

# Quantifying moral foundations from various topics on Twitter conversations

Rishemjit Kaur<sup>1</sup> and Kazutoshi Sasahara<sup>\*2,3</sup>

<sup>1</sup>CSIR-Central Scientific Instruments Organisation, Chandigarh, India

<sup>2</sup>Department of Complex Systems Science, Nagoya University, Nagoya, Japan

<sup>3</sup>Center for Complex Networks and Systems Research, Indiana University, Bloomington, IN, USA

## Abstract

Moral foundations theory explains variations in moral behavior using innate moral foundations: Care, Fairness, Ingroup, Authority, and Purity, along with experimental supports. However, little is known about the roles of and relationships between those foundations in everyday moral situations. To address these, we quantify moral foundations from a large amount of online conversations (tweets) about moral topics on the social media site Twitter. We measure moral loadings using latent semantic analysis of tweets related to topics on abortion, homosexuality, immigration, religion, and immorality in general, showing how the five moral foundations function in spontaneous conversations about moral violating situations. The results indicate that although the five foundations are mutually related, Purity is the most distinctive foundation and Care is the most dominant foundation in everyday conversations on immorality. Our study shows a new possibility of natural language processing and social big data for moral psychology.

**Keywords:** moral foundations theory; natural language processing; latent semantic analysis; social big data; Twitter.

---

\*Correspondence should be addressed to K.S. (sasahara@nagoya-u.jp)

# 1 Introduction

Social media communication is a regular part of our daily life and a large amount of linguistic information is posted and digitally recorded every day. Our digital behavioral traces allow us to quantify human behavior in a natural setting to complement experimental data. The availability of social big data for social science research is now widely acknowledged [5, 13], and computational social science provides new insights into human nature in the digital era [18, 19, 21].

In this paper, we study moral behavior using social big data from a popular microblogging platform, Twitter. One of the most influential theories about moral behavior is Moral Foundations Theory (hereafter, MFT) proposed by Haidt and Joseph [9]. MFT can explain variations in moral behavior on the basis of the following innate moral foundations:

- Care: disliking the pain of others and feeling of protecting the vulnerable;
- Fairness: doing the right thing or justice based on shared rules;
- Ingroup: being loyal to social groups, including family and nation;
- Authority: respecting and obeying tradition and legitimating authority;
- Purity: feeling an antipathy for disgusting things and contamination.

According to Haidt et al. [8], Care and Fairness, which focus on individuals' rights and freedom, are grouped together and referred to as 'individualizing foundations,' whereas Ingroup, Authority, and Purity are mainly centered around binding people together for the welfare of community or group and are called 'binding foundations.' MFT is especially successful in explaining political ideology and cross-cultural differences in moral behavior. For example, previous studies have shown that political progressives or 'liberals' stress only Care and Fairness in moral judgment, whereas political conservatives equally stress all moral foundations [9]. Another study has demonstrated that people in collectivist societies are more sensitive to violations of the community-related moral foundations, whereas people in Western societies are more likely to discriminate between care-related violations and convention-related violations [9]. While moral foundations have been confirmed experimentally, more recent studies have used natural language processing (NLP) techniques and social big data [4, 16, 17]: they measured the moral loadings from written expressions on a particular real-world event (2013 government shutdown in the US) from Twitter, showing that Purity is related to social distance. The findings were then validated experimentally. These studies have led to a new

method with empirical evidence for MFT, yet we know little about the roles and relationships between moral foundations in everyday moral situations.

To address these issues, we quantified moral foundations from various topics in everyday conversations on Twitter. More specifically, we examined the relationships of five moral foundations by measuring moral loadings from a large amount of posted messages or tweets that are related to moral violating situations. Moreover, we analyzed tweets about moral topics such as abortion, homosexuality, immigration, and religion for insights into the roles of moral foundations in various topics.

