

Moods and Moral Values in Blog Posts

M. Généreux†

genereux@clul.ul.pt

R.P. Evans††

R.P.Evans@brighton.ac.uk

†Centro de Linguística da Universidade de Lisboa
Lisbon, Portugal

††Natural Language Technology Group
Brighton, United Kingdom

Résumé :

Dans cet article, nous utilisons une approche basée sur les corpus pour explorer la relation entre les valeurs morales et le bien-être humain dans les médias sociaux. Dans la foulée, nous examinons s’il y a un effet significatif relié au genre et à l’âge en ce qui concerne l’intérêt porté à la morale et quelles sont les valeurs que l’on priorise. Nos résultats supportent en partie l’idée d’une relation directe entre la moralité et le bien-être ainsi qu’un effet du genre et de l’âge sur l’intérêt des gens vis-à-vis des questions morales, mais aucun effet réel sur la priorité des valeurs.

Mots-clés : blogues, bien-être, moralité

Abstract :

In this paper we use a corpus-based approach on blog posts to investigate the relationship between moral values and human well-being in social media. In the process, we look at whether there is a significant gender and age effect with regards to interest in morality and value priority. We find some support for a direct relationship between morality and well-being and a gender/age effect on the extent to which people are interested in moral questions, but no real effect on value priority.

Keywords : blog posts, well-being, morality

1 Introduction

In his book *The Moral Landscape* [4], Sam Harris defends the idea that human (and animal) well-being is the direct consequence of people living ethical lives by prioritizing specific moral values. In other words, well-being is the direct consequence of what kind of moral principles or actions we apply for ourselves and others. This way of thinking implies that to experience well-being we should prioritize moral values that impact positively on us and others. To put it more simplistically, if we want to be happy, we should be good. This paper looks for supporting evidence for such claims using a corpus-based approach, by looking into the relationship between moods (as a proxy for well-being) and value priorities in blog posts.

We make the assumption that people’s mood as well as the type of moral values they prioritize can be detected through their writing, and moreover that moral values and mood expression can

be detected through lexical content alone, something which has already been verified empirically for mood [5]. The questions that interest us are to what extent a relationship exists between the morality and mood of a blogger, and to what extent such a relationship varies in bloggers of different age and gender, as reflected through their blog posts.

2 Models for moods and moral values

Our model or typology for moods¹ is shown in figure 1. This typology is based on a model of emotion as a multicomponent process [9]. In this model, the distribution of the affective states is the result of analysing similarity judgements by humans for a set of emotion terms using cluster-analysis and multidimensional scaling techniques to map out the structure as a two-dimensional space. For present purposes, we are primarily interested in the distinction between positive moods, represented by the left half (quadrants 3 and 4), and negative moods, represented by the right half (quadrants 1 and 2)². The positioning of words in this space is somewhat ‘fuzzy’; an affective state such as “angry” to describe facial expression in speech may have a slightly different location than an “angry” blog post. Nevertheless this model has been shown to correlate well with automatic measurement of affect in texts [3], so it is well suited for our purpose.

Our reference model for value types is based on [11], as shown on Figure 2. Schwartz derived ten motivational types of values from the universal requirements of human existence and further research offers considerable evidence to support the comprehensiveness of the model [12, 13, 1]. This model synthesizes the results of

1. We are treating “mood”, “affective state” and “emotion” as the same thing.

2. The division between active and passive moods is shown only for completeness, as it is part of Scherer’s model.

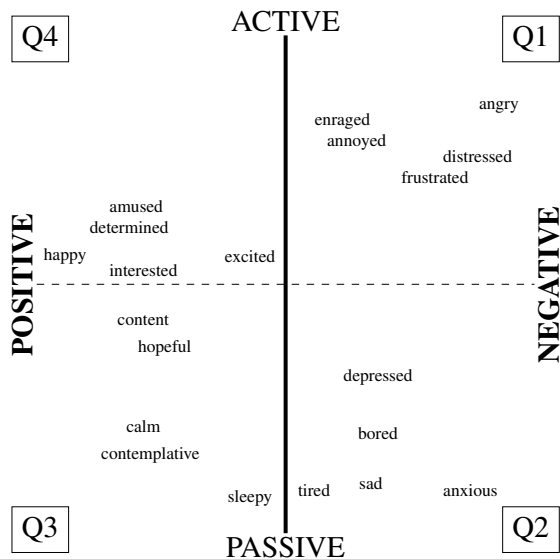


FIGURE 1 – Typology of affective states based on [9].

