Thematic Analysis of Words that Invoke Values in the Net Neutrality Debate

Kenneth R. Fleischmann,[1] Yasuhiro Takayama,[2] An-Shou Cheng,[3] Yoichi Tomiura,[4] Douglas W. Oard,[5] Emi Ishita[4]

1: University of Texas at Austin, USA; 2: National Institute of Technology, Tokuyama College, Japan; 3: National Sun Yat-Sen University, Taiwan; 4: Kyushu University, Japan; 5: University of Maryland, College Park, USA

Abstract
This paper describes an initial analysis of the association of specific vocabulary choices with the invocation of human values in testimonies prepared for public hearings about Net neutrality in the United States. Motivation for this work comes from an interest in understanding what people value and how they express those values in writing. Related work includes research on human values from fields ranging from social psychology to advertising to human-computer interaction. First, human annotators used closed coding to identify human values in testimonies based on a prior meta-analysis of human values. Next, a "values dictionary" was automatically learned that identifies words that are strongly associated with sentences that human annotators coded as being related to specific values. Finally, an open-ended thematic analysis was conducted. The contribution of the paper is to enhance our understanding of how human values are expressed, as well as to introduce and evaluate a new automated tool for facilitating social science research.
Keywords: human values, Net neutrality, thematic analysis, automated content analysis, natural language processing.
Citation: Editor will add citation with page numbers in proceedings and DOI.

Contact: kfleisch@ischool.utexas.edu

1 Motivation

What do people value, and how do they express those values in text? Understanding users' values is critical for ensuring that technology can be designed to be sensitive to their values [11]. Research on natural language processing tasks such as sentiment analysis [14, 23] has demonstrated the capability to detect some private states [32] in texts. In this paper, we focus specifically on human values. According to Cheng and Fleischmann [3], "values serve as guiding principles of what people consider important in life" (p. 2). Research on human values has spanned a wide range of fields, including anthropology, sociology, psychology, science and technology studies, information studies, business, and computer science. The next section reviews the history of research on human values.

In prior work, we have built systems that learned to assign human values to sentences. The key idea in the design of our initial system was to have human coders annotate each sentence in several documents with zero or more values, and then to learn the association between values and sequences of words. We then evaluated the accuracy of our systems by checking how well the resulting learned classification model works on sentences from other documents that had not been used to train the machine. This approach works well, correctly guessing with which values a well-qualified human annotator would annotate a sentence nearly as often would a second well-qualified human annotator [28]. In this paper we now explore the patterns that our systems learn, with an eye towards learning something about how people invoke values when making arguments. To do this, we developed a new classifier in which we require the classifier to associate each word with one or more human values, regardless of the context in which that word occurs. This results in a codebook (which we call a "values dictionary") that, for each human value, tells us which words are most strongly associated with that value. We then manually categorize the strongest learned associations; the resulting categories in turn allow us to characterize the actions of our classifier in ways that yield insights into how people invoke or reflect specific values.

2 Related Work

What motivates human attitudes and behaviors? Although there are certainly many factors at play, human values play an important role in explaining attitudes and behaviors. Whereas attitudes are specific to a particular issue or situation, values transcend those specific circumstances and apply to many aspects of everyday life [26]. Research has shown that human values can be measured through instruments such as surveys, revealing differences between countries and subpopulations in terms of their values [13, 24, 25, 26, 27]. These differences are in turn correlated with differences in attitudes and behavior, such as attitudes favoring immigration, which are positively correlated with the value of universalism and negatively correlated with the values of security and conformity [26]. Similarly, in the Park51 debate (commonly referred to as the "Ground Zero Mosque"), support for the project was positively correlated with universalism and negatively correlated with security [30]. Finally, Cheng and colleagues [4] found that, in the Net neutrality debate, support of Net neutrality was positively correlated with valuing innovation and negatively correlated with valuing wealth.

Values are tightly connected to how people use technology. Friedman, Howe, and Felten [10] found that the value of privacy was supported by technologies such as cookie management, while Friedman and colleagues [12] found that the value of privacy was challenged by technologies such as webcams in public spaces. One example of such studies is exploration of the use of technologies by the homeless, and how they are connected to value differences [18, 19, 20, 22, 33]. Information technology developers typically have positive intentions, but these intentions may not lead to software that is sensitive to users' values unless software developers have a good understanding of what their users value [7, 9]. Thus, values are a critical concept for understanding how users accept and adopt new information technologies.