## **2 Method**

### **2.1 Data collection**

We collected tweets related to moral concerns—abortion (‘abortion’), homosexuality (‘homosexuality’ OR ‘homosexual’), immigration (‘immigration’ OR ‘immigrant’), religion (‘religion’ OR ‘religious’), and immorality (‘immorality’ OR ‘immoral’) that were posted between March 1 and April 24, 2016 using the Twitter Search API (<https://dev.twitter.com/docs/api/>). The search queries written in brackets were selected to observe daily conversations on moral topics. English tweets were used for our text analysis. The sizes of the datasets are 1516119 tweets for abortion, 456674 tweets for homosexuality, 2102886 tweets for immigration, 4628102 tweets for religion, and 217975 tweets for immorality, respectively. The topic datasets (i.e., tweets with ‘abortion’, ‘homosexuality’, ‘immigration’, or ‘religion’) were used only to identify keywords and context words associated with these topics. Except these, our analysis has been done with the immorality dataset (i.e., tweets with ‘immorality’).

### **2.2 Data pre-processing**

Each tweet consists of multiple fields such as text, retweeted status, language, and other meta-data. We extracted ‘text’ field and ‘retweeted.status.text’ field if the tweet has been re-tweeted. To clean the texts, we removed stop words (e.g., ‘a’, ‘an’, ‘the’) that are defined in NLTK library (<http://www.nltk.org>), URLs, screen names (e.g., @BarackObama), special characters (e.g., ! and \$), numbers, leading and trailing white-spaces, and short words with length less than three (e.g., th, gf) from tweets. After that, hashtag symbol was removed from tweets (e.g., ‘#harm’ was replaced with ‘harm’). Furthermore, the words used in queries for Twitter Search API were removed because we are interested in popular words that characterize moral topics except those used in queries. For example, in the

case of the immorality dataset, ‘immoral and ‘immorality were removed and thus not considered as keywords and context words for subsequent analyses. Further, we removed the duplicate tweets from our dataset. Finally, we tokenized tweet texts by splitting on white-spaces. Suppose a tweet “@CharlesMBlow 50% marginal taxrates aren’t immoral. Letting the majority of public school kids live in poverty is. <https://t.co/grEgeu9lPj>, by following the steps stated above we are left with ‘marginal’, ‘taxrates’, ‘letting’, ‘majority’, ‘public’, ‘school’, ‘kids’, ‘live’, ‘poverty.’ This pre-processing was applied to the topic datasets and the immorality dataset after which we refer to them as the topic corpora and the immorality corpus, respectively.

### **2.3 Quantification of moral foundations**

To quantify the moral foundations from the Twitter corpora, we used a method proposed by Dehghani et al. [4] with several modifications. Simply speaking, the method is based on the latent semantic analysis (LSA) [3, 12]: using a bag-of-words model, a corpus is represented by a word-context matrix and by reducing its dimensionality we get low-dimensional word vectors, in which similar meaning words are represented by similar vectors. With this method, we can measure moral foundations from everyday tweets by comparing the words in the moral foundations dictionary (described below) and words in the Twitter corpora. The details are described below.

#### **2.3.1 Moral foundations dictionary**

The moral foundations dictionary (hereafter, the MF dictionary) was created by Graham and Haidt [7]. It is available online at <http://moralfoundations.org>. The MF dictionary lists the words and word stems associated with five moral foundations—Care (called ‘Harm’ in the MF dictionary), Fairness, Ingroup, Authority and Purity, along with general words associated with morality and immorality. These are further divided into two categories, ‘virtue’ and ‘vice.’ Virtue words are foundation-supporting words (e.g. safe\* and shield for Care virtue), whereas vice words are foundation-violating words (e.g. kill and ravage for Care vice). To limit our analysis of moral violating situations, we used 149 vice words in the MF dictionary.

#### **2.3.2 Selection of keywords and context words**

To construct a word-context matrix used for the succeeding analysis, we selected keywords and context words based on the tf-idf score from the topic corpora and the immorality corpus (Fig. 1). From each of the Twitter corpora, we created a