cross-cultural research in more than 200 samples from over 60 countries. Competing value types emanate in opposing directions from the center; complementary types are in close proximity going around the circle³. Subsequently, Sagiv and Schwartz [8] carried out a study to investigate the relationships of value priorities to measures of subjective well-being (1261 subjects). Their results provided some of the first evidence for what we are trying to show : a direct association between basic value priorities (as measured by lexical frequency) and the affective component of subjective well-being (as measured by self-annotated moods). *Achievement*, *self-direction* and *stimulation* values correlated positively with well-being on the pure affective index, while *tradition*, *conformity* and *security* values correlated negatively. The positive correlations are shown in the dotted pattern in figure 2 and negative correlations in the checkered pattern. One of the three types of measures of subjective well-being used in the study was the *Bradburn affect scale*, a measurement index using an affective vocabulary similar to the moods in figure 1. Our first experiment will try and replicate these findings in blog posts, using Scherer's model to measure well-being. As we already mentioned, earlier studies have verified that this model is a good description

3. In [8], the reliabilities for the value types across samples were : universalism 0.73 ; benevolence 0.68 ; tradition 0.49 ; conformity 0.64 ; security 0.64 ; power 0.66 ; achievement 0.65 ; hedonism 0.71 ; stimulation 0.61 ; self-direction 0.58.

of the distribution of affect in a two-dimensional space : in other words, quadrants in the model are clearly separable. For our purpose, it is sufficient that we can separate *positive* (i.e. quadrants 3 and 4) from *negative* (i.e. quadrants 1 and 2) moods.




FIGURE 2 – Theoretical model of relations among ten motivational types of values based on [8]. The positive correlations are shown in the dotted pattern and negative correlations in the checkered pattern.

3 Corpora

For our experiments we will use two corpora. The first corpus is called *LiveJournal* and was harvested through the LiveJournal⁴ website during the year 2005 for a total of 28,672 blog posts (mood-annotated by authors), belonging to one of the four quadrants from figure 1⁵ :

Quadrant 1 (4777) annoyed (3607) frustrated (445) angry (329) distressed (187) enraged (75) jealous (29) envious (24) bitter (20) disgusted (19) impatient (16) hateful (6) tense (5) alarmed (4) defiant (4) suspicious (3) discontented (2) insulted (1) startled (1)

4. <http://www.livejournal.com>

5. The location of top 5 (underlined) moods is shown in figure 1 and numbers indicate how many posts in the corpus were tagged with this mood. The LiveJournal corpus was not cleaned and may include ill-formed English (or no English at all), advertising and other boiler-plate material.

Quadrant2 (7971) tired (2372) bored (1360) depressed (907) sad (821) anxious (654) lonely (392) disappointed (355) worried (303) apathetic (273) gloomy (244) uncomfortable (157) embarrassed (102) miserable (14) ashamed (7) desperate (6) dissatisfied (1) doubtful (1) droopy (1) hesitant (1)

Quadrant 3 (7964) content (1838) calm (1401) contemplative (1302) sleepy (1266) hopeful (552) relaxed (354) pleased (324) satisfied (286) pensive (236) peaceful (235) impressed (129) serious (13) glad (9) serene (9) astonished (3) at ease (2) reverent (2) attentive (1) friendly (1) longing (1)

Quadrant4 (7960) happy (3009), amused (2726), excited (1689), determined (480), interested (13), triumphant (13), adventurous (5), ambitious (5), delighted (5), expectant (5), enthusiastic (4), joyous (3), aroused (2), passionate (1)

As self-annotation of moods on LiveJournal was optional, we can have a certain degree of confidence that the mood indicated by the author truly reflects his/her true mood. We filtered out posts with less than 1000 characters, leaving 3139, 5522, 5999 and 5350 for quadrants 1, 2, 3 and 4 respectively. The average post length is 508 words. The corpus is available for download⁶.

