Investigating Human Values in Online Communities

Anonymous ACL submission

Abstract

Studying human values is instrumental for cross-cultural research, enabling a better understanding of preferences and behaviour of society at large and communities therein. To study the dynamics of communities online, we propose a method to computationally analyse values present on Reddit. Our method allows analysis at scale, complementing survey based approaches. We train a value relevance and a value polarity classifier, which we thoroughly evaluate using in-domain and out-of-domain human annotations. Using these, we automatically annotate over six million posts across 12k subreddits with Schwartz values. Our analysis unveils both previously recorded and novel insights into the values prevalent within various online communities. For instance, we discover a very negative stance towards conformity in the Vegan and AbolishTheMonarchy subreddits. Additionally, our study of geographically specific subreddits highlights the correlation between traditional values and conservative U.S. states. Though our work, we demonstrate how our dataset and method can be used as a complementary tool for qualitative study of online communication.

1 Introduction

Human Values have been a useful analysis lens for social sciences scholars (Ponizovskiy et al., 2020; Boyd et al., 2021; Schwartz, 1994). Applied to communities or individuals, they are used to study political affiliations, cultural integration, human disagreement, economic growth, and human development, among others (Inglehart, 2020). Many frameworks exist for studying human values, including Rockeach values (Rokeach, 1973), Hofstede's cultural dimensions (Hofstede, 1984), and Moral Foundations Theory (Graham et al., 2013). However, studies using these value frameworks often struggle with sample sizes and rely on self-reported surveys to calculate the values of communities (Weld et al., 2023), leading to concerns about representation and the generalisation of the results to populations beyond the ones studied (Gerlach and Eriksson, 2021; Bašnáková et al., 2016; Fontaine et al., 2008).

Val	Subreddits
AC	r/startups, r/resumes, r/xboxachievements
BE	r/Adoption, r/BPDlovedones, r/Petloss
CO	r/policebrutality, r/HOA, r/BadNeighbors,
HE	r/FreeCompliments, r/transpositive, r/-cozy
PO	r/debtfree, r/geopolitics, r/dividends
SE	r/GunsAreCool, r/worldevents, r/CombatFootage
SD	r/antidepressants r/DebateReligion, r/TrueUnpopularOpinion, r/nutrition
ST	r/crossdressing, r/Hobbies, r/NailArt
TR	r/religion, r/AskAPriest, r/atheism
UN	r/AskFeminists, r/IsraelPalestine, r/climatechange

Table 1: Subreddits with the highest expression of each of the ten Schwartz values. The stance of Green subreddits towards the value is positive (above 0.2), whereas Red indicates negative stance (below -0.2). Blue represents neutral. AC=achievement, BE=benevolence, CO=conformity, HE=hedonism, PO=power, SE=security, SD=self-direction, ST=stimulation, TR=tradition, UN=universalism.

Social media platforms provide unadulterated access to vast and diverse expressions of human thoughts and opinions in the form of posts, discussions, and comments. This rich data source is invaluable for investigating various aspects of human society (Newell et al., 2016; Agrawal et al., 2023; Zomick et al., 2019; Turcan and McKeown, 2019). However, analysing large amounts of social media data qualitatively remains challenging.

To overcome this challenge, in this work, we provide a method for computationally analysing human values in language used on social media at scale. Our method and dataset supplies scholars with a tool for extracting high level values based insights from a large corpus of social media text,

Value	Description
Power	Social status and prestige, control or dominance over people and resources
Achievement	Personal success through demonstrating competence according to social standards.
Hedonism	Pleasure and sensuous gratification for oneself.
Stimulation	Excitement, novelty, and challenge in life.
Self-direction	Independent thought and action-choosing, creating, exploring.
Universalism	Understanding, appreciation, tolerance, and protection for the welfare of all people and for nature.
Benevolence	Preservation and enhancement of the welfare of people with whom one is in frequent personal contact.
Tradition	Respect, commitment, and acceptance of the customs and ideas that traditional culture or religion provide.
Conformity	Restraint of actions, inclinations, and impulses likely to upset or harm others and violate social expectations or norms.
Security	Safety, harmony, and stability of society, of relationships, and of self.

Table 2: The ten Schwartz values and their meaning; descriptions from Schwartz (1994)

which can be used to identify interesting phenomena to qualitatively study when studying online behaviour and communities. Specifically, we follow Schwartz's Theory of Human Values (Schwartz, 1994), due to its wide adoption as a value framework in the social sciences as well as in Natural Language Processing, and apply it on Reddit. Reddit is conceptually based on named communities, or, **subreddits**, and studying the values they exhibit is particularly interesting for studying online behaviour as the discussions are already segregated by topic or sub-communities. For instance, by examining the `r/teenagers` subreddit, we can gain insights into adolescent perspectives, while the `r/ConservativeValues` allows us to study the impact of conservative values on worldviews.

To extract values from Reddit, we train supervised value extraction models to classify the *presence* and *polarity* of Schwartz values in text and apply them to Reddit posts. We then evaluate the models, validating their effectiveness and addressing limitations in prior work. Subsequently, we use these models to infer the expressed values of 9 million user posts and comments across 11,616 most popular subreddits on Reddit.

Through such a computational value analysis, we can analyse digital trace and behaviour data at scale, complementing social science studies by highlighting patterns as well as outliers needing further study. As an instance of this, our analysis demonstrates that people contributing to `r/feminism` exhibit high values in *self-direction*, as substantiated by previous studies. We also demonstrate the prevalence of high traditional values in conservative US states. In sum, our **contributions** are:

- A flexible method of conducting large-scale value relevance and polarity analysis to complement social science research on online communities;
- Our analysis of values at a large scale confirms previously recorded phenomena and unveils novel insights, indicating that our method can complement traditional methods in understanding societal phenomena;
- We release our dataset of 12k subreddits with their corresponding Schwartz values for future analysis.[1]

Although our findings offer unique insights into the values of individuals from various cultures, the focus of our paper is the analysis of values present in *online communities*. Hence the subject of our study is the communities themselves and how they operate, rather than the broader cultural backgrounds of their members, which we do not measure directly. Our methods are designed to complement, not replace, traditional survey-based approaches for studying values, by providing additional perspectives from digital trace data.

2 Related Work and Background

2.1 Schwartz's Values Framework

Values represent a crucial aspect of human nature. According to Schwartz (2012), *"A person's value priority or hierarchy profoundly affects his or her attitudes, beliefs, and traits, making it one core component of personality."* Schwartz (1994) define values as (1) concepts or beliefs, that (2) pertain to desirable end states or behaviors, (3) transcend specific situations, (4) guide selection or evaluation of behavior and events, and (5) are ordered by relative importance. Based on this, they outline ten basic human values: *Security, Conformity, Tradition, Benevolence, Universalism, Self-direction, Stimulation, Hedonism, Achievement, Power*. We provide descriptions of each of the values in Tab. 2. The values were originally applied to measure differences in values across cultures (Schwartz, 1994).

[1] `https://github.com/[anonymised]/[anonymised]`

The framework is suggested as a tool to study populations rather than individual people, hence it serves as a suitable tool for the analysis of online communities. We leverage the fact that a person's identity and values are often reflected in the linguistic choices they make (Jaffe, 2009; Norton, 1997) to analyse values embedded in text.

2.2 Values and Natural Language Processing

Recently, there have been a number of studies exploring values and morals using NLP. In analysing values and language on social media, Ponizovskiy et al. (2020) released a Schwartz value dictionary. Boyd et al. (2021) show the promise of free-text survey response and Facebook data for value extraction. There have also been studies exploring morals and norms in text. Trager et al. (2022) release the Moral Foundations Reddit Corpus with 16k Reddit comments; Roy et al. (2021) similarly study morality in political tweets. Havaldar et al. (2024) study the presence of values in a geolocated Twitter corpus using a lexicon-based approach, finding a lack of correlation with survey data. The closest to our study is van der Meer et al. (2023), who train a value extraction model based on datasets from Qiu et al. (2022); Kiesel et al. (2022). Using the model, they analyse values at an individual user level to understand disagreement in online discussions. However, a key limitations in their work is only detecting the *presence* of a value, neglecting the *polarity* of the discussion towards that value. In our work, we overcome this limitation by training an additional value **stance** model. Further, we perform our analysis at the community level to understand the values of communities rather than individuals, bringing it closer to the original framework outlined by Schwartz designed to understand cultural values.

2.3 Studying Online Communities

Much work has been done to explore Reddit and its user base. Reddit users' personalities have been studied (Gjurković and Šnajder, 2018) as well as their mental health (Zomick et al., 2019; Turcan and McKeown, 2019; Chancellor et al., 2018). Previous work has also studied Reddit by focusing on events affecting community dynamics (Newell et al., 2016; Agrawal et al., 2023). For instance, Chandrasekharan et al. (2022) examine the effect of content moderation on two controversial subreddits. Soliman et al. (2019) study the characteristics and differences between left-leaning and right-leaning political subreddits. Closer to our work, Weld et al. (2023) craft a taxonomy of community values and associate them with specific subreddits. However, their methodology, while high quality, involves collecting values through self-reporting questionnaires of limited size and suffers from selection bias, a limitation we address in this paper. Finally, studies have also explored modelling community norms (Park et al., 2024) and detecting their violations (Cheriyan et al., 2021).

3 Method

This section details the collection and processing of our Reddit dataset, and our approach to training and evaluating the Schwartz values extractor model.

3.1 Data Collection

We download an image of Reddit posts and comments authored between January and August 2022[2] using Pushshift's API.[3] We filter out posts and comments with fewer than ten words or with fewer than ten upvotes to reduce possible noise from low-quality text. We then merge the lists of posts and comments by subreddit, not further distinguishing between them for our analysis, and filter out subreddits that are tagged NSFW,[4] have fewer than 5,000 subscribers, or fewer than 250 content examples. Finally, we down-sample large subreddits to 1,000 random samples due to computational constraints and remove non-English content. This process results in a dataset $\mathcal{D}$ of 11,616 unique subreddits. See Tab. 3 in App. B for dataset statistics.

3.2 Value Extraction

For extracting Schwartz values and their polarity from text, we first extract the *relevance* of a post or comment with each Schwartz value, then use a stance model to extract the *polarity* of the sentence towards the relevant value.

Value Relevance For extracting relevance, our approach is similar to van der Meer et al. (2023) in training a neural Schwartz Values relevance classifier. We use a DeBERTa model (He et al., 2021), over RoBERTa, due its improved performance and speed. As a single post can potentially express several values simultaneously, the model is trained in a *multi-label setting*, predicting a vector of 10

[2] System limitations pose restrictions on the date range we can process – an image of a single month is over 300GB.