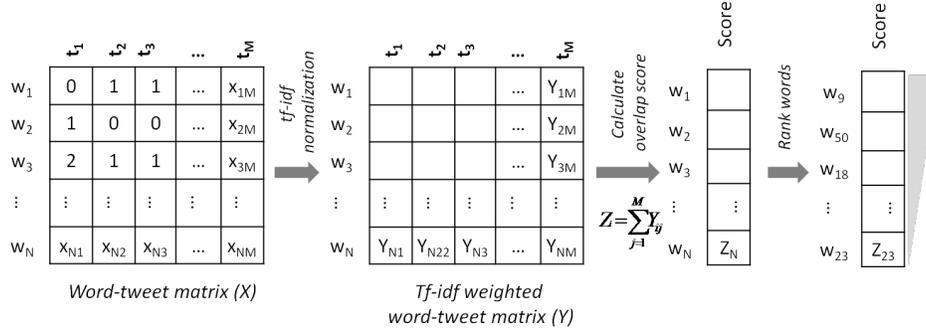


Figure 1: Selection of keywords and context words

word-tweet matrix  $X$ , respectively, in which each row denotes word  $w$  and each column denotes tweet  $t$ , and the element  $X_{ij}$  denotes the frequency of word  $w_i$  in tweet  $t_j$ . Then, this matrix was converted to a tf-idf weighted matrix  $Y$  by the following equation:

$$tf-idf(w_i, t_j) = tf(w_i, t_j)(\log(M + 1) - \log(df(w_i))), \quad (1)$$

where  $tf(w_i, t_j)$  is the frequency of word  $w_i$  in tweet  $t_j$ ,  $M$  is the total number of tweets and  $df(w_i)$  is the number of tweets containing word  $w_i$ . The overlap score of a word  $w_i$  in  $Y$  is computed as a measure of word importance [14]:

$$Score(w_i) = \sum_{j=1}^M (tf-idf(w_i, t_j)). \quad (2)$$

According to (2), the words were ranked in the decreasing order and the top  $N_1$  words were selected as keywords and the top  $N_2$  words as context words for a word-context matrix in order to use the subsequent analysis. Using the same settings of Dehgani et al. [4], we set  $N_1 = 2000$  and  $N_2 = 20000$ .

### 2.3.3 Word-context matrix and singular value decomposition

In using keywords as columns and context words as rows, we created a word-context matrix  $C$ , in which the element  $C_{ij}$  represents the number of co-occurrence of word  $w_i$  and context word  $w_j$ . In contrast to Dehgani et al. [4], we further converted it to the positive pointwise mutual information (PPMI) based matrix so as to not assign a higher weight to a popular general word that is irrelevant to moral issues [11]:

$$PPMI(v_i, v_j) = \max(\log_2(P(v_i, v_j)/P(v_i)P(v_j)), 0), \quad (3)$$

Here,  $v_i$  and  $v_j$  are the words in  $C$ ,  $P(v_i)$  is the occurrence probability of a word  $v_i$  and  $P(v_i, v_j)$  is the joint occurrence probability of  $v_i$  and  $v_j$ . Using the PPMI-based word-context matrix  $C$ , we applied singular value decomposition (SVD) to achieve a lower dimensional representation of keywords:  $C = U \Sigma V^*$  where  $U$  is a left singular matrix,  $\Sigma$  is a diagonal matrix and  $V^*$  is the right singular matrix. The top  $k$  dimensions (in our case  $k = 100$ ) of  $U$  matrix are retained, and each row of this reduced matrix represents a word in  $k$  dimensions. The matrices  $\Sigma$  and  $V^*$  are discarded. In this way, SVD converts the high-dimensional and sparse word-context matrix into a lower dimensional, real-valued matrix, which represents the semantic relationships between words [11].

### 2.3.4 Construction of context vectors

The resulting word vector space is linear. Therefore, it is possible to approximate a text by adding corresponding word vectors [4]. In doing this, we constructed tweet context vectors, topic context vectors, and moral foundation (MF) context vectors.