A wide range of methods have been applied to study human values. As noted above, the most commonly used approach has been surveys. However, Fleischmann and colleagues [8] point out that although surveys have many strengths as a way of measuring human values, they are often subject to limitations such as response bias, retrospective bias, self-selection, and difficulty in accessing some populations. Content analysis is less subject to these issues, although there are also limitations to that method, such as the differing interpretations that different annotators might apply to the same text. Cheng and colleagues [4] apply content analysis to the Net neutrality debate, providing an example of coding values in formal texts. Koepfler and Fleischmann [19] apply content analysis to the use of Twitter by the homeless, providing an example of coding values in informal texts. Crowdsourcing can also be used to broaden the range of participation in the coding effort [30]. It is then possible to automate either manual [15, 28] or crowdsourced [31] content analysis, providing benefits in terms of scalability and consistency. This paper describes one such effort to automate content analysis and then derives new insights into how human values are expressed in text.

3 Methods

In prior work we have described the process that we used to manually annotate 102 prepared testimonies related to Net neutrality, the complete set of prepared testimonies for Senate, House, and FCC Hearings on Net neutrality [2]. Briefly, we started with a meta-inventory of human values that had originally been developed by Cheng and Fleischmann [3] from many pre-existing values inventories and we iteratively tailored a domain-specific values inventory through four rounds of coding. This process resulted in a set of six value categories relevant to the Net neutrality debate that human coders could reliably code: freedom, honor, innovation, justice, social order, and wealth [2].

Specifically, the third author coded all 102 prepared testimonies for values at the sentence level. Coding involved detecting explicit and implicit invocations of values, and included statements that expressed positive, negative, and neutral sentiment toward those values. A sentence could reflect multiple values, or none. Two coders independently coded twenty of these testimonies. Substantial agreement was achieved for freedom, innovation, social order, and wealth, and moderate agreement was achieved for honor and justice [2].

We have previously demonstrated the ability to train classifiers for the assignment of values to sentences by using a suite of support vector machines, one for each value [28, 29]. The resulting systems agree with human annotations nearly as often as another human annotator would for five of the six values (specifically, all values other than honor, which is the value least often annotated in our collection). For the analysis in this paper we therefore report results over five values, omitting honor.

Support vector machines can be effective, but they reason very differently than people do and thus they are not well suited to the development of explanations for their results. For this paper, the second author developed a new classifier with a simple and easily interpreted classification rule. For each single word or each pair of words we learn an entry in a "values dictionary" that simply tells us which values to assign. Then whenever we see that word or pair of adjacent words in a sentence, we assign the corresponding values to that sentence. To avoid guessing every value for every sentence, the values dictionary must be sparse – it must assign no values to most words. Given enough time, we could try every possible values dictionary (i.e., every combination of values for every word – for example, maybe when we see "financial" in a document we should guess the values wealth and social order?) and then pick the best one (by checking to see how well each values dictionary would work on some subset of our annotated interviews that we use for "training"). In practice, we use a simple machine learning technique, Metropolis with simulated annealing [17], to efficiently explore the space of reasonable possibilities. This technique starts from a reasonable guess (the values we see most often with each word in annotated training sentences) and then iteratively tries small variations in the entries of the values dictionary until no further improvement can be found (e.g., maybe seeing "financial" in a document should only cause us to guess the value wealth).

Because single words may not be sufficiently informative, we also learn to associate adjacent two-word pairs with values. Then when we see a two-word pair (e.g., financial market) we guess the values associated with that pair (e.g., wealth) rather than the values associated with each word in the pair individually. Because we can accumulate more evidence for training if we treat variants of the same word in the same way, we make only a single entry in our values dictionary for each word stem (e.g., market is the stem of market, markets, and marketing) or pair of stems; this is equivalent to making identical entries for any words or word pairs that share the same stem(s). For the sake of simplicity and brevity, this paper only reports results from the single word associations, while future work may further explore the results from two-word pairs.

Rather remarkably, this simple approach yields results that are about as accurate as those achieved by a support vector machine. One way to measure the degree of correctness of a classifier is to compute the balanced harmonic mean of recall (are missed detections rare?) and precision (are hallucinated detections rare?). This is the so-called F_1 measure. Our new method achieves $F_1 = 0.7114$, which is about the same as the $F_1 = 0.7068$ that our earlier support vector machine implementation achieved [28]. These results are averages across 8,660 sentences based on 102-fold document cross-validation (in which we use every sentence in some set of 101 documents for training and then test on the sentences in the 102nd document, repeating the process 102 times with different sets of training and test documents). Unlike a support vector machine, however, our new process produces a values dictionary that we can now analyze.