The second corpus is the *Blog Authorship Corpus (BAC)*⁷ [10]. This corpus consists of posts from 19,320 bloggers gathered from *blogger.com* in August 2004. The corpus incorporates a total of 681,288 posts and over 140 million words - or approximately 35 posts and 7250 words per person. Each blog indicates the blogger's self-provided age, so that we have three groups : 8240 "10s" blogs (ages 13-17), 8086 "20s" blogs (ages 23-27) and 2994 "30s" blogs (ages 33-48). For each age group there are an equal number of male and female bloggers. Each blog in the corpus includes at least 200 occurrences of common English words, so each post is more likely to include garbage-free texts than the LiveJournal corpus. This corpus is a good candidate for our study because it is rather clean, balanced and provides a larger sample of posts for each individual blogger than the LiveJournal corpus (for which there is only one post per individual blogger).

6. <http://www.clul.ul.pt/bigfiles/LJcorpus.tar.gz>

7. <http://u.cs.biu.ac.il/~koppel/ BlogCorpus.htm>

4 Experiments

In the following two experiments we are interested in two questions : (1) the relationships that may or may not exist between moral values and moods, and by extension, the well-being of people (2) the effect of age and gender on moral values interest and priority.

Experiment 1

To answer the first question, a corpus-based approach, reminiscent of techniques used in sentiment analysis [6], comes down to looking at the distribution of terms associated with a certain moral value (e.g. *benevolence*) in a mood-annotated blog corpus. We investigate this potential relationship for the LiveJournal corpus. In what follows we present the list of terms used to "model" each type of moral value⁸ :

POWER : power (3.9), authority (4.1), wealth (9.8)

ACHIEVEMENT : achievement (5), success (3.4), ambition (7.9), influence (5.2)

UNIVERSALISM : universal (5.7), understanding (3.3) appreciation (4.6), tolerance (7.2) wisdom (6.3) equality (6.3)

BENEVOLENCE : kindness (8.2), helpful (5.8), honesty (6.6), forgiving (9.8), loyal (7.9), responsible (3.4)

STIMULATION : stimulation (8.2), excitement (6.4), novelty (8.7), challenge (4.4)

SELF-DIRECTION : autonomy (4.7), creativity (5.0), freedom (4.8), independence (5.0), curiosity (8.0)

TRADITION : tradition (5.3), humility (9.1), devotion (6.2), moderation (6.7)

CONFORMITY : conformity (5.6), politeness (8.4), obedience (9.1)

SECURITY : security (3.7), safety (2.7), stability (4.0), clean (3.9)

HEDONISM : pleasure (4.7), enjoyment (5.8), gratification (9.8)

The number associated with each term is the inverse document frequency (idf)⁹ value that we used as a weight for each term in the computation of the ten scores per document (one for each

8. From table 1 on p. 179 of [8].

9. Idf values are taken from the MEAD (<http://www.summarization.com/mead/>) summarization system default database.

VALUE	P/N	POS	NEG	St. error
achievement ✓	1.41	1.07	0.76	±0.09
stimulation ✓	1.29	1.98	1.53	±0.18
power	1.27	6.16	4.84	±0.34
hedonism	1.27	1.70	1.34	±0.16
universalism	1.25	2.07	1.65	±0.13
self-direction ✓	1.25	1.69	1.36	±0.12
tradition	1.00	0.54	0.54	±0.09
security ✓	0.96	5.19	5.40	±0.26
benevolence	0.89	1.86	2.08	±0.14
conformity ✓	0.38	0.06	0.16	±0.04
Average	1.10	2.23	1.97	-

TABLE 1 – Value-bearing content per 10,000 words for each half of the LiveJournal corpus. ✓ indicates agreement with figure 2. P/N is number of Positives divided by number of Negatives.

moral value), to mitigate the effect of chance occurrence. To give an example of the calculation of a score, a blog post with two occurrences of the word “obedience” and one occurrence of the word “security” and “safety” would get a score of 18.2 (2×9.1) for the moral value *conformity*, 6.4 ($3.7 + 2.7$) for the moral value *security* and zero for all other eight values. All scores are subsequently normalized to account for the fact that moral values are modelled with a different number of terms, and that documents have different length. This normalization allows for comparison across values but also across blog types (polarity, gender and age). We do not take into account phenomena such as irony, understatement, reported speech as well as lexical variations and negation¹⁰. Studies like [7] showed that it is indeed difficult to compute the scope of *not*, which is moreover not always negative in context : e.g. *not only*, *not just*, *not to mention*, etc.