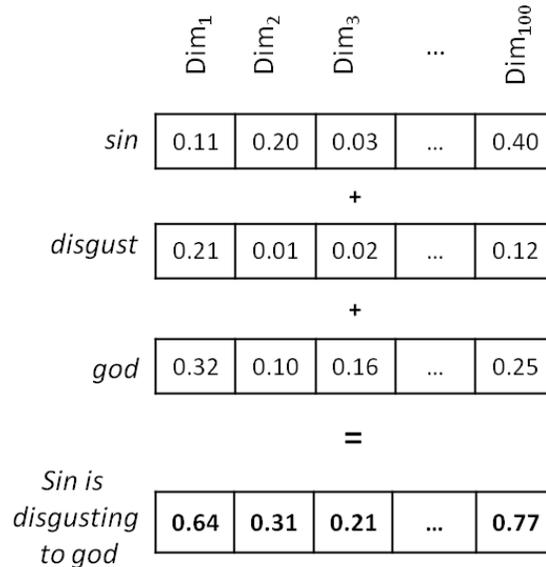


Figure 2: Example of the construction of a tweet context vector

Fig. 2 shows an example of how to construct a tweet context vector. Given a tweet, ‘Sin is disgusting to god. the tweet context vector is obtained by the addition of the vectors of corresponding words ‘sin, ‘disgust, and ‘god. Note that the stop

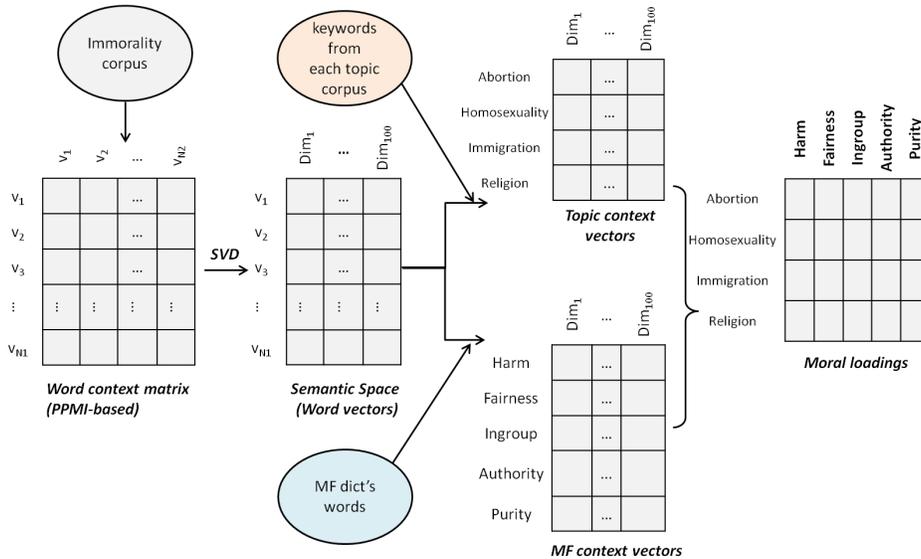


Figure 3: Measurement of moral loadings for topics

words ‘is and ‘to’ were removed during pre-processing. Similarly, we constructed MF context vectors by adding the vectors corresponding to the MF dictionary’s words present in the immorality corpus, and topic context vectors by adding the vectors corresponding to topic specific words. As mentioned earlier, we have used the topic corpus to select keywords describing a specific topic (e.g., topic ‘religion’ may be described by various words such as ‘God’, ‘church’, ‘religious’, ‘Islam’, ‘Christianity’, ‘Hindu’). For comparison, we created the topic context vectors using 10 keywords and 100 keywords, respectively.

### 2.3.5 Measurement of moral loadings

Of the three kinds of context vectors described above, one possible way to quantify moral foundations from texts is to measure moral loadings [4], which is defined as similarity to MF context vectors. For example, the moral loading of a tweet is measured by the cosine similarity between the tweet context vector and the MF context vector. The cosine value of 1 signifies synonymy of expressions, while the cosine value of 0 indicates that these are semantically unrelated. In this manner, if tweet context vectors (TV) are represented by  $\langle TV_1, TV_2, TV_3, \dots, TV_M \rangle$  and MF context vectors (MV) by  $\langle MV_{\text{Care}}, MV_{\text{Fairness}}, MV_{\text{Ingroup}}, MV_{\text{Authority}}, MV_{\text{Purity}} \rangle$ , then the moral loadings are summarized in a  $M \times 5$  matrix whose element represents the similarity between a tweet and each moral foundation. Similarly, we can

compute moral loadings for topics (i.e., abortion, homosexuality, immigration, religion), which results in a 4 (number of topics)  $\times$  5 (number of moral foundations) matrix representing similarity between topics and moral foundations. This procedure is illustrated in Fig. 3.