[3] https://pushshift.io/

[4] Not Safe For Work

independent probabilities for each input. We fine-tune the classifier on the concatenation of two supervised Schwartz values datasets, ValueNet (Qiu et al., 2022) and ValueArg (Kiesel et al., 2022). Given a labelled triplet (s, v, y), with s being a string of text, v the name of one of the ten Schwartz values (Tab. 2), and $y \in \{0, 1\}$, we construct the input $x =$ `[CLS]` v `[SEP]` s `[SEP]` and train the model to predict $p(y|x)$. Our classifier's performance on this dataset is similar to the figures reported by van der Meer et al. (2023), namely a macro-averaged F_1 score of 0.76 on the merged ValueArg and ValueNet datasets. Similar to van der Meer et al. (2023), we collapse the two non-neutral labels (-1 and 1) in ValueNet into a single positive class.[5] This model is trained, therefore, to predict the *presence* of a value v within a string s.

Value Stance Prior work on value extraction neglegts to model the polarity of content towards values, substantially limiting the insights one can draw. To detect the polarity of a comment or a post s towards a value v, we train a separate stance model on ValueNet's non-neutral labels. Specifically, given a triplet (s, v, y) where v is a value expressed by s and $y \in \{-1, 1\}$, we fit the stance model to predict $p_{\text{stance}}(y|x)$, where x is constructed as described above. We refer readers to App. C.1 for additional training details about the classifiers.

When extracting values of comments or posts from Reddit, we first use the relevance model to predict the probabilities for all Schwartz values for a given input instance, resulting in a vector $\boldsymbol{u}_{\text{rel}} \in [0, 1]^{10}$. Specifically, given a string s, we construct $\boldsymbol{u}_{\text{rel}} = [p(y^1|x^1), ..., p(y^{10}|x^{10})]$ by replacing v^k in the construction of $x^k =$ `[CLS]` v^k `[SEP]` s `[SEP]` with each one of the ten possible Schwartz values. That is, each entry in the vector is the independent probability that s expresses the value v^k, supporting a multi-label approach. Then, for each k with $\boldsymbol{u}^k_{\text{rel}} > 0.5$ (i.e., input text s expresses the value v^k with a probability greater than 0.5), we predict $p_{\text{stance}}(y^k|x^k)$, (i.e., the polarity of input text s towards the value v^k). Thus, we construct a vector of probabilities $\boldsymbol{u}_{\text{stance}}$ of dimensionality 10, where each entry $\boldsymbol{u}^k_{\text{stance}}$ is either in $[-1, 1]$ or Null (if s does not express the value v^k greater than 0.5). Tab. 9 in App. B.3 contains a sample of Reddit content and the associated Value Extractor model predictions.

[5]We do this to follow the format of the ValueArg dataset, which only contains annotations for value relevance.

3.3 Evaluation of Value Extraction

An important limitation of van der Meer et al. (2023) is the lack of direct evaluation of the value extraction model. The need for this becomes particularly pronounced considering the large domain shift of applying the model—which, similarly to us, was trained on debating data—to analyze content on platforms like Reddit. Therefore, we conduct a thorough evaluation to test the model's capabilities of extracting values from Reddit content.

Ideally, we would want to assess the model's performance using a large, annotated dataset of randomly sampled Reddit posts. However, this is impractical because randomly selected posts are unlikely to contain any Schwartz values. Consequently, annotating these posts would yield a dataset with only a few positive examples. To address this challenge, we evaluate the Value Extraction models directly.

Relevance Model Evaluation We evaluate the relevance model by first using it to label 10,000 posts and 10,000 comments for predicting the presence of values. Thereafter, for each value v, we sample three posts or comments: one with high model confidence for the presence of the value (above 0.8), one with medium confidence (0.4-0.6), and one with low confidence (below 0.2). Three annotators then rank these comments/posts based on which are more related to value v, regardless of polarity. We repeat this process five times per value, totalling 50 rankings per annotator. See App. A for annotation guidelines and dataset samples.

For annotator agreement, we calculated averaged Spearman's ρ, which looks at correlation amongst ranks. The agreement we found was 0.63, which is reasonable for such a subjective task.[6] Certain values (e.g., *security*, *universalism*) showed better agreement than others (e.g., *self-direction*). A full breakdown is available in Tab. 5, App. B. For assessing model performance, we assigned a gold pseudo-ranking to each sample by averaging the annotators' rankings. We evaluate the relevance model's performance by comparing the model's predicted ranking to this gold standard. The average Spearman's ρ we obtain is 0.51 (again, with values such as *security* outperforming other values such as *self-direction*). Looking at the top ranked content, the NDCG@1 we obtain is an impressive 0.87, demonstrating that high certainty predictions

[6]Kiesel et al. (2022) found the averaged agreement amongst annotators for value annotation to be $\alpha = 0.49$.

made by our classifier are highly relevant to their corresponding value. With this, we can extract comments or posts relevant to certain values from the larger set. See App. B.2 for examples of model misclassification and value breakdown.

Stance Model Evaluation We first sample 20 posts or comments per value v where the relevance model predicted the presence of v with high confidence. Two annotators then label the samples for stance, choosing between positive, negative, or neutral/unrelated (i.e., where no stance is clearly expressed or the relevance model misclassified the sample). The annotators achieved a Cohen's Kappa score of 0.47, representing moderate agreement, highlighting the challenging nature of the task. However, post discussion on the disagreements, the annotators were able to converge on decisions for each sample, thus obtaining gold labels for the 200 samples. Finally, we apply the stance model to predict the stances of all positive or negative samples. The model predicts the stance towards values with an F_1 score of 0.72. Given the subjective nature of the task, where one has to predict stance towards an abstract concept like human values and one where annotators often disagreed, we believe the model performs reasonably well. To highlight some of the trickier cases where the model is failing, Tab. 8 in App. B.2 contains examples of model misclassification and breakdown into values.

3.4 Assigning Values to Subreddits

We assign each subreddit $\mathcal{S} \in \mathcal{D}$ a single vector of Schwartz probabilities $\boldsymbol{u}_{\text{rel}}(\mathcal{S})$ and a single vector of stances $\boldsymbol{u}_{\text{stance}}(\mathcal{S})$. Given $\mathcal{S}$ with content $(c_i)|i \in \mathcal{S}$, where c_i is either a post or a comment, we predict $\boldsymbol{u}_{\text{rel}}(c_i)$ and $\boldsymbol{u}_{\text{stance}}(c_i)$ from c_i using the process above. Finally, we calculate $\boldsymbol{u}_{\text{rel}}(\mathcal{S})$ by averaging over the predicted vectors $\boldsymbol{u}_{\text{rel}}(c_i)$:

$$\boldsymbol{u}_{\text{rel}}(\mathcal{S}) = \frac{1}{|\mathcal{S}|} \sum_{i \in \mathcal{S}} \boldsymbol{u}_{\text{rel}}(c_i)$$

We construct $\boldsymbol{u}_{\text{stance}}(\mathcal{S})$ similarly, considering only non-Null entries. That is, each entry in $\boldsymbol{u}^k_{\text{stance}}(\mathcal{S})$ is computed as

$$\boldsymbol{u}^k_{\text{stance}}(\mathcal{S}) = \frac{1}{|\mathcal{S}^k|} \sum_{i \in \mathcal{S}^k} \boldsymbol{u}^k_{\text{stance}}(c_i),$$

where $\mathcal{S}^k = \{i \in \mathcal{S} | \boldsymbol{u}^k_{\text{stance}}(c_i) \neq \text{Null}\}$.

4 Experiments

Our experiments serve two primary goals: first, to validate our approach by comparing the values extracted using our method to existing knowledge. Second, to demonstrate the extensibility of our method to new topics that have not been previously investigated in the social sciences. We conduct qualitative and quantitative evaluations of our approach in §4.1 and §4.2, to validate our method. Then, to uncover interesting phenomena and demonstrate the utility of our method, we investigate subreddits with differing opinions on controversial topics and compare our findings with existing studies. Finally, in §4.4, we correlate values extracted using our method to values gained through traditional approaches like surveys and questionnaires.

4.1 Qualitative Analysis

We start by assessing how well our values extraction model performs on our Reddit dataset.

High relevance Intuitively, certain Schwartz values are expected to be distinctly present in specific communities; e.g. *tradition* in religion-related communities. Therefore, we find subreddits with particularly strong signals of individual Schwartz values. For each value, we sort the subreddits' values probability vectors $\boldsymbol{u}_{\text{rel}}(\mathcal{S})$ by their entry corresponding to the particular value. Tab. 1 lists a sample of the top subreddits for each value; Tab. 10 in App. B lists the top 30 subreddits per value. Many of the subreddits collected for each value seem to be intuitively related to it (e.g., `r/resumes` with *achievement*, `r/conservation` with *Universalism*). The results are encouraging, demonstrating the effectiveness of our approach and its potential for conducting interesting analyses.

Strong stances Equivalently, we can also investigate which subreddits express the strongest *stances* towards each value. For each value, we sort the subreddits' stance vectors $\boldsymbol{u}_{\text{stance}}(\mathcal{S})$ by their entry corresponding to the particular value. Tab. 11 in App. B enumerates the 10 subreddits with the highest *positive* and *negative* stance towards each value. Again, while some of the listed subreddits are intuitive (e.g., `r/migraine` having a negative "hedonism" stance and `r/raisingkids` having a strong positive stance towards "achievement"), other subreddits are more surprising (`r/TheHague` with a positive stance towards "hedonism").

Value magnitude We hypothesise that online com-

munities differ not only in the set of values they express but also in the total **magnitude** of expressed values. Online communities pertaining to more objective topics (e.g., linear algebra) should express fewer Schwartz values than communities dedicated to subjective or controversial topics (e.g., politics). To test this, for each subreddit $\mathcal{S} \in \mathcal{D}$ we calculate $\text{mag}(\mathcal{S}) = |\boldsymbol{u}_{\text{rel}}(\mathcal{S})|_2$, the total magnitude of values expressed in the subreddit. The 20 subreddits with the highest and lowest value magnitudes are detailed in Tab. 12 in App. B. Subreddits with particularly high magnitudes are those communities that focus on debating, discussion, or emotional narratives. Examples of these include `r/changemyview`, `r/Adoption`, and `r/PoliticalDiscussion`. Conversely, subreddits with notably low magnitudes are generally more objective and neutral in tone. These include `r/crystalgrowing` and `r/whatisthisfish`. Some subreddits that exhibit low-value magnitudes are dedicated to sharing photos taken by community members, e.g., `r/astrophotography`.

4.2 Quantitative Analysis

We hypothesise that similar subreddit communities share similar values, and systematically investigate this using empirical evidence. We define three measures of similarity between subreddits:

Value similarly. We define $\sigma_{\text{val}}(\mathcal{S}_1, \mathcal{S}_2) = \cos\left(\boldsymbol{u}_{\text{rel}}(\mathcal{S}_1)||\boldsymbol{u}_{\text{stance}}(\mathcal{S}_1), \boldsymbol{u}_{\text{rel}}(\mathcal{S}_2)||\boldsymbol{u}_{\text{stance}}(\mathcal{S}_2)\right)$, the cosine similarity between the concatenation of the relevance and stance vectors of the two subreddits.[7]

Semantic similarity. We scrape the **public description** of each subreddit from its page (see Fig. 4 in App. B), and embed these natural language descriptions using a sentence transformer[8] to construct a semantic embedding vector $\boldsymbol{e}_{\mathcal{S}}$ for each subreddit $\mathcal{S}$[9]. We now define $\sigma_{\text{sem}}(\mathcal{S}_1, \mathcal{S}_2) = \cos(\boldsymbol{e}_{\mathcal{S}1}, \boldsymbol{e}_{\mathcal{S}2})$.