Table 1 shows the value-bearing content scores for POSitive and NEGative blog posts from the LiveJournal corpus. Positive scores include all blog posts annotated with a mood from quadrants 3 and 4, while negative scores include posts from quadrants 1 and 2. For example, positive posts have an average score of 6.16 (per 10,000 words) for the value “power”. Moral values with the highest ratio Pos/Neg are presented first in the table : a ratio > 1 indicates a correlation of the associated moral value with positive moods. The differences between the means of scores shown here are all significant at $p < .0001$. The results from table 1 are consistent with those of

10. For example, *I hate conformity* or *I do not like conformity* is also counted as one occurrence of *conformity*.

VALUE	10s	20s	30s
power	5.44 ±.20	9.62 ±.30	12.48 ±.64
security	5.00 ±.20	6.54 ±.18	9.29 ±.44
universalism	1.78 ±.09	3.32 ±.11	4.58 ±.36
self-direction	1.92 ±.11	3.60 ±.16	4.43 ±.29
stimulation	1.71 ±.11	3.47 ±.15	3.86 ±.26
benevolence	1.90 ±.10	2.79 ±.11	3.65 ±.23
hedonism	1.26 ±.10	2.56 ±.14	3.32 ±.30
achievement	1.22 ±.08	2.20 ±.11	2.61 ±.17
tradition	0.63 ±.06	1.19 ±.09	2.04 ±.50
conformity	0.20 ±.04	0.54 ±.12	0.40 ±.08
Total	21.06	35.82	46.65

TABLE 2 – Value-bearing content per 10,000 words for the BAC by age.

[8] highlighted in figure 2 above. Only *tradition* cannot be classified clearly as it has a ratio of one between positive and negative posts. This suggests that when blog posts are self-annotated for moods, they represent a reliable indicator of a person moral character, as moods and moral values (as found in blog posts) correlate in almost exactly the same way as when they are measured more directly in studies from psychology. We can also see that positive posts tend to make use more frequently of terms linked to morality (2.23 to 1.97), which might suggest that people in a good/positive mood (happy people ?) are more prone to discuss morality. Finally, in [8], the moral values *power*, *hedonism*, *universalism* and *benevolence* could not be conclusively placed in one of the poles : our results suggest that *power*, *hedonism* and *universalism* are favoured by happy (positive) people while *benevolence* is mostly a preoccupation for sad (negative) people.

Experiment 2

We now turn to our second research question, which investigates the effect of age and gender on the “moral vocabulary” and moral value priority of bloggers. Previous studies provide some empirical support for the *Gender-as-Culture* hypothesis [2]. The BAC is well-suited for this task, as it is annotated for age and gender, and comprises a few blog post entries for each blogger, which allows for a more complete picture of the blogger personality. The value-bearing content scores for each blogger are computed as previously.

Table 2 shows the value-bearing content scores for each age category, and reveals that older

VALUE	female	male
power	5.92 \pm 0.20	10.64 \pm 0.20
security	6.20 \pm 0.15	6.42 \pm 0.17
self-direction	2.56 \pm 0.12	3.46 \pm 0.12
universalism	2.46 \pm 0.09	3.26 \pm 0.10
stimulation	2.63 \pm 0.17	2.93 \pm 0.17
benevolence	2.40 \pm 0.10	2.69 \pm 0.11
achievement	1.43 \pm 0.10	2.26 \pm 0.10
hedonism	2.02 \pm 0.18	2.22 \pm 0.20
tradition	0.90 \pm 0.14	1.26 \pm 0.15
conformity	0.26 \pm 0.20	0.50 \pm 0.20
Total	22.77	35.63

TABLE 3 – Value-bearing content per 10,000 words for the BAC by gender.

people are increasingly preoccupied with moral questions than younger people (from a score of 21.06 to 46.65). For each moral category (with the except of *conformity*), there is a monotonic increase in the score with age. With respect to value priorities, it seems that people retain the same moral views over time. Nevertheless, they become less altruistic (*benevolence* drops two ranks) and embrace a more holistic view of life (*universalism* steps up two ranks).