## 3 Results

### 3.1 Immorality in everyday tweets

We analyzed the immorality corpus in terms of moral violating situations and found that 81.1% of the MF dictionary's vice words were present in the corpus. Fig. 4 shows examples of the MF dictionary's vice words present in the corpus, in which the size of a word is proportional to its occurrence frequency and colors represent different moral foundations. For example, words such as 'war', 'killing', 'violence\*', 'attack' are most frequently occurring words from Care foundation, and words such as 'sin', 'sick', 'disgust', 'dirt', 'adultery\*' were present in the Purity foundation. Note that the MF dictionary's words such as 'spurn', 'favoritism', 'jilt\*', 'obstruct', 'blemish' are not shown in this figure because either these words did not pass the keyword selection described before or these were not present in our corpus. It is worth noticing that 'Illegal' is the most frequent word in the five foundations. This suggests that illegal issues would be the most immoral in everyday moral situations.

Table 1 shows examples of the tweets with the highest moral loadings from each moral foundation. Tweet #1 has the highest similarity with the MF context vector for Care. This could be because of the presence of the word 'kills' in the tweet. Tweet #4 is an interesting example where the highest correlation of 0.553 is obtained with Authority, the second highest of 0.348 was with Ingroup, and the third highest was 0.346 and with Fairness. If we look at the MF dictionary, the word 'treasonous' belongs to both Ingroup and Authority, and the word 'illegal' belongs to Authority. Thus, this tweet has a higher correlation with Authority than Ingroup. This tweet also has a high correlation with Fairness even though it does not include any word from the Fairness category in the MF dictionary. It is because of the presence of words such as 'patriotic' (0.33), 'government' (0.23), etc. in the tweet which are highly correlated with Fairness in terms of cosine similarity. This is an example case where a standard word counting approach [2,7] may fail because the MF dictionary includes only a limited number of words (e.g., 149 words in the vice category), and other everyday moral words are not considered. In contrast, the matrix  $U$  resulting from the moral loading method provides an extended version of the MF dictionary, which we will detail later. Hence, it may be applicable to a wider class of texts.

To examine which moral foundations are actually used as well as how often

Table 1: Tweets showing the highest moral loadings with each foundation

#	Tweet	Care	Fairness	Ingroup	Authority	Purity
1	@harikriss it should it's immoral and kills innocent people	<b>0.772</b>	-0.016	0.199	0.063	0.113
2	@goat777face @ColinUlster96 @dkm49321 Yeah, not being to discriminate is so unbelievably immoral and unjust.	0.184	<b>0.712</b>	0.177	0.311	0.006
3	@DanWosHere @malcolmtyson @JulianBurnside Foul immoral traitors will go to any lengths to defend the enemy they've allied with.	0.225	0.189	<b>0.544</b>	0.336	0.186
4	@Scribbles646 @Snowden So, we agree arming Al Qaeda is treasonous. Exposing the illegal/immoral actions of a government is patriotic.	0.218	0.346	0.348	<b>0.553</b>	0.007
5	#PresstitutesDay Sick of these indecent, perverted, shameless, wicked, sinful, immoral, lewd, self-indulgent anti nationals.	0.134	0.249	0.238	0.044	<b>0.684</b>



Table 2: Number of tweets present in each moral foundation

Foundation	Number of tweets
Care	21135
Fairness	15731
Ingroup	6665
Authority	4932
Purity	14587

Table 3: Cosine similarities between MF context vectors

Foundation	Care	Fairness	Ingroup	Authority	Purity
Care	-	0.113	<b>0.394</b>	0.133	0.229
Fairness	-	-	0.147	0.223	0.14
Ingroup	-	-	-	<b>0.598</b>	0.239
Authority	-	-	-	-	0.081
Purity	-	-	-	-	-

have compelling correlations with either individualizing or binding foundations.

### 3.2 Extended MF dictionary

Using the immorality corpus, we extended the MF dictionary by adding the semantically related words using the resulting word-context matrix. That is, we again calculated cosine similarities between five MF context vectors and word vectors and then listed the top 100 most similar words for each foundation. In this way, the extended MF dictionary had 500 total words.