Community similarity For each pair of subreddits, we define their community similarity to be the *overlap coefficient* between the users of the two subreddits. That is,

$$\sigma_{\text{com}}(\mathcal{S}_1, \mathcal{S}_2) = \frac{|U(\mathcal{S}_1) \cap U(\mathcal{S}_2)|}{\min(|U(\mathcal{S}_1)|, |U(\mathcal{S}_2)|)}$$

Where $U(\mathcal{S})$ is the set of the subreddit's members.[10]

To answer the question of if similar subreddits share similar values, we correlate σ_{val} with σ_{sem} and σ_{com} as follows. First, for each subreddit $\mathcal{S}$ we find $\overline{\mathcal{S}}_{\text{sem}} = \text{argmax}_{\mathcal{S}'}\left[\sigma_{\text{sem}}(\mathcal{S}, \mathcal{S}')\right]$ and $\overline{\mathcal{S}}_{\text{com}} = \text{argmax}_{\mathcal{S}'}\left[\sigma_{\text{com}}(\mathcal{S}, \mathcal{S}')\right]$. If similar subreddits share a similar set of values, we should expect $\mathbb{E}_{\mathcal{S}}[\sigma_{\text{val}}(\mathcal{S}, \overline{\mathcal{S}}_{\text{sem}})]$ and $\mathbb{E}_{\mathcal{S}}[\sigma_{\text{val}}(\mathcal{S}, \overline{\mathcal{S}}_{\text{com}})]$ to be significantly larger than $\mathbb{E}\left[\sigma_{\text{val}}(\cdot, \cdot)\right]$. That is, the Schwartz values of two similar subreddits should be significantly closer to each other than the Schwartz values of two random subreddits. Our results, computed using empirical expectations over our dataset, confirm this hypothesis. Significant differences exist between the expected Schwartz values of similar and random subreddits. The expected values for both semantic and community similarity is 0.81 versus 0.64 for random subreddits. The z-test scores are 73.2 and 74.4, respectively, indicating a high level of statistical significance.

4.3 Controversial Topics

Different online communities have differing viewpoints on widely discussed issues based on their values. To test whether we can also observe this in subreddits, we extend the analysis to **controversial topics**, based on Wikipedia's list of controversial topics.[11] For a selected topic, we identify subreddits with differing viewpoints on the topic, e.g., `r/Communism` and `r/Capitalism`.

To establish the utility of our method and dataset, we analyse both topics studied by prior work in the social sciences, as well as previously unexplored topics that could be further studied. An overview of the results on all controversial topics can be found in Figure 1. An overview of closely related subreddits (in terms of values) to each of the subreddits discussed here can be found in Appendix B.5.

Feminism In our investigation involving the Reddit communities of `r/Feminism` and `r/MensRights` (prior work (Khan, 2020; Witt, 2020) characterized `r/MensRights` as a non-feminist community based on the content shared by its members), we find both subreddits to be extremely anti-conformative (Figure 1a). The values of *universalism*, *benevolence*, and *self-direction* were found

[7] We experimented with other formulations for calculating $\sigma_{\text{val}}(\mathcal{S}_1, \mathcal{S}_2)$, and arrived to similar results.

[8] `https://www.sbert.net/`. We use the `all-mpnet-base-v2` pretrained model.

[9] A limitation of this approach is that, occasionally, the public description is rather vague. However, a manual analysis we conducted reveals that this is rare.

[10] As Reddit does not make this information publicly available, we estimate $U(\mathcal{S})$ by defining it as the set of users that had posted or commented in the subreddit within our dataset.

[11] `https://en.wikipedia.org/wiki/Wikipedia:List_of_controversial_issues`

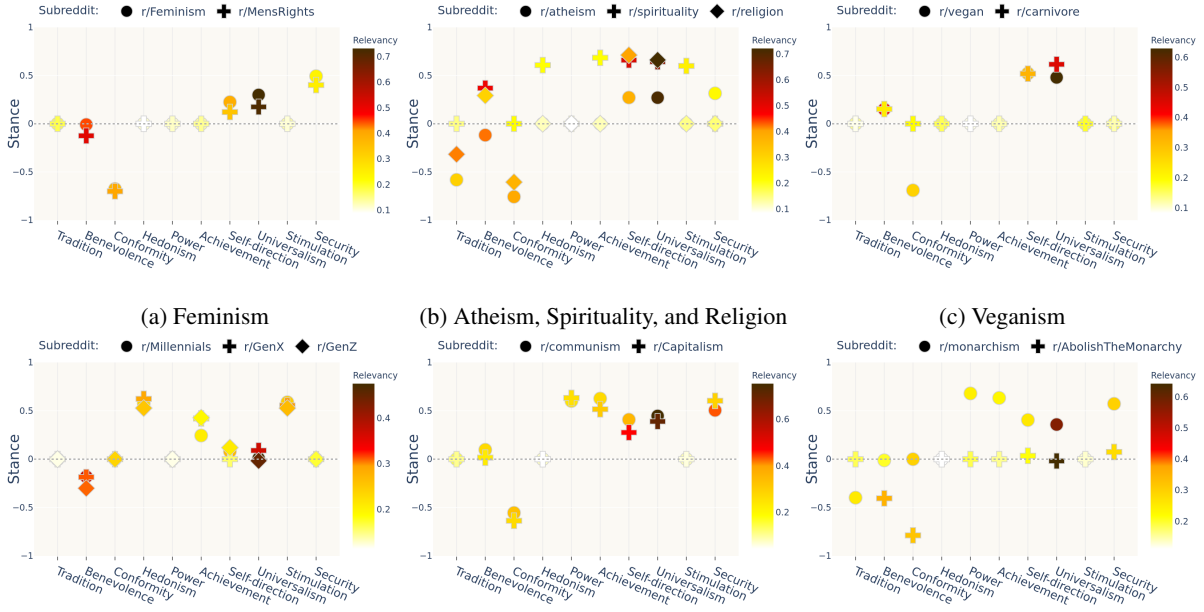

(a) Feminism (b) Atheism, Spirituality, and Religion (c) Veganism

(d) Generations (e) Communism vs Capitalism (f) Monarchism

Figure 1: Radar plots displaying the ten Schwartz values of subreddits dedicated to controversial topics.

to be highly relevant for both. *self-direction* as a value being more core to feminists, as compared to non-feminists has also been found in prior work exploring their values (Zucker and Bay-Cheng, 2010).

Religion There are multiple large-scale studies aiming to establish the connection between religion and Schwartz values. Saroglou et al. (2004) present a meta-analysis across 15 countries, finding that religious people prefer conservative values, e.g., *tradition*, and dislike values related to openness (*stimulation*, *self-direction*). We see similar trends in our comparison of `r/atheism`, `r/spirituality`, and `r/religion` (Figure 1b). While `r/atheism` and `r/religion` express *tradition*, interestingly, both subreddits have a negative stance towards it. Similarly, lower scores for *power*, *hedonism*, and *achievement* for those communities also align with previous findings. Interestingly, `r/spirituality` has high relevance and a very positive stance towards some of those values, demonstrating the focus on the self.

Veganism Figure 1c compares the values for `r/vegan` and `r/carnivore`. The largest difference can be observed in the value of *conformity*. `r/Vegan` has a very negative stance towards it, representing them challenging the status quo. Holler et al. (2021) review studies related to the values of vegetarians/vegans and omnivores, and find that vegetarians were found to have a stronger relation to *universalism*. We find a similar pattern with `r/vegan` having a slightly higher relevance but `r/carnivore` having a more positive stance. They also found them to have a greater emphasis on *self-direction*, which aligns with our findings.

Generations Lyons et al. (2007) study the differences in values across generations. Similarly, we compare `r/Millenials`, `r/GenZ`, and `r/GenX` in Figure 1d. While the values of the different generations' subreddits are highly aligned, small differences can be observed, such as the more positive stance towards *achievement* in GenZ and GenX than Millenials. This contradicts Lyons et al. (2007), which found Millenials to be more achievement focused than GenX.

Communism vs capitalism Here, we describe the difference in values for different economic ideologies, i.e., communism vs capitalism in Figure 1e. While Schwartz and Bardi (1997) describe the effect of communism on Eastern Europe (e.g., high importance to conservatism and hierarchy values), they did not include the values held by the people supporting the ideology on a theoretical level. We find that contributors to `r/communism` hold a more positive stance towards *achievement*, whereas contributors to `r/Capitalism` hold high relevance values for *self-direction* and *security*. Both subreddits have high relevance with *universalism*.

Monarchism We further present values on monarchism in Figure 1f, as another example of a controversial topic yet to be studied in terms of values. We find contributors to `r/AbolishThe-Monarchy` to be against conformity, with a negative stance towards *benevolence*. The contributors to `r/monarchism` have a high relevance to *tradition*, but a negative stance towards it. They also converse a lot more positively about *power*, *achievement*, and *security*.

4.4 Correlation with Surveys

Finally, to understand how well aligned values expressed in online communities are with values of real-world communities, we compare the Schwartz values extracted from Reddit with those obtained from traditional questionnaire methods.

US States First, we use our method to investigate the premise that conservative US states are more traditional than liberal states. We correlate the *tradition* relevance value extracted using our method from the states' subreddits with responses to a survey on conservative ideologies across U.S. states (Center, 2014b), finding a Spearman's R of 0.55 (p-value < 0.0001). Moreover, when correlating these values with a survey on state religiosity levels (Center, 2014a), we find a Spearman's R of 0.63 (p-value < 0.0001). Fig. 5 in App. B.4 displays the *tradition* value of the 50 states subreddits, colour-coded by the recent US election results – "red" states clearly tend to have higher *tradition* values than "blue" states.

Country-level values Next, we extend this correlation experiment to additional Schwartz values and countries other than the USA. Schwartz and Cieciuch (2022) used questionnaires to analyse the Schwartz values associated with 49 countries,[12] which we correlate with values extracted from subreddits related to these countries using our classifier. We start by manually identifying a set of subreddits relevant to each of the 49 countries.[13] We successfully found at least one subreddit for 41 countries. We also determined which language is the most likely to be used online by citizens of the country (and not, for example, by tourists).[14] See Tab. 13 in App. B for a comprehensive list of countries and their associated languages and subreddits. Next, we collected posts written in the country's language[15] from the country's subreddits. We then randomly sampled 2,000 posts and 2,000 comments per country, excluding countries with fewer than 250 total samples. Listing 3 lists the 32 countries that passed this filter. Finally, we used Google Translate API to translate the posts into English[16] and applied our approach to extract Schwartz values for each country.