Our goal in this paper is to use the resulting values dictionary to now explore what our system has learned about how words are used to reflect values. We have applied thematic analysis [1] to identify themes within each of these sets of words. The first and third authors independently conducted open coding of the words (actually, word stems) related to each value, and identified categories that help to organize those lists. The first author compared the two lists and reconciled them, and mutual agreement was reached on the resulting composite organization, which reflected portions of both of the coders' categories.

4 Results

This section presents the results from analysis of each of the five values by grouping their associated word lists into categories. As an example, Table 1 provides the word groupings for the value freedom.

Table 1: Word Groupings for the Value of Freedom

Choice	Restriction	Centralization	Decentralization	Private Sector	Public Sector
choice deregulation freedom freely liberty open permission unfettered unimpeded unregulated	barrier censorship control gatekeeper impede inhibit interference obstacle preclude	concentration consolidation bundle	decentralized dissemination	anti-competitive anticompetitive competition duopoly pro-competitive	democracy democrat king

Please note that the annotators also put some words into a junk category, not shown here, if they were not able to identify a theme that fit that word. For example, for freedom, 38 of the 70 words (54%) were put in the junk category. Examples of junk category words for freedom include "edit" and "greatly." Such cases are natural in any technique that is based on statistical associations – some associations are interesting, while others arise simply from chance (or possibly from factors that might be interesting, but that we have not yet been able to recognize).

Table 2 displays the resulting categories for each of the five values. There are a total of 19 categories, with two categories (private sector and public sector) occurring within the word lists for three of the values (freedom, social order, and wealth).

Table 2: Themes Identified for Each of the Five Values

Freedom	Innovation	Justice	Social Order	Wealth
Choice Restriction Centralization Decentralization Private Sector Public Sector	Innovator Technology Creativity Evolution Entrepreneur	Fairness Unfairness Censorship Openness	Private Sector Public Sector	Riches Poverty Buyer Seller Private Sector Public Sector

5 Discussion

Interestingly, polar opposite pairs, which we refer to as value conflicts [9], are immediately perceptible for four of the five values. Specifically, freedom involves choice vs. restriction, centralization vs. decentralization, and private sector vs. public sector. Justice involves fairness vs. unfairness and censorship vs. openness. Social order involves private sector vs. public sector. Wealth includes riches vs. poverty, buyer vs. seller, and private sector vs. public sector. The only exception is innovation, which appears to have only the one-sided categories innovator, technology, creativity, evolution, and entrepreneur.

For freedom, it is not surprising that choice vs. restriction was one of the salient value conflicts, since it is a fundamental aspect of freedom, particularly in relation to the Net neutrality debate. Centralization vs. decentralization is also a logical pairing, as it has clear implications for freedom. Finally, the private sector vs. public sector conflict is also a common one in general (as illustrated in how frequently it occurred) and specifically within the issue of freedom, where some people would see the public sector as embodying and protecting freedom while others would see the private sector as doing so.

For justice, we found a value conflict between fairness and unfairness, one of the fundamental concepts for justice. Another dimension was censorship vs. openness, which is a particularly salient dimension of the Net neutrality debate, because many of the debates involving Net neutrality involve the issue of censorship, with the concern that corporations might attempt to censor the Web content they provide.

For social order, we again found the private sector vs. public sector conflict, which makes sense given that each of these sectors of society can contribute in different ways to social order.

For wealth, we found a value conflict between riches and poverty, which is the most fundamental aspect of financial inequality. Another value conflict was buyer vs. seller, which is the most fundamental pairing of actors in business transactions. Finally, we again found the private sector vs. public sector conflict, which fits into wealth since budgetary considerations are critical within both.

Contrastingly, for innovation, we did not find value conflicts. Rather, we found the rather one-sided categories of innovator, technology, creativity, evolution, and entrepreneur. Innovator and entrepreneur are two of the critical roles within innovation (reflecting innovation in both technology and in business practices). Technology is the main focus of the Net neutrality debate. Creativity is a critical component of innovation, as being creative is a necessary but not sufficient for being innovative. Finally, evolution denotes change, a critical aspect of innovation.