We compared the original and extended dictionaries and found several similarities and dissimilarities. While most words are commonly present in the same foundation category (e.g., ‘kill’, ‘kills’, ‘attack\*’ are involved in Care), others present in different foundation categories (e.g., ‘insurgent’ is in Authority in the original dictionary but it appears in Care in the extended dictionary). Some words appear in more than one foundation in the extended dictionary (e.g., ‘terroris\*’, which is present in Ingroup in the original dictionary, appears both in Care and Ingroup). Importantly, words not present in the original dictionary have been added in the extended dictionary (e.g., ‘bombing’, ‘isis’, ‘torture’, ‘murder’, ‘nazis’, ‘stereotypes’ in Care, ‘racist’, ‘elitists’, ‘racism’ in Ingroup and ‘lust’, ‘flesh’, ‘devil’ in Purity). Furthermore, some morally neutral words, e.g., ‘Kashmir’ (the name of a place in India), is present in Care, which in recent years has seen political and social instability leading to violent clashes and a loss of life and property. The

presence of words such as ‘homosexual’ and ‘lesbians’ along with other words of Purity (e.g., ‘disgust’, ‘sinful’, ‘sin’) reflects the attitude of people towards these topics similar to ‘abortion’ in Care. In this way, the extended MF dictionary reflects actuarial word usages in everyday conversations.

To further examine how these foundations relate to each other, we mapped the extended dictionary’s word vectors along with moral foundations vectors on a 2D plane using principal component analysis (PCA) [10]. Fig. 5 shows that except Purity, the words belonging to other four foundations overlap with each other along the PC1 axis. Thus, the PC1 axis differentiates Purity and non-Purity foundations, which can be explained in terms of ‘person-based attributes’ vs. ‘situation-based attributes’ [1]. The other four foundations are differentiated along the PC2 axis, although with some overlaps. Care and Fairness are placed at the extreme ends along the PC2 axis, indicating that these two foundations are dissimilar to each other. This is similar to the result found in the previous section. This result indicates that Care and Fairness are both individualizing tendencies but function in different manners. Overall, both axes well differentiate Purity, Fairness, and Care from each other, but there is a significant overlap between Authority and Ingroup. It can also be seen that Ingroup has an overlap with Care, and Authority overlaps with Fairness to some extent. These results imply that the five moral foundations are not mutually exclusive—at least in everyday conversations on immorality.

### 3.3 Roles of moral foundations in various topics

To investigate the roles of moral foundations on different topics, we measured the moral loadings between different topics and moral foundations. Table 4 shows the result when we used the top 10 keywords according to score (eq. (2)) for topic context vectors. We see that topics ‘abortion and ‘religion are correlated to Care, whereas topics ‘immigration and ‘homosexuality have the highest similarities with Ingroup and Purity, respectively. The similarity of homosexuality with Purity can be explained as it has been found by Pizarro et al. [15] that purity violations evoke feeling of disgust. Disgust is also considered to be positively correlated with the negative attitudes towards homosexuals [20, 22]. Graham et al. [6] mentioned that immigrants are a trigger for Purity violations, but our results show that its highest correlation is with Ingroup violations.

The differences between moral concerns of Republicans and Democrats about ‘abortion were discussed by Sagi et al. [16]. Their results showed that Democrats were mostly concerned with Fairness, whereas Republicans were concerned with the Purity aspects of abortion. In our case, abortion has the highest correlation with Care.