The above-mentioned process generates a country-value matrix of dimensions 32×10 of Schwartz values obtained from Reddit. We compare this matrix to the 49×10 the country-value described in Schwartz and Cieciuch (2022) by removing rows corresponding to countries that do not exist in Listing 3. We then correlate the columns of the two now-identical matrices using Spearman's ρ. In line with previous studies (Nasif et al., 1991; Joseph et al., 2021), we find no correlation between the values extracted from SM and values extracted from questionnaires—an average Spearman's ρ of -0.03. Tab. 4 in App. B.4 reports the results in full. This lack of correlation may arise from the distinct demographics of Reddit users compared to questionnaire respondents or the unique dynamics of online interaction, which tends to be more reactionary compared to survey responses (Joseph et al., 2021). Our results confirm findings from previous studies that dynamics of online behaviour differs from offline, and should not directly be used as a proxy without further qualitative investigation.

5 Conclusion

In this paper, we apply Schwartz's Theory of Human Values to Reddit at scale, offering scholars a complementary tool for studying online communities. We train and thoroughly evaluate supervised value extraction models to detect the presence and polarity towards human values expressed in language used on social media. Using these, we conduct extensive analysis of nine million posts across 12k popular subreddits. Through our analysis, we both confirm previous findings from the social sciences, as well as uncover novel insights into the values expressed within various online communities, shedding light on existing patterns and outliers that warrant further investigation.

[12] While the survey is conducted with various populations and regions, here we use the term "countries" for simplicity.

[13] We searched for country-related subreddits online and filtered based on their public description and front-page posts, disregarding ones that tourists or non-residents majorly use.

[14] We acknowledge that this is a simplification, as in many countries, more than one language is commonly spoken.

[15] We used `lingua` (`https://github.com/pemistahl/lingua-py`) to detect the language of each post.

[16] We manually scanned a sample of the translations and deemed them as having satisfactory quality.

Limitations

The main limitations of our work stem from the nature of the task itself. Given the inherent subjectivity and complexity of assigning human values to texts, a certain level of noise—aleatoric uncertainty—is unavoidable. This unpredictability can be amplified by epistemic uncertainty, which arises from the limitations of the models we trained. Although we have thoroughly validated the labels generated by the model (§3.3), the predictions can sometimes be noisy or inconsistent.

Another limitation of our approach is the consolidation of the values of posts in a subreddit into a single vector. While this is necessary for understanding the values of communities, it discounts the values the individual posts of users that constitute these communities. This also opens up much room for future research, particularly in studying the internal dynamics of online communities and understanding the role that values play at the individual post level.

Lastly, while our approach can be leveraged to identify phenomena worthy of further investigation, it lacks the means to explain these observations fully. Values are an abstract concept that is challenging to quantify and analyse, with the values frameworks themselves drawing criticism (Jackson, 2020). We argue that interdisciplinary collaboration with experts in psychology and sociology is essential to understand these phenomena and their implications properly. Such collaborations will not only enrich our understanding of online communities but also contribute to developing more robust and nuanced machine learning models in the future, given the inclusion of such text in the training data.

Broader Impact and Ethical Considerations

We believe our work goes beyond the methodological toolkit of social scientists. The accessibility, large scale, and relatively high quality of Reddit data have positioned it as a valuable resource for training data for diverse NLP tasks (Overbay et al., 2023; Blombach et al., 2020; Huryn et al., 2022). It opens the door to training models that are better attuned to the specific needs of various communities, potentially protected ones. Recognizing the values of online communities, often utilized as training data for LM development, is crucial for making informed decisions about incorporating the data and contributing to the creation of models that genuinely reflect the diverse perspectives within these communities (Arora et al., 2023). These also percolate into downstream applications of LMs (Jakesch et al., 2023) and have a broader impact on their use.

As for ethical considerations, we made specific efforts to ensure that our work does not impair the privacy and anonymity of Reddit users (Sugiura et al., 2017). We refrain from attributing values to individual users and instead study communities as a whole by aggregating and condensing individual data points into a single vector. Nevertheless, many online communities were created to serve as a safe space for vulnerable individuals, where they share highly sensitive and private information. Therefore, research on Reddit and other online communities should make utmost efforts to respect these spaces and handle their data with care.

References

Pratik Agrawal, Tolga Buz, and Gerard de Melo. 2023. Wallstreetbets beyond gamestop, yolos, and the moon: The unique traits of reddit's finance communities. In *AMCIS 2022*.

Arnav Arora, Lucie-aimée Kaffee, and Isabelle Augenstein. 2023. Probing pre-trained language models for cross-cultural differences in values. In *Proceedings of the First Workshop on Cross-Cultural Considerations in NLP (C3NLP)*, pages 114–130, Dubrovnik, Croatia. Association for Computational Linguistics.

Jana Bašnáková, Ivan Brezina, and Radomír Masaryk. 2016. Dimensions of culture: The case of slovakia as an outlier in hofstede's research. *Ceskoslovenska Psychologie*, 60(1).

Andreas Blombach, Natalie Dykes, Philipp Heinrich, Besim Kabashi, and Thomas Proisl. 2020. A corpus of German Reddit exchanges (GeRedE). In *Proceedings of the 12th Language Resources and Evaluation Conference*, pages 6310–6316, Marseille, France. European Language Resources Association.

Ryan Boyd, Steven Wilson, James Pennebaker, Michal Kosinski, David Stillwell, and Rada Mihalcea. 2021. Values in words: Using language to evaluate and understand personal values. *Proceedings of the International AAAI Conference on Web and Social Media*, 9(1):31–40.

Pew Research Center. 2014a. How religious is your state?

Pew Research Center. 2014b. Political ideology by state.

Stevie Chancellor, Andrea Hu, and Munmun De Choudhury. 2018. Norms matter: Contrasting social support around behavior change in online weight loss

communities. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*, CHI '18, page 1–14, New York, NY, USA. Association for Computing Machinery.

Eshwar Chandrasekharan, Shagun Jhaver, Amy Bruckman, and Eric Gilbert. 2022. Quarantined! examining the effects of a community-wide moderation intervention on reddit. *ACM Trans. Comput.-Hum. Interact.*, 29(4).

Jithin Cheriyan, Bastin Tony Roy Savarimuthu, and Stephen Cranefield. 2021. Norm violation in online communities – a study of stack overflow comments. In *Coordination, Organizations, Institutions, Norms, and Ethics for Governance of Multi-Agent Systems XIII*, pages 20–34, Cham. Springer International Publishing.

Johnny R. J. Fontaine, Ype H. Poortinga, Luc Delbeke, and Shalom H. Schwartz. 2008. Structural equivalence of the values domain across cultures: Distinguishing sampling fluctuations from meaningful variation. *Journal of Cross-Cultural Psychology*, 39(4):345–365.

Philipp Gerlach and Kimmo Eriksson. 2021. Measuring cultural dimensions: external validity and internal consistency of hofstede's vsm 2013 scales. *Frontiers in Psychology*, 12:662604.

Matej Gjurković and Jan Šnajder. 2018. Reddit: A gold mine for personality prediction. In *Proceedings of the Second Workshop on Computational Modeling of People's Opinions, Personality, and Emotions in Social Media*, pages 87–97, New Orleans, Louisiana, USA. Association for Computational Linguistics.

Jesse Graham, Jonathan Haidt, Sena Koleva, Matt Motyl, Ravi Iyer, Sean P. Wojcik, and Peter H. Ditto. 2013. Chapter two - moral foundations theory: The pragmatic validity of moral pluralism. In Patricia Devine and Ashby Plant, editors, *Advances in Experimental Social Psychology*, volume 47 of *Advances in Experimental Social Psychology*, pages 55–130. Academic Press.

Shreya Havaldar, Salvatore Giorgi, Sunny Rai, Thomas Talhelm, Sharath Chandra Guntuku, and Lyle Ungar. 2024. Building knowledge-guided lexica to model cultural variation. In *Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, pages 211–226, Mexico City, Mexico. Association for Computational Linguistics.

Pengcheng He, Jianfeng Gao, and Weizhu Chen. 2021. Debertav3: Improving deberta using electra-style pretraining with gradient-disentangled embedding sharing.

Geert Hofstede. 1984. *Culture's consequences: International differences in work-related values*, volume 5. sage.

Sophie Holler, Holger Cramer, Daniela Liebscher, Michael Jeitler, Dania Schumann, Vijayendra Murthy, Andreas Michalsen, and Christian S Kessler. 2021. Differences between omnivores and vegetarians in personality profiles, values, and empathy: a systematic review. *Frontiers in psychology*, 12:579700.

Daniil Huryn, William M. Hutsell, and Jinho D. Choi. 2022. Automatic generation of large-scale multi-turn dialogues from Reddit. In *Proceedings of the 29th International Conference on Computational Linguistics*, pages 3360–3373, Gyeongju, Republic of Korea. International Committee on Computational Linguistics.

Ronald Inglehart. 2020. *Modernization and postmodernization: Cultural, economic, and political change in 43 societies*. Princeton university press.

Terence Jackson. 2020. The legacy of geert hofstede.

Alexander Jaffe. 2009. Stance: Sociolinguistic perspectives. In *Stance: Sociolinguistic Perspectives*.

Maurice Jakesch, Advait Bhat, Daniel Buschek, Lior Zalmanson, and Mor Naaman. 2023. Co-writing with opinionated language models affects users' views. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*, CHI '23, New York, NY, USA. Association for Computing Machinery.

Kenneth Joseph, Sarah Shugars, Ryan Gallagher, Jon Green, Alexi Quintana Mathé, Zijian An, and David Lazer. 2021. (mis)alignment between stance expressed in social media data and public opinion surveys. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 312–324, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.

Abeer Khan. 2020. Reddit mining to understand gendered movements. In *EDBT/ICDT Workshops*.

Johannes Kiesel, Milad Alshomary, Nicolas Handke, Xiaoni Cai, Henning Wachsmuth, and Benno Stein. 2022. Identifying the human values behind arguments. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 4459–4471, Dublin, Ireland. Association for Computational Linguistics.

Sean T Lyons, Linda Duxbury, and Christopher Higgins. 2007. An empirical assessment of generational differences in basic human values. *Psychological reports*, 101(2):339–352.

Ercan G Nasif, Hamad Al-Daeaj, Bahman Ebrahimi, and Mary S Thibodeaux. 1991. Methodological problems in cross-cultural research: An updated review. *MIR: Management International Review*, pages 79–91.

Edward Newell, David Jurgens, Haji Saleem, Hardik Vala, Jad Sassine, Caitrin Armstrong, and Derek Ruths. 2016. User migration in online social networks: A case study on reddit during a period of

community unrest. In *Proceedings of the International AAAI Conference on Web and Social Media*, volume 10, pages 279–288.

Bonny Norton. 1997. Language, identity, and the ownership of english. *TESOL Quarterly*, 31(3):409–429.

Keighley Overbay, Jaewoo Ahn, Fatemeh Pesaran zadeh, Joonsuk Park, and Gunhee Kim. 2023. mRedditSum: A multimodal abstractive summarization dataset of Reddit threads with images. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 4117–4132, Singapore. Association for Computational Linguistics.

