and Moral Values Through Lyrics. *ISMIR* (2022).
- [62] Vjosa Prentiqi, Kyriaki Kalimeri, and Charalampos Saitis. 2023. Soundscapes of morality: Linking music preferences and moral values through lyrics and audio. *Plos one* 18, 11 (2023), e0294402.
- [63] Shalini Priya, Ryan Sequeira, Joydeep Chandra, and Sourav Kumar Dandapat. 2019. Where should one get news updates: Twitter or Reddit. *Online Social Networks and Media* 9 (2019), 17–29.
- [64] Liang Qiu, Yizhou Zhao, Jinchao Li, Pan Lu, Baolin Peng, Jianfeng Gao, and Song-Chun Zhu. 2022. ValueNet: A New Dataset for Human Value Driven Dialogue System. In *Proceedings of the 36th AAAI Conference on Artificial Intelligence (AAAI '22)*. 11183–11191. <https://doi.org/10.1609/aaai.v36i10.21368>
- [65] Nino Scherrer, Claudia Shi, Amir Feder, and David Blei. 2024. Evaluating the moral beliefs encoded in llms. *Advances in Neural Information Processing Systems* 36 (2024).
- [66] H Andrew Schwartz, Johannes C Eichstaedt, Margaret L Kern, Lukasz Dzurzynski, Stephanie M Ramones, Megha Agrawal, Achal Shah, Michal Kosinski, David Stillwell, Martin EP Seligman, et al. 2013. Personality, gender, and age in the language of social media: The open-vocabulary approach. *PloS one* 8, 9 (2013), e73791.
- [67] Shalom H. Schwartz. 2012. An Overview of the Schwartz Theory of Basic Values. *Online readings in Psychology and Culture* 2, 1 (2012), 1–20. <https://doi.org/10.9707/2307-0919.1116>
- [68] Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajwal Bhargava, Shruti Bhosale, et al. 2023. Llama 2: Open foundation and fine-tuned chat models. *arXiv preprint arXiv:2307.09288* (2023).
- [69] Jackson Trager, Alireza S Ziabari, Aida Mostafazadeh Davani, Preni Golazian, Farzan Karimi-Malekabadi, Ali Omrani, Zhihe Li, Brendan Kennedy, Nils Karl Reimer, Melissa Reyes, et al. 2022. The Moral Foundations Reddit Corpus. *arXiv preprint arXiv:2208.05545* (2022).
- [70] Laura Weidinger, John Mellor, Maribeth Rauh, Conor Griffin, Jonathan Uesato, Po-Sen Huang, Myra Cheng, Mia Glaese, Borja Balle, Atoosa Kasirzadeh, et al. 2021. Ethical and social risks of harm from language models. *arXiv preprint arXiv:2112.04359* (2021).
- [71] Steven R. Wilson, Yiting Shen, and Rada Mihalcea. 2018. Building and Validating Hierarchical Lexicons with a Case Study on Personal Values. In *Proceedings of the 10th International Conference on Social Informatics (SocInfo '18)*. Springer, St. Petersburg, Russia, 455–470. [https://doi.org/10.1007/978-3-030-01129-1\\_28](https://doi.org/10.1007/978-3-030-01129-1_28)
- [72] Carole-Jean Wu, Ramya Raghavendra, Udit Gupta, Bilge Acun, Newsha Ardalani, Kiwan Maeng, Gloria Chang, Fiona Aga, Jinshi Huang, Charles Bai, et al. 2022. Sustainable AI: Environmental Implications, Challenges and Opportunities. *Proceedings of Machine Learning and Systems* 4 (2022), 795–813.
- [73] Jiacheng Ye, Jiahui Gao, Qintong Li, Hang Xu, Jiangtao Feng, Zhiyong Wu, Tao Yu, and Lingpeng Kong. 2022. ZeroGen: Efficient Zero-shot Learning via Dataset Generation. In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, Yoav Goldberg, Zornitsa Kozareva, and Yue Zhang (Eds.), Association for Computational Linguistics, Abu Dhabi, United Arab Emirates, 11653–11669. <https://doi.org/10.18653/v1/2022.emnlp-main.801>
- [74] Jingyan Zhou, Minda Hu, Junan Li, Xiaoying Zhang, Xixin Wu, Irwin King, and Helen Meng. 2023. Rethinking Machine Ethics—Can LLMs Perform Moral Reasoning through the Lens of Moral Theories? *arXiv preprint arXiv:2308.15399* (2023).