When we used the top 100 keywords according to score (eq. (2)) for construct-

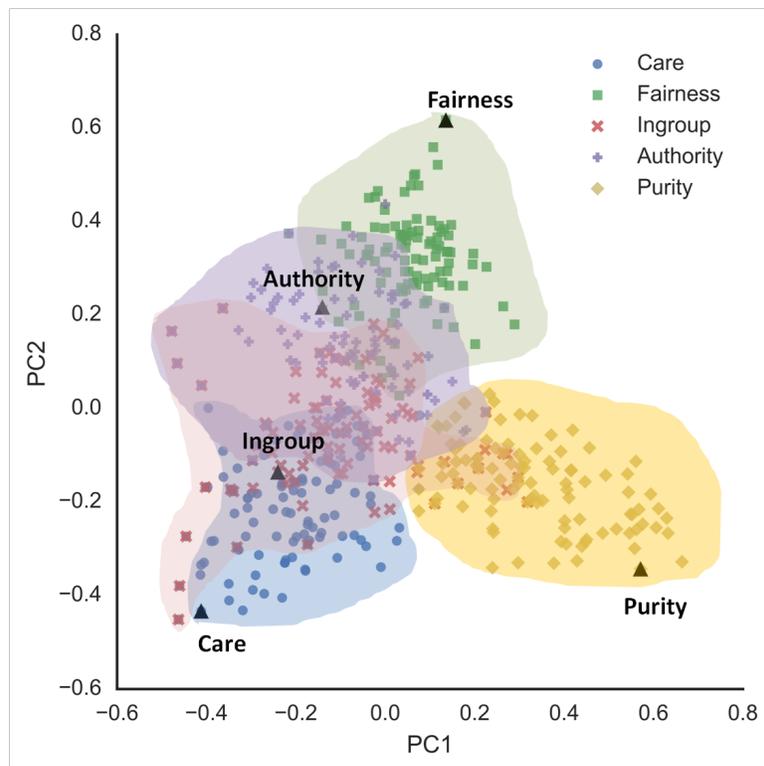


Figure 5: PCA plot of the extended MF dictionary's words along with five moral context vectors

ing topic context vectors from topic corpora, all topics were maximally correlated with Care (Table 5). This result suggests that although the most popular keywords (e.g. top 10 words) are related to the corresponding topics, many words ranked below the top 10 were related to Care. This is partly supported by the fact that the maximum number of tweets in our dataset were predominantly related to Care as shown Table 2.

## 4 Discussion

Most social psychology data lie in the bounding box of language, culture, ethnicity, and even environment. Social media sites like Facebook and Twitter, however, are useful places for behavioral data mining wherein people without the care of language, ethnicity, and boundaries are free to speak their mind without the need to think of any consequence thereof. Therefore, social big data is suitable for

Table 4: Moral loadings for topics (based on top 10 keywords)

	Care	Fairness	Ingroup	Authority	Purity
Abortion	<b>0.32</b>	0.15	0.16	0.16	0.14
Homosexuality	0.08	0.09	-0.01	0.04	<b>0.35</b>
Immigration	0.39	0.18	<b>0.51</b>	0.27	0.05
Religion	<b>0.43</b>	0.31	0.22	0.09	0.19

Table 5: Moral loadings for topics (based on top 100 keywords)

	Care	Fairness	Ingroup	Authority	Purity
Abortion	<b>0.39</b>	0.2	0.15	0.1	0.16
Homosexuality	<b>0.27</b>	0.2	0.08	0.06	<b>0.26</b>
Immigration	<b>0.39</b>	0.24	<b>0.36</b>	0.16	0.12
Religion	<b>0.33</b>	0.25	0.14	0.06	<b>0.23</b>

observing and measuring human behavior in a natural setting. This complements data from social psychology experiments.

We have analyzed spontaneous moral conversations on Twitter where people wrote about myriad topics such as abortion, homosexuality, immigration, religion, and immorality in general. By quantifying moral loadings from tweets, we observed that while the five moral foundations are mutually related, Care is the most dominant foundation and Purity is the most distinctive foundation—at least in on-line conversations about immorality. These findings need to be further tested with a larger size of corpora from various social media sites. With the MF dictionary, a simple word counting approach is often used to measure moral charge from texts, but the LSA-based approach used here has some advantages. It allows us to capture the meanings of moral words as vectors, and thus there are a wide variety of methods in vector semantics. For example, topic modeling is a promising application of NLP to gain new insights into moral psychology. Here, we fixed the size of a word-context matrix according to the purpose, but the estimation of the appropriate size is another important issue that future studies need to address. We are aware that translated versions of the MF dictionary are necessary for cross-cultural comparisons of moral conversations. This is also an important future direction. Although several forthcoming issues remain, the current study demonstrates a new possibility of NLP and social big data for moral psychology by quantifying moral diversity in everyday conversations.

## Acknowledgment

This work was supported by JSPS KAKENHI grant no.15H03446. We thank M. Karasawa for discussions.