Chan Young Park, Shuyue Stella Li, Hayoung Jung, Svitlana Volkova, Tanushree Mitra, David Jurgens, and Yulia Tsvetkov. 2024. Valuescope: Unveiling implicit norms and values via return potential model of social interactions.

Vladimir Ponizovskiy, Murat Ardag, Lusine Grigoryan, Ryan Boyd, Henrik Dobewall, 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):885–902.

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 AAAI Conference on Artificial Intelligence*, volume 36, pages 11183–11191.

Milton Rokeach. 1973. *The Nature of Human Values.* The Nature of Human Values. Free Press, New York, NY, US.

Shamik Roy, Maria Leonor Pacheco, and Dan Goldwasser. 2021. Identifying morality frames in political tweets using relational learning. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 9939–9958, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.

Vassilis Saroglou, Vanessa Delpierre, and Rebecca Dernelle. 2004. Values and religiosity: A meta-analysis of studies using schwartz's model. *Personality and individual differences*, 37(4):721–734.

Shalom H. Schwartz. 1994. Are there universal aspects in the structure and contents of human values? *Journal of Social Issues*, 50(4):19–45.

Shalom H. Schwartz. 2012. An overview of the schwartz theory of basic values. *Online Readings in Psychology and Culture*, 2:11.

Shalom H Schwartz and Anat Bardi. 1997. Influences of adaptation to communist rule on value priorities in eastern europe. *Political psychology*, 18(2):385–410.

Shalom H Schwartz and Jan Cieciuch. 2022. Measuring the refined theory of individual values in 49 cultural groups: psychometrics of the revised portrait value questionnaire. *Assessment*, 29(5):1005–1019.

Ahmed Soliman, Jan Hafer, and Florian Lemmerich. 2019. A characterization of political communities on reddit. In *Proceedings of the 30th ACM Conference on Hypertext and Social Media*, HT '19, page 259–263, New York, NY, USA. Association for Computing Machinery.

Lisa Sugiura, Rosemary Wiles, and Catherine Pope. 2017. Ethical challenges in online research: Public/private perceptions. *Research Ethics*, 13(3-4):184–199.

Jackson Trager, Alireza S. Ziabari, Aida Mostafazadeh Davani, Preni Golazizian, Farzan Karimi-Malekabadi, Ali Omrani, Zhihe Li, Brendan Kennedy, Nils Karl Reimer, Melissa Reyes, Kelsey Cheng, Mellow Wei, Christina Merrifield, Arta Khosravi, Evans Alvarez, and Morteza Dehghani. 2022. The moral foundations reddit corpus.

Elsbeth Turcan and Kathy McKeown. 2019. Dreaddit: A Reddit dataset for stress analysis in social media. In *Proceedings of the Tenth International Workshop on Health Text Mining and Information Analysis (LOUHI 2019)*, pages 97–107, Hong Kong. Association for Computational Linguistics.

Michiel van der Meer, Piek Vossen, Catholijn Jonker, and Pradeep Murukannaiah. 2023. Do differences in values influence disagreements in online discussions? In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 15986–16008, Singapore. Association for Computational Linguistics.

Galen Weld, Amy X. Zhang, and Tim Althoff. 2023. Making online communities 'better': A taxonomy of community values on reddit.

Taisto D. Witt. 2020. *Affect, identity, and conceptualizations of feminism and feminists on the antifeminist r/MensRights subreddit*. Ph.D. thesis, University of British Columbia.

Jonathan Zomick, Sarah Ita Levitan, and Mark Serper. 2019. Linguistic analysis of schizophrenia in Reddit posts. In *Proceedings of the Sixth Workshop on Computational Linguistics and Clinical Psychology*, pages 74–83, Minneapolis, Minnesota. Association for Computational Linguistics.

Alyssa N Zucker and Laina Y Bay-Cheng. 2010. Minding the gap between feminist identity and attitudes: The behavioral and ideological divide between feminists and non-labelers. *Journal of personality*, 78(6):1895–1924.

A Annotation Guidelines

```
For each three consecutive rows in
the table (colour-coded), start by
reading the three sentences in
column B. Next, review the Schwartz
value associated with the three
sentences (listed in column C) and
ensure you understand its meaning,
using the figure at the top of the
table as a reference. Then, order
the sentences according to the
extent to which they express the
value, ignoring the stance of the
sentences towards the value. In
other words, our focus is solely on
whether the value is expressed, not
on how it is expressed. Even if a
sentence violates the value, it
still expresses it. For instance,
the sentence ``I never went to
church'' expresses the value of
``Tradition,'' despite
contradicting it.


This task is subjective, and in
many cases, there is no single
correct answer. If you are
uncertain about a particular
instance, respond to the best of
your ability.
```

Listing 1: Guidelines for the task of annotating the dataset used to evaluate the relevancy model.

```
For each row in the table, start by
reading the sentence in column A.
Next, review the Schwartz value
associated with the sentence in
column B and ensure you understand
its meaning, using the figure at
the top of the table as a
reference. Then, determine the
stance of the sentence towards the
value. If the sentence supports the
value (i.e., the value is expressed
in a positive way, or the sentence
describes a situation where the
value is maintained), select
``positive''. If the sentence
violates the value (i.e., the value
is expressed negatively, or the
sentence describes a situation
where the value is disregarded or
broken), select ``negative''. If
the sentence simply does not
express the value, select ``N\A''.
For instance, the sentence ``I
never went to church'' has a
negative stance towards the value
of ``Tradition''. Conversly, the
sentence ``family comes first for
her'' has a positive stance towards
the value.
```

	Avg	std	Percentile			Total
			25	50	75	
c.p.s	765.4	264.8	517	918	996	8,888,535
w.p.c	62.8	148.0	15	26	58	558,327,230

Table 3: Statistics of the Reddit dataset we analyse. c.p.s is content item (posts and comments) per subreddit, and w.p.c is word per content item.

AC	BE	CO	HE	PO	SE	SD	ST	TR	UN
-0.16	-0.31	-0.04	0.04	-0.21	0.10	-0.12	0.16	0.22	0.01

Table 4: Spearman correlation between the Schwartz values obtained from a questionnaire and the values extracted from Reddit.

```
This task is subjective, and in
many cases, there is no single
correct answer. If you are
uncertain about a particular
instance, respond to the best of
your ability.
```

Listing 2: Guidelines for the task of annotating the dataset used to evaluate the stance model.

As described in §3.3, three authors of this paper annotated Reddit content to evaluate the Value Extraction Models. To evaluate the relevance model, each annotator ranked triplets (s_1, s_2, s_3) according to the extent each sentence s_i expresses a given value v. In total, each annotator ranked 50 triplets, or 5 triplets per value. Fig. 2 contains a print screen of the annotation task, whereas Listing 1 specifies the annotation guidelines.

To evaluate the stance model, each annotator was given a sentence s and a value v, and was tasked to determine the stance of s towards v (one of "positive", "negative", or "irrelevant/neutral", indicating that sentence s does not express the value v). Each annotator annotated a total of 200 sentences, or 20 per value. Fig. 3 contains a print screen of the annotation task, whereas Listing 2 specifies the annotation guidelines.

B Additional Material

B.1 Dataset Statistics

Tab. 3 displays the statistics of the Reddit dataset we analyse.

B.2 Annotation and Evaluation

The per-value agreements between the annotators are detailed in Tab. 5 (relevance model) and Tab. 6 (stance model). See additional details in §3.3 and App. A. Moreover, Tab. 7 and Tab. 8 present cases

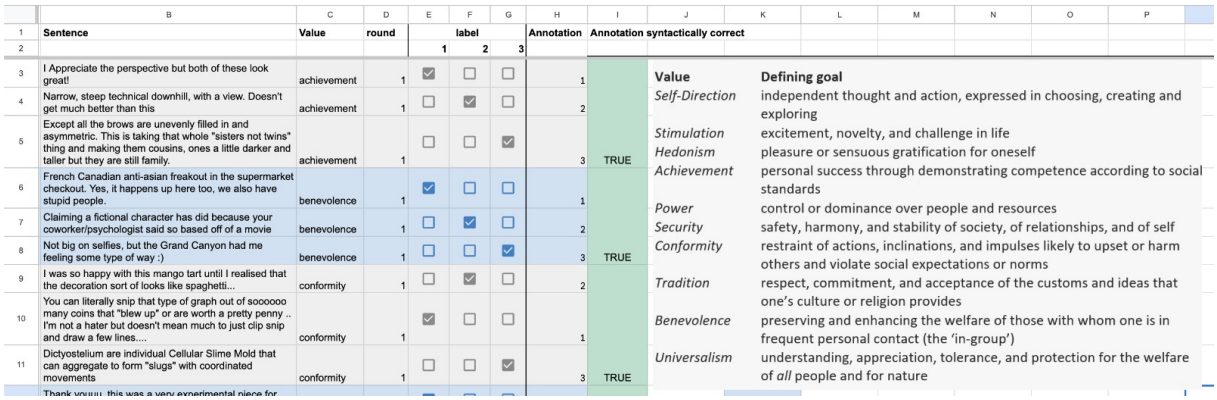

	Sentence	Value	round	label 1	label 2	label 3	Annotation	Annotation syntactically correct
3	I Appreciate the perspective but both of these look great!	achievement	1	☑	☐	☐	1	
4	Narrow, steep technical downhill, with a view. Doesn't get much better than this	achievement	1	☐	☑	☐	2	
5	Except all the brows are unevenly filled in and asymmetric. This is taking that whole "sisters not twins" thing and making them cousins, ones a little darker and taller but they are still family.	achievement	1	☐	☐	☑	3	TRUE
6	French Canadian anti-asian freakout in the supermarket checkout. Yes, it happens up here too, we also have stupid people.	benevolence	1	☑	☐	☐	1	
7	Claiming a fictional character has did because your coworker/psychologist said so based off of a movie	benevolence	1	☐	☑	☐	2	
8	Not big on selfies, but the Grand Canyon had me feeling some type of way :)	benevolence	1	☐	☐	☑	3	TRUE
9	I was so happy with this mango tart until I realised that the decoration sort of looks like spaghetti...	conformity	1	☐	☑	☐	2	
10	You can literally snip that type of graph out of soooooo many coins that "blew up" or are worth a pretty penny .. I'm not a hater but doesn't mean much to just clip snip and draw a few lines....	conformity	1	☑	☐	☐	1	
11	Dictyostelium are individual Cellular Slime Mold that can aggregate to form "slugs" with coordinated movements	conformity	1	☐	☐	☑	3	TRUE

Value	Defining goal
Self-Direction	independent thought and action, expressed in choosing, creating and exploring
Stimulation	excitement, novelty, and challenge in life
Hedonism	pleasure or sensuous gratification for oneself
Achievement	personal success through demonstrating competence according to social standards
Power	control or dominance over people and resources
Security	safety, harmony, and stability of society, of relationships, and of self
Conformity	restraint of actions, inclinations, and impulses likely to upset or harm others and violate social expectations or norms
Tradition	respect, commitment, and acceptance of the customs and ideas that one's culture or religion provides
Benevolence	preserving and enhancing the welfare of those with whom one is in frequent personal contact (the 'in-group')
Universalism	understanding, appreciation, tolerance, and protection for the welfare of *all* people and for nature

Figure 2: Print screen of the annotation task of the relevance model. The three annotators were tasked with ranking triplets of sentences according to the extent to which they express the value of column C.