Table 3 shows the value-bearing content scores by gender. Given the total scores (23 versus 36), the male discourse appears more preoccupied with morality than the female, and this is particularly obvious for *power* and *achievement*, but also *conformity*. Females prioritize *security* and males *power*, while *stimulation* has less priority for males than females. Considering common-sense psychology, these results are not surprising, say archetypal¹¹.

5 Conclusion

This paper has explored the relationship that may exist between a set of moral values and the well-being of people. We have hypothesized that the moral values talked about by bloggers in their posts is a reliable proxy of their moral character. We have found support for the claim that *achievement*, *stimulation* and *self-direction* are linked to happiness, as opposed to *security* and *conformity*. This is also supported in the psychology literature.

We have also looked into the effect of age and gender on the “moral vocabulary” and moral

value priority of bloggers. Our findings are that older people are more preoccupied by morality than younger people, and that men do seem more talkative about these moral issues than women. Despite small differences, value priorities remain fairly stable among “positive” or “negative” people and across gender and age.

Given the relation between the “moral vocabulary” and age or gender, we think that these results can advantageously assist the type of automatic genre and age identification as put forwards in [8] and based on style and content. The relation between well-being and moral values may also be a good starting point to build an automated system to help draw a psychological profile of people according to their writing. Moreover, since the layout of figure 2 appears to mimic quite closely the political divide between left and right, it would be very interesting to see what the relation between left and right blog posts (instead of positive and negative) and moral values is.

Références

- [1] A. Bardi, R.M. Calogero, and B. Mullen. 2008. *A new archival approach to the study of values and value-behavior relations : Validation of the value lexicon*. Journal of Applied Psychology, 93 :483-497.
- [2] A. Mulac, J.J. Bradac and P. Gibbons. 2001. *Empirical support for the gender-as-culture hypothesis : An inter-cultural analysis of male/female language differences*. Human Communication Research 27 :121-152.
- [3] M. Génereux and R. Evans. 2006. *Towards a validated model for affective classification of texts*. In Sentiment and Subjectivity in Text, Workshop at the Annual Meeting of the Association of Computational Linguistics (ACL 2006), Sydney, Australia, July 22, 2006.
- [4] S. Harris. 2010. *The Moral Landscape. How science can determine human values*. Free Press.
- [5] G. Mishne. 2005. Experiments with mood classification in blog posts. In Proceedings of the 1st Workshop on Stylistic Analysis Of Text For Information Access (Style 2005), Brasil, 2005.
- [6] B. Pang and L. Lee. 2008. *Opinion mining and sentiment analysis. Foundations and Trends in Information Retrieval, 2* :1-135.
- [7] L. Jia, C. Yu and W. Meng. 2009. *The effect of negation on sentiment analysis and retrieval effectiveness*. Proceeding of the 18th ACM conference on Information and knowledge management (CIKM), p. 1827-1830, 2009.
- [8] L. Sagiv and S. H. Schwartz. 2000. *Value priorities and subjective well-being : direct relations and congruity effects*. European Journal of Social Psychology, 30 :177-198.
- [9] K. R. Scherer. 1984. *Emotion as a multicomponent process : A model and some cross-cultural data*. P.

11. As pointed out by an anonymous reviewer.

- Shaver (Ed.) Review of Personality and Social Psychology, 5 :37-63.
- [10] J. Schler, M. Koppel, S. Argamon, and J. Pennebaker. 2006. *Effects of age and gender on blogging*. In Proceedings of 2006 AAAI Spring Symposium on Computational Approaches for Analyzing Weblogs.
- [11] S.H. Schwartz. 1992. *Universals in the content and structure of values : Theoretical advances and empirical tests in 20 countries*. Advances in experimental social psychology, 25 :1-65.
- [12] S.H. Schwartz. 1994. *Are there universal aspects in the content and structure of values ?* Journal of Social Issues, 50 :19-46.
- [13] S.H. Schwartz. 1995. *Identifying culture-specifics in the content and structure of values*. Journal of Cross-Cultural Psychology, 26 :92-116.