## References

- [1] A. Chakroff and L. Young. Harmful situations, impure people: An attribution asymmetry across moral domains. *Cognition*, 136:30–37, Mar. 2015.
- [2] S. Clifford and J. Jerit. How words do the work of politics: Moral foundations theory and the debate over stem cell research. *The Journal of Politics*, 75(3):659–671, 2013.
- [3] S. Deerwester, S. T. Dumais, G. W. Furnas, T. K. Landauer, and R. Harshman. Indexing by latent semantic analysis. *Journal of the American Society for Information Science*, 41(6):391, 1990.
- [4] M. Dehghani, K. Johnson, J. Hoover, E. Sagi, J. Garten, N. J. Parmar, S. Vaisey, R. Iliev, and J. Graham. Purity homophily in social networks. *Journal of Experimental Psychology: General*, 145(3):366–375, 2016.
- [5] S. A. Golder and M. W. Macy. Digital footprints: opportunities and challenges for online social research. *Annual Review of Sociology*, 40(1):129–152, 2014.
- [6] J. Graham, J. Haidt, S. Koleva, M. Motyl, R. Iyer, S. P. Wojcik, and P. H. Ditto. Moral foundations theory: The pragmatic validity of moral pluralism. *Advances in Experimental Social Psychology, Forthcoming*, 47:55–130, 2013.
- [7] J. Graham, J. Haidt, and B. A. Nosek. Liberals and conservatives rely on different sets of moral foundations. *Journal of Personality and Social Psychology*, 96(5):1029–1046, 2009.
- [8] J. Haidt. Morality. *Perspectives on Psychological Science*, 3(1):65–72, 2008.
- [9] J. Haidt. *The righteous mind: Why good people are divided by politics and religion*. Pantheon, 2012.
- [10] L. Jolliffe. *Principal Component Analysis*. Wiley Online Library, 2002.

- [11] D. Jurafsky and J. H. Martin. *Speech and Language Processing*. Pearson, 2014.
- [12] T. K. Landauer and S. T. Dumais. A solution to Plato’s problem: The latent semantic analysis theory of acquisition, induction, and representation of knowledge. *Psychological Review*, 104(2):211, 1997.
- [13] D. Lazer, A. Pentland, L. A. Adamic, S. Aral, A.-L. Barabási, D. Brewer, N. A. Christakis, N. Contractor, J. Fowler, M. Gutmann, T. Jebara, G. King, M. W. Macy, D. Roy, and M. Van Alstyne. Computational social science. *Science*, 323(5915):721–723, 2009.
- [14] C. D. Manning, P. Raghavan, and H. Schütze. *Introduction to information retrieval*. Cambridge University Press, New York, NY, USA, 2009.
- [15] D. Pizarro, Y. Inbar, and C. Helion. On disgust and moral judgment. *Emotion Review*, 3(3):267–268, 2011.
- [16] E. Sagi and M. Dehghani. Measuring moral rhetoric in text. *Social Science Computer Review*, 32(2):132–144, 2014.
- [17] E. Sagi and M. Dehghani. Moral rhetoric in Twitter: A case study of the US Federal Shutdown of 2013. In *Proceedings of 36th Annual Conference of the Cognitive Science Society*, pages 1347–1352, 2014.
- [18] K. Sasahara. Visualizing collective attention using association networks. *New Generation Computing (in press)*.
- [19] K. Sasahara, Y. Hirata, Y. Hirata, M. Toyoda, M. Toyoda, M. Kitsuregawa, M. Kitsuregawa, and K. Aihara. Quantifying collective attention from Tweet stream. *PLoS ONE*, 8(4):e61823, 2013.
- [20] K. B. Smith, D. Oxley, M. V. Hibbing, J. R. Alford, and J. R. Hibbing. Disgust sensitivity and the neurophysiology of Left-Right political orientations. *PLoS ONE*, 6(10):e25552, 2011.
- [21] Y. Takeichi, K. Sasahara, R. Suzuki, and T. Arita. Concurrent bursty behavior of social sensors in sporting events. *PLoS ONE*, 10(12):e0144646, 2015.
- [22] J. A. Terrizzi, N. J. Shook, and W. L. Ventis. Disgust: A predictor of social conservatism and prejudicial attitudes toward homosexuals. *Personality and Individual Differences*, 49(6):587–592, 2010.