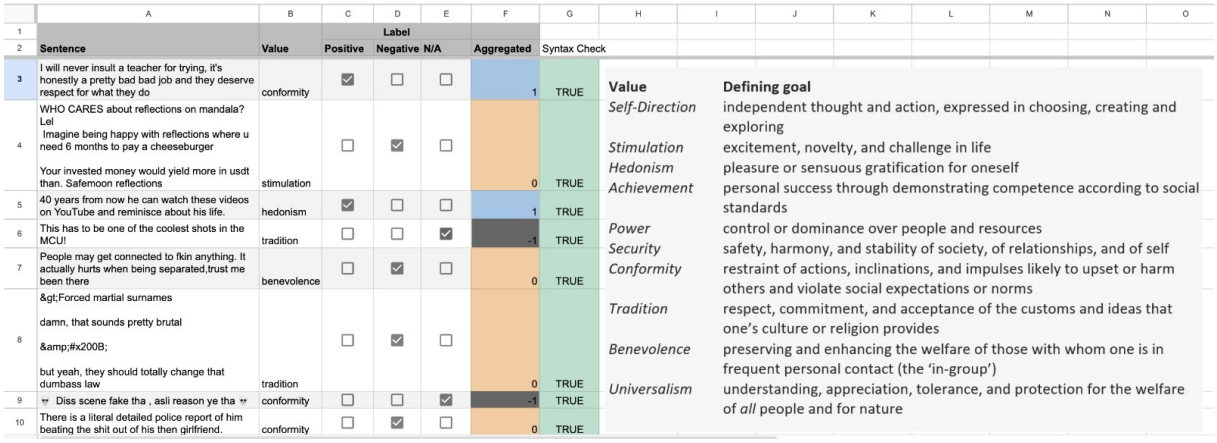

	Sentence	Value	Label: Positive	Label: Negative	Label: N/A	Aggregated	Syntax Check
3	I will never insult a teacher for trying, it's honestly a pretty bad bad job and they deserve respect for what they do	conformity	☑	☐	☐	1	TRUE
4	WHO CARES about reflections on mandala? Lel Imagine being happy with reflections where u need 6 months to pay a cheeseburger Your invested money would yield more in usdt than. Safemoon reflections	stimulation	☐	☑	☐	0	TRUE
5	40 years from now he can watch these videos on YouTube and reminisce about his life.	hedonism	☑	☐	☐	1	TRUE
6	This has to be one of the coolest shots in the MCU!	tradition	☐	☐	☑	-1	TRUE
7	People may get connected to fkin anything. It actually hurts when being separated,trust me been there	benevolence	☐	☑	☐	0	TRUE
8	>Forced martial surnames damn, that sounds pretty brutal but yeah, they should totally change that dumbass law	tradition	☐	☑	☐	0	TRUE
9	👀 Diss scene fake tha , asli reason ye tha 👀	conformity	☐	☐	☑	-1	TRUE
10	There is a literal detailed police report of him beating the shit out of his then girlfriend.	conformity	☐	☑	☐	0	TRUE

Value	Defining goal
Self-Direction	independent thought and action, expressed in choosing, creating and exploring
Stimulation	excitement, novelty, and challenge in life
Hedonism	pleasure or sensuous gratification for oneself
Achievement	personal success through demonstrating competence according to social standards
Power	control or dominance over people and resources
Security	safety, harmony, and stability of society, of relationships, and of self
Conformity	restraint of actions, inclinations, and impulses likely to upset or harm others and violate social expectations or norms
Tradition	respect, commitment, and acceptance of the customs and ideas that one's culture or religion provides
Benevolence	preserving and enhancing the welfare of those with whom one is in frequent personal contact (the 'in-group')
Universalism	understanding, appreciation, tolerance, and protection for the welfare of *all* people and for nature

Figure 3: Print screen of the annotation task of the stance model. Two annotators were tasked to determine the stance of the sentence in column A towards the value in column B.

	AC	BE	CO	HE	PO	SE	SD	ST	TR	UN
Spearman ρ	0.73	0.66	0.66	0.6	0.56	0.8	0.46	0.5	0.6	0.73
NDCG@1	0.89	0.81	0.93	0.96	0.76	0.93	0.74	0.86	0.96	0.84

Table 5: Average Spearman's ρ and NDCG@1 between the three annotators of the relevance model, per-value breakdown. AC=achievement, BE=benevolence, CO=conformity, HE=hedonism, PO=power, SE=security, SD=self-direction, ST=stimulation, TR=tradition, UN=universalism.

AC	BE	CO	HE	PO	SE	SD	ST	TR	UN
0.51	0.57	0.61	0.77	0.30	0.45	-0.27	0.77	0.67	0.26

Table 6: Cohen's Kappa between the two annotators of the stance model, per-value breakdown.

Val	Sentence	Annotated Average Rank
SD	idk it brings some discussions to the community, can't be bad	2.33
PO	You really gave it all you had very impressive content	2.66
CO	Turkish guy debates Marxism and Freud with a street upper	2.33
BE	Claiming a fictional character has did because your coworker/psychologist said so based off of a movie	2.33
PO	What is each of your favorite song from Mercurial World to perform live?	2.66

Table 7: Instances where the relevance model predicted that the sentence expresses a value with high confidence (above 0.8), whereas the annotators concluded that the sentence does not express the value.

Val	Sentence	Predicted (wrong) Stance
PO	Crypto traders using the war to make money. This is a new low	Pos
HE	God forbid you pay someone for their content that you enjoy.	Neg
CO	SLPT: drive at speed in reverse gear through speed cameras to get fines repaid by the police	Pos
SD	If the island keeps drifting, the UK would be part of the US again?	Neg
UN	The double standards of western media towards third world countries	Pos
SD	The coffee is probably making it worse. People forget it's a drug with consequences.	Neg

Table 8: Instances where the stance model misclassified a sentence.

where the relevance model and stance models, respectively, misclassified instances.

B.3 Prediction Examples

Tab. 9 contains examples of posts and comments from Reddit and the values and stances that the relevance and stance models predicted for each instance.

B.4 Experiments

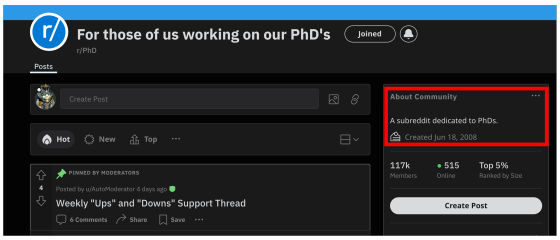

Figure 4: Location of the "Public description" attribute on a subreddit page.

```
Turkey, South Africa, Estonia,
Romania, USA, Costa Rica, Poland,
New Zealand, Brazil, Greece,
Finland, Ukraine, Croatia,
Colombia, Slovakia, United Kingdom,
Iceland, Czech Republic, Italy,
Sweden, China, Ecuador, Germany,
Peru, Indonesia, Serbia, Spain,
India, Australia, Philippines,
Portugal, France,
```

Listing 3: List of countries used in the correlation experiment in §4.4.

Tab. 10 Subreddits with the highest signal for each one of the ten Schwartz values. The stance of Green subreddits towards the value is positive (above 0.2), whereas Red indicates negative stance (below -0.2).

Tab. 11 Subreddits expressing the strongest positive and negative stances for each value.

Tab. 12 Subreddits with the highest and lowest total magnitude of Schwartz values, calculated according to $\text{mag}(\mathcal{S}) = |\boldsymbol{u}_{\text{rel}}(\mathcal{S})|_2$.

Fig. 4 Location of the "Public description" attribute on a subreddit page, used to calculate $\sigma_{\text{sem}}(\mathcal{S}_1, \mathcal{S}_2)$ in §4.2.

Fig. 5 US states sorted by their *tradition* values extracted from Reddit, colour coded based on the recent (2020) US election results (Blue – democratic majority. Red – republican majority.

Listing 3 List of countries used in the correlation experiment in §4.4.

B.5 Controversial Topics – Extended

Additionally to the analysis presented in Section 4.3, Table 14 displays the ten most similar subreddits in terms of values. We use value similarity as described in Section 4.2 to determine the closest subreddits to each of the controversial topic subreddits.

C Reproducibility

C.1 Training the Schwartz Values Extractor

We trained the relevance model and the stance model for ten epochs with early stopping, using a learning rate of $5 \cdot 1e^{-5}$, batch size of 32, Adamw optimiser with default parameters and linear learning rate scheduler. We trained the models on a single TitanRTX GPU for about 5 hours for the relevance model, and 2 hours for the stance model.

Sentence	Predicted Value	Predicted Stance
The green party was literally removed from the ballot in key battleground states in the 2020 election.	Power	Negative
When you're trained by a leader you become a leader.	Power	Positive
I've been learning jazz for about 8 months and still can't write a single good peace, is that normal?Just wondering if Im not deaf because every time I'm trying to write something I end up raging. In general I make music for 4-5 years if that matters.	Achievement	Negative
My first ever attempt at finger crochet (or any crochet for that matter)! I'm really proud of myself and just wanted to share somewhere :)	Achievement	Positive
I have never had a birthday celebrationSo my parents are Jehovah's witnesses I'm not but I have never received a birthday present and this year I'm really depressed I want to celebrate but I'm broke	Hedonism	Negative
To love more is Beautiful.I wish you Wonderful evening and happy weekend	Hedonism	Positive
This video gave me anxiety. It's so inefficient with its movements that it made me squirm.	Stimulation	Negative
This awe inspiring film is the culmination of 3 months working underwater, diving, exploring and filming the most remote corners of the Great Barrier Reef. It showcases some of the most stunning scenes, schooling sharks and most vibrant coral reef the ocean has to offer. Hope you enjoy!	Stimulation	Positive
Having to lie so that you don't have to work Saturdays is funny and wholesome smh	Self-direction	Negative
If you want to be better at squatting, this is fine. If you are going for a specific aesthetic, evaluate/critique your physique for any lagging parts areas, then choose the exercises that help bring them up.	Self-direction	Positive
"The only moral abortion is MY abortion" or "Grifters gonna Grift"	Universalism	Negative
I REALLY THINK ALL OF US HUMANS NEED TO BE REPRESENTED PROPERLY, REGARDLESS OF WHAT LANGUAGE PACKS WE HAVE INSTALLED.	Universalism	Positive
And then blaming the teacher when half the class fails.	Benevolence	Negative
please ban mentions of hasan off of LSF. his mental health is being affected by the harassment from this sub.	Benevolence	Positive
Respect my religion while I shove it down your throat....	Tradition	Negative
Korobeiniki: known in the West as "Tetris music"; in Russian, itś a folk song about courtship.	Tradition	Positive
There is a literal detailed police report of him beating the shit out of his then girlfriend.	Conformity	Negative
Indian minister drinks dirty water from 'holy' river polluted with sewage to show locals it's safe... ends up in hospital days later	Conformity	Positive
This is in Mexico. The guy can literally pay the cops to look the other way, and he obviously has the money to do so	Security	Negative
India also built the largest border walls with pakistan, that too in the Himalayan foothills and thar desert. Really helped in curbing cross border terror though.	Security	Positive

Table 9: Examples of posts and comments from Reddit and the values and stances that the relevance and stance models predicted for each instance.

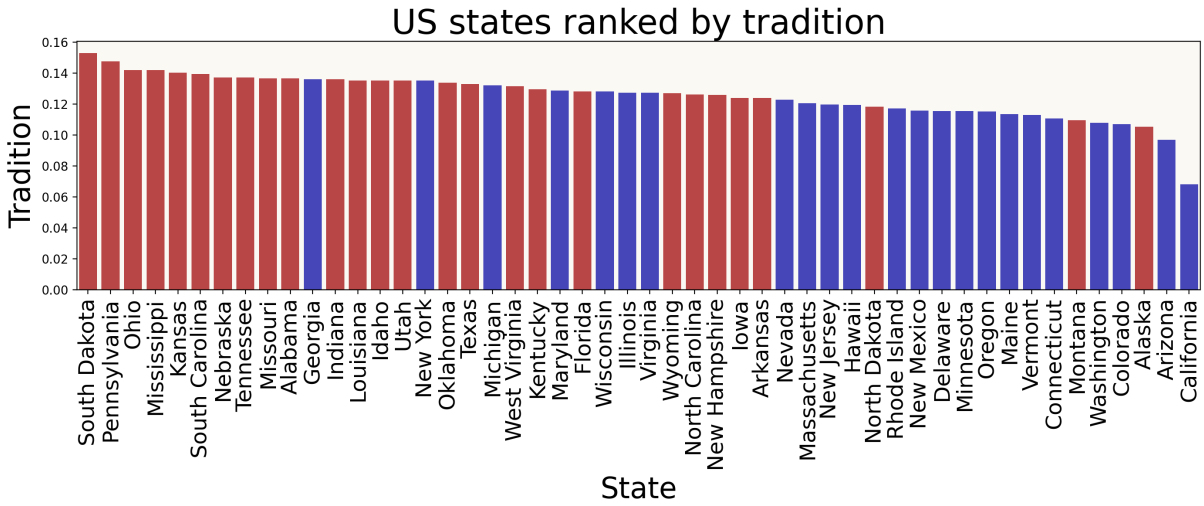

Figure 5: US states sorted by their *tradition* values extracted from Reddit, colour coded based on the recent US election results (Blue – democratic majority. Red – republican majority).

Value	Subreddits
Tradition	r/religion, r/Anglicanism, r/TraditionalCatholics, r/AskAChristian, r/AcademicBiblical, r/Catholic, r/AskBibleScholars, r/Episcopalian, r/AskAPriest, r/Antitheism, r/OrthodoxChristianity, r/Catholicism, r/antitheistcheesecake, r/agnostic, r/atheistmemes, r/CatholicMemes, r/Bible, r/RadicalChristianity, r/atheism, r/religiousfruitcake, r/Gnostic, r/Christianity, r/christianmemes, r/OpenChristian, r/paganism, r/excatholic, r/atheismindia, r/EXHINDU, r/ChristianUniversalism, r/bahai
Benevolence	r/coparenting, r/Stepmom, r/BipolarSOs, r/theotherwoman, r/stepparents, r/AdultChildren, r/ADHD_partners, r/CaregiverSupport, r/SingleParents, r/Petloss, r/Adoption, r/emotionalabuse, r/EstrangedAdultChild, r/AgingParents, r/FriendshipAdvice, r/AnxiousAttachment, r/Adopted, r/attachment_theory, r/AsOneAfterInfidelity, r/NarcAbuseAndDivorce, r/regretfulparents, r/survivinginfidelity, r/Fosterparents, r/AvoidantAttachment, r/AlAnon, r/Custody, r/BPDlovedones, r/Codependency, r/emotionalneglect, r/SuicideBereavement
Conformity	r/AgainstHateSubreddits, r/reclassified, r/TheseFuckingAccounts, r/AmIFreeToGo, r/ModSupport, r/modsbeingdicks, r/ModsAreKillingReddit, r/policebrutality, r/modhelp, r/WatchRedditDie, r/Bad_Cop_No_Donut, r/Chiraqhits, r/JustUnsubbed, r/legaladviceofftopic, r/circlebroke2, r/HOA, r/BadNeighbors, r/redditrequest, r/law, r/fuckHOA, r/40kOrkScience, r/legaladvice, r/LegalAdviceIndia, r/ACAB, r/Ask_Lawyers, r/legaladvicecanada, r/CapitolConsequences, r/supremecourt, r/NarcAbuseAndDivorce, r/JustNoTruth
Hedonism	r/crossdressing, r/transadorable, r/Autumn, r/TheMidnight, r/dykesgonemild, r/cakedecorating, r/FreeCompliments, r/GothStyle, r/NailArt, r/cozy, r/PlusSizeFashion, r/MTFSelfieTrain, r/oldhagfashion, r/femboy, r/transpositive, r/christmas, r/malepolish, r/RainbowEverything, r/gaybrosgonemild, r/HappyTrees, r/selfie, r/cookiedecorating, r/weddingdress, r/genderfluid, r/happy, r/Watercolor, r/AltJ, r/AnyaTaylorJoy, r/BigNoseLadies, r/Madonna
Power	r/FundRise, r/AskEconomics, r/finance, r/acorns, r/HistoricalWhatIf, r/geopolitics, r/strongblock, r/UWMCShareholders, r/debtfree, r/qyldgang, r/rocketpool, r/Anchor, r/IndianStreetBets, r/Economics, r/dividends, r/VoteBlue, r/EuropeanFederalists, r/defi, r/ValueInvesting, r/AlibabaStock, r/CredibleDefense, r/Stellar, r/DalalStreetTalks, r/CapitalismVSocialism, r/PoliticalDiscussion, r/LETFs, r/suzerain, r/AllCryptoBets, r/SatoshiBets, r/startups
Achievement	r/passive_income, r/FundRise, r/ValueInvesting, r/OptionsMillionaire, r/EngineeringResumes, r/xboxachievements, r/algotrading, r/FuturesTrading, r/resumes, r/abmlstock, r/AllCryptoBets, r/cryptostreetbets, r/dividends, r/SatoshiBets, r/Forex, r/EANHLfranchise, r/IndianStockMarket, r/quant, r/Blogging, r/aabbstock, r/CryptoMars, r/Daytrading, r/ETFs, r/startups, r/SEO, r/SaaS, r/biotech, r/SecurityAnalysis, r/Wishstock, r/Xelastock
Self-direction	r/CapitalismVSocialism, r/ScientificNutrition, r/Abortiondebate, r/DebateAnarchism, r/BasicIncome, r/DebateAVegan, r/Neuropsychology, r/changemyview, r/AskEconomics, r/AskPsychiatry, r/AskDID, r/nutrition, r/AskSocialScience, r/healthcare, r/askpsychology, r/radicalmentalhealth, r/ketoscience, r/DebateReligion, r/Marxism, r/PsychedelicTherapy, r/psychoanalysis, r/intentionalcommunity, r/TrueUnpopularOpinion, r/HealthInsurance, r/Microbiome, r/AskLibertarians, r/sorceryofthespectacle, r/actual_detrans, r/ScienceBasedParenting, r/DebateAChristian
Universalism	r/DebateAVegan, r/IsraelPalestine, r/DebateEvolution, r/DebateAnarchism, r/DebateReligion, r/changemyview, r/Abortiondebate, r/CapitalismVSocialism, r/AskSocialScience, r/EndlessWar, r/LeftWingMaleAdvocates, r/DebateAnAtheist, r/TrueUnpopularOpinion, r/PoliticalDiscussion, r/Ask_Politics, r/IntellectualDarkWeb, r/chomsky, r/AskFeminists, r/Marxism, r/DebateCommunism, r/israelexposed, r/DebateAChristian, r/AskAnthropology, r/ControversialOpinions, r/media_criticism, r/askphilosophy, r/Socialism_101, r/climatechange, r/conservation, r/honesttransgender
Stimulation	r/Autumn, r/crossdressing, r/TheMidnight, r/transadorable, r/FreeCompliments, r/dykesgonemild, r/cakedecorating, r/xboxachievements, r/PlusSizeFashion, r/GothStyle, r/cozy, r/MTFSelfieTrain, r/christmas, r/AltJ, r/Madonna, r/weddingdress, r/NailArt, r/transpositive, r/happy, r/Hobbies, r/HappyTrees, r/BigNoseLadies, r/oldhagfashion, r/selfie, r/gaybrosgonemild, r/RainbowEverything, r/BeachHouse, r/RoomieOfficial, r/Watercolor, r/FlorenceAndTheMachine
Security	r/CredibleDefense, r/ww3, r/syriancivilwar, r/geopolitics, r/EndlessWar, r/warinukraine, r/AfghanConflict, r/UkrainianConflict, r/UkraineConflict, r/CombatFootage, r/LessCredibleDefence, r/GunsAreCool, r/war, r/FutureWhatIf, r/nuclearweapons, r/WarCollege, r/UkraineInvasionVideos, r/UkraineRussiaReport, r/RussiaUkraineWar2022, r/UkraineWarReports, r/N_N_N, r/UkraineWarVideoReport, r/HistoryWhatIf, r/worldevents, r/ukraine, r/chomsky, r/HistoricalWhatIf, r/ukraina, r/anonymous, r/IndianDefense

Table 10: Subreddits with the highest signal for each one of the ten Schwartz values. The stance of Green subreddits towards the value is positive (above 0.2), whereas Red indicates negative stance (below -0.2).

Value	Positive Stance	Negative Stance
Tradition	r/Ankrofficial, r/lds, r/CharliDamelioMommy, r/Christian, r/AskAPriest, r/Bible, r/bahai, r/Quakers, r/PrismaticLightChurch, r/OrthodoxChristianity	r/SuperModelIndia, r/Jewdank, r/EXHINDU, r/DesiMeta, r/linguisticshumor, r/exmuslim, r/AsABlackMan, r/Satan, r/IndoEuropean, r/AfterTheEndFanFork
Benevolence	r/freebsd, r/RandomKindness, r/Terraform, r/Petloss, r/nextjs, r/Wetshaving, r/AllCryptoBets, r/NixOS, r/vancouverhiking, r/ansible	r/FromDuvalToDade, r/CrimeInTheD, r/NBAYoungboy, r/40kOrkScience, r/LILUZIVERTLEAKS, r/DuvalCounty, r/Phillyscoreboard, r/Chiraqhits, r/SummrsXo, r/CARTILEAKS
Conformity	r/Ankrofficial, r/nanotrade, r/NervosNetwork, r/Vechain, r/steroids, r/USCIS, r/treelaw, r/Stellar, r/cancun, r/JapanFinance	r/Animewallpaper, r/kencarson, r/FromDuvalToDade, r/LilDurk, r/Cookierun, r/freddiegibbs, r/SummrsXo, r/DestroyLonely, r/Gunna, r/okbuddydaylight
Hedonism	r/eastside, r/RedditPHCyclingClub, r/OaklandFood, r/carcamping, r/CryptoMars, r/VeganBaking, r/TheHague, r/ZZZ_Official, r/pottedcats, r/ambientmusic	r/depression, r/TIHI, r/Shark_Park, r/willowbramley, r/Phillyscoreboard, r/migraine, r/BodyDysmorphia, r/SuicideWatch, r/2meirl4meirl, r/anhedonia
Power	r/Yotsubros, r/Fitness, r/CryptoMoonShots, r/infertility, r/AllCryptoBets, r/ketoscience, r/steroids, r/ProgrammingLanguages, r/cryptostreetbets, r/cooperatives	r/masterhacker, r/Stake, r/uknews, r/RustConsole, r/uspolitics, r/OPBR, r/BidenIsNotMyPresident, r/capitalism_in_decay, r/Patriot911, r/occupywallstreet
Achievement	r/infertility, r/theravada, r/edrums, r/raisingkids, r/CryptoMars, r/Yotsubros, r/cozy, r/gaidhlig, r/PrismaticLightChurch, r/learnrust	r/DreamStanCringe, r/FGOmemes, r/AntiTrumpAlliance, r/BidenWatch, r/misanthropy, r/Patriot911, r/TRUTHsocialWatch, r/Instagram, r/TwitterCringe, r/Negareddit
Self-direction	r/Mosses, r/jungle, r/LandscapingTips, r/icecreamery, r/esp32, r/SatoshiBets, r/rust_gamedev, r/openstreetmap, r/QuantumComputing, r/cryptostreetbets	r/Negareddit, r/misanthropy, r/PeopleFuckingDying, r/BoomersBeingFools, r/AmericanFascism2020, r/RepublicanValues, r/FoxFiction, r/ParlerWatch, r/libsofreddit, r/FragileWhiteRedditor
Universalism	r/SatoshiBets, r/CryptoMars, r/AllCryptoBets, r/nextjs, r/AskAstrophotography, r/GardenWild, r/vancouverhiking, r/dungeondraft, r/EatCheapAndVegan, r/GraphicsProgramming	r/CrimeInTheD, r/FromDuvalToDade, r/Phillyscoreboard, r/DaDumbWay, r/DuvalCounty, r/Chiraqhits, r/BruceDropEmOff, r/punchableface, r/ConservativeRap, r/NYStateOfMind
Stimulation	r/CryptoMars, r/cryptostreetbets, r/JoshuaTree, r/wonderdraft, r/reenactors, r/AllCryptoBets, r/estoration, r/yerbamate, r/GiftIdeas, r/OaklandFood	r/depression, r/Shark_Park, r/heck, r/TwitterCringe, r/depressed, r/CommercialsIHate, r/Demps, r/D_Demps, r/christenwhitmansnark, r/Sleepparalysis
Security	r/BoringCompany, r/haskell, r/ProgrammingLanguages, r/rust, r/psychoanalysis, r/crypto, r/steroids, r/AskComputerScience, r/infertility, r/DebateEvolution	r/holdmycosmo, r/davidbowiecirclejerk, r/RoastMyCat, r/Chiraqhits, r/DaDumbWay, r/okbuddydaylight, r/bottomgear, r/SuicideWatch, r/CrimeInTheD, r/Bombing

Table 11: Subreddits expressing each value's strongest positive and negative stances.

Magnitude	Subreddits
Maximal	r/DebateAnarchism, r/Abortiondebate, r/therapyabuse, r/CapitalismVSocialism, r/changemyview, r/AvoidantAttachment, r/LeftWingMaleAdvocates, r/coparenting, r/ADHD_partners, r/DebateAVegan, r/Ask_Politics, r/IsraelPalestine, r/PoliticalDiscussion, r/AskSocialScience, r/NarcAbuseAndDivorce, r/AskDID, r/attachment_theory, r/Adoption, r/kpoprants, r/TrueUnpopularOpinion
Minimal	r/vegan1200isplenty, r/caloriecount, r/Watchexchange, r/Brogress, r/crystalgrowing, r/sneakermarket, r/gundeals, r/buildapcsales, r/whatisit, r/NMSCoordinateExchange, r/BulkOrCut, r/astrophotography, r/legodeal, r/whatisthisthing, r/whatsthisfish, r/filmfashion, r/TipOfMyFork, r/1500isplenty, r/safe_food, r/Repbudgetfashion

Table 12: Subreddits with the highest and lowest total magnitude of Schwartz values, calculated according to $\text{mag}(\mathcal{S}) = |\boldsymbol{u}_{\text{rel}}(\mathcal{S})|_2$.

Country	Language	Subreddits
Australia	English	r/australia, r/Australia_, r/australian, r/sydney, r/melbourne
Brazil	Portuguese	r/brasil, r/brasilivre, r/GTAorBrazil, r/brasilia, r/saopaulo
China	Chinese	r/China, r/China_irl, r/real_China_irl
Colombia	Spanish	r/Colombia, r/ColombiaReddit, r/Bogota
Costa Rica	Spanish	r/costarica, r/Ticos
Croatia	Croatian	r/croatia, r/zagreb
Czech Republic	Czech	r/czech
Ecuador	Spanish	r/ecuador
Estonia	Estonian	r/Eesti
Faroe Islands	All	r/Faroeislands
Finland	Finnish	r/Suomi
France	French	r/france, r/paris
Georgia	Georgian	r/Sakartvelo, r/tbilisi
Germany	German	r/de, r/deuchland, r/Munich, r/berlin
Ghana	English	r/ghana
Greece	Greek	r/greece
Hong Kong	Chinese	r/HongKong
Iceland	Icelandic	r/Iceland
India	All	r/india, r/IndiaSpeaks, r/unitedstatesofindia, r/askindia
Indonesia	Indonesian	r/indonesia
Israel	Hebrew	r/Israel, r/telaviv, r/jerusalem
Italy	Italian	r/italy, r/Italia
Japan	Japanese	r/japan, r/tokyo, r/newsokur
New Zealand	English	r/newzealand, r/auckland
Oman	Arabic	r/Oman
Peru	Spanish	r/PERU
Philippines	English	r/Philippines
Poland	Polish	r/Polska
Portugal	Portuguese	r/portugal, r/lisboa, r/porto
Romania	Romanian	r/Romania, r/bucuresti, r/iasi, r/cluj, r/timisoara
Serbia	Serbian	r/serbia
Slovakia	Slovak	r/Slovakia
South Africa	All	r/southafrica, r/capetown, r/johannesburg
South Korea	Korean	r/hanguk
Spain	Spanish	r/Espana, r/spain, r/Madrid, r/Barcelona
Sweden	Swedish	r/Sverige, r/sweden, r/stockholm, r/svenskpolitik
Turkey	Turkish	r/Turkey, r/istanbul, r/ankara, r/Ismir
Ukraine	Ukrainian	r/ukraine, r/Ukraine_UA
United Kingdom	All	r/unitedkingdom, r/london, r/manchester, r/CasualUK, r/unitedkingdom, r/askuk
USA	All	All states subreddits
Vietnam	Vietnamese	r/VietNam

Table 13: Countries, languages and subreddits used in the correlation study.

subreddit	closest subreddits
r/vegan	r/intersex, r/transgenderUK, r/TransSpace, r/Transmedical, r/ABCDesis, r/LGBTindia, r/DeepThoughts, r/nonduality, r/exvegans, r/BlockedAndReported
r/carnivore	r/carnivorediet, r/Psoriasis, r/PsoriaticArthritis, r/rheumatoid, r/acne, r/Hashimotos, r/Testosterone, r/PlasticSurgery, r/lupus, r/kratom
r/communism	r/InformedTankie, r/CommunismWorldwide, r/CPUSA, r/socialism, r/StupidpolEurope, r/LatinAmerica, r/sendinthetanks, r/ROI, r/Africa, r/myanmar
r/Capitalism	r/georgism, r/Unions, r/SandersForPresident, r/labor, r/NewDealAmerica, r/BernieSanders, r/WorkersStrikeBack, r/ndp, r/theydidthemath, r/MayDayStrike
r/monarchism	r/leftistvexillology, r/UsefulCharts, r/RoughRomanMemes, r/IndiaPlace, r/pureasoiaf, r/MedievalHistory, r/shittyhalolore, r/MemriTVmemes, r/ancientrome, r/darkwingsdankmemes
r/AbolishTheMonarchy	r/Sham_Sharma_Show, r/forwardsfromgrandma, r/WordAvalanches, r/Malaphors, r/Lal_Salaam, r/Metal, r/BanVideoGames, r/QanonKaren, r/PoliticalHumor, r/insanepeoplefacebook
r/Millennials	r/GenZ, r/PlusSizedAndPregnant, r/ThailandTourism, r/indiasocial, r/Liverpool, r/TallGirls, r/Psychedelics, r/Chandigarh, r/Cardiff, r/ChronicIllness
r/GenZ	r/Millennials, r/indiasocial, r/TallGirls, r/Kibbe, r/Zillennials, r/aggretsuko, r/precure, r/nerdfighters, r/KoeNoKatachi, r/comiccon
r/GenX	r/lesbianfashionadvice, r/CosplayNation, r/Xennials, r/bigboobproblems, r/femininebоys, r/PlusSize, r/distantsocializing, r/CasualPH, r/bald, r/TallGirls
r/Feminism	r/antiwoke, r/MensLib, r/ControversialOpinions, r/EnoughIDWspam, r/prochoice, r/prolife, r/fourthwavewomen, r/TransSpace, r/TrueUnpopularOpinion, r/AntiVegan
r/MensRights	r/antifeminists, r/AntiHateCommunities, r/AsABlackMan, r/LeftWingMaleAdvocates, r/aznidentity, r/EnoughPCMSpam, r/FragileWhiteRedditor, r/BlatantMisogyny, r/TheLeftCantMeme, r/ForwardsFromKlandma
r/atheism	r/exmuslim, r/antitheistcheesecake, r/extomatoes, r/Antitheism, r/atheismindia, r/mormon, r/progressive_islam, r/Judaism, r/RadicalChristianity, r/religiousfruitcake
r/spirituality	r/SpiritualAwakening, r/awakened, r/pureretention, r/Stoicism, r/Shamanism, r/starseeds, r/Mediums, r/Soulnexus, r/Semenretention, r/SASSWitches
r/religion	r/AskAChristian, r/ChristianUniversalism, r/RadicalChristianity, r/bahai, r/hinduism, r/shia, r/Bible, r/progressive_islam, r/extomatoes, r/mormon

Table 14: For each subreddit in the controversial topics analysis, the 10 most similar subreddits in terms of values.