It is not particularly surprising that innovation was the one value that did not consist of value conflicts, given the national context of the corpus. The corpus involved testimonies prepared for hearings held by the Senate, House, and FCC, part of the legislative and executive branches of the United States federal government. Although more research would be needed to understand what differences might occur within different national contexts, it seems reasonable to assume that this finding for innovation may be nation-specific, as the U.S. historically embraces innovation in a way not seen in cultures more resistant to technological change, as evidenced by democratic interventions such as science shops in Northern Europe [6], as well as in societies where conformity is traditionally valued over individuality, as in East and Southeast Asia [16].

6 Contribution

This paper makes two key contributions to the information field. First, it presents, demonstrates, and evaluates (qualitatively) the effectiveness of a new method that combines manual and automatic content analysis as well as both closed and open coding. This method can be applied to a wide range of purposes, such as, for example, as a tool for training annotators, or as in this study, as a new tool for furthering social science analysis. Although the effectiveness of this method for other domains beyond Net neutrality and for other phenomena beyond values would need to be tested, the success in deploying the method in this study is a promising sign of its potential for broader use.

In addition, this paper contributes to our understanding of the value conflicts involved in the Net neutrality debate, including how different values may be framed differently by different actors [5]. Previous research has explored value framing in relation to nuclear power [30] and homelessness [21]. This study expands to a new domain, Net neutrality, and also employs a novel method that could be used in those as well as other domains. As such, the paper has the potential to influence both social science and computational research.

References

1. Braun, V., & Clarke, V. (2006). Using thematic analysis in psychology. *Qualitative Research in Psychology, 3*, 77-101.
2. Cheng, A.-S. (2012). *Values in the Net Neutrality Debate: Applying Content Analysis to Testimonies from Public Hearings*. Doctoral Thesis, University of Maryland.
3. Cheng, A.-S. & Fleischmann, K.R. (2010). Developing a meta-inventory of human values. *Proceedings of the 73rd Annual Meeting of the American Society for Information Science and Technology*, Pittsburgh, PA.
4. Cheng, A.-S., Fleischmann, K.R., Wang, P., Ishita, E., & Oard, D.W. (2012). The role of innovation and wealth in the net neutrality debate: A content analysis of human values in Congressional and FCC hearings. *Journal of the American Society for Information Science and Technology*, 63(7), 1360-1373.
5. Chong, D., & Druckman, J.N. (2007). Framing theory. *Annual Review of Political Science, 10*, 103-126.
6. Farkas, N. (2002). *Bread, Cheese, and Expertise: Dutch Science Shops and Democratic Institutions*. Doctoral Thesis, Rensselaer Polytechnic Institute.
7. Fleischmann, K.R. 2014. *Information and Human Values*. San Rafael, CA: Morgan & Claypool.
8. Fleischmann, K.R., Oard, D.W., Cheng, A.-S., Wang, P., & Ishita, E. (2009). Automatic classification of human values: Applying computational thinking to information ethics. *Proceedings of the 72nd Annual Meeting of the American Society for Information Science and Technology*, Vancouver, BC, Canada.
9. Fleischmann, K.R. & Wallace, W.A. (2010). Value conflicts in computational modeling. *Computer*, 43(7), 57-63.
10. Friedman, B., Howe, D.C., and Felten, E. (2002). Informed consent in the Mozilla browser: Implementing value-sensitive design. *Proceedings of the 35th Hawai'i International Conference on Systems Sciences*, Waikoloa Village, HI.
11. Friedman, B., Kahn, P.H., Jr., & Borning, A. (2006). Value sensitive design and information systems. In P. Zhang & D. Galletta (eds.), *Human-Computer Interaction in Management Information Systems: Foundations*. New York: ME Sharpe.
12. Friedman, B., Kahn, P.H., Jr., Hagman, J., Severson, R.L., & Gill, B. (2006). The watcher and the watched: Social judgments about privacy in a public place. *Human-Computer Interaction, 21*, 235-272.
13. Hofstede, G. (2001). *Culture's consequences: Comparing values, behaviors, institutions, and organizations across nations*. Thousand Oaks, CA: Sage.
14. Hopkins, D.J. and King, G. (2010). A method of automated nonparametric content analysis for social science, *American Journal of Political Science* 54(1), 229-247.

15. Ishita, E., Oard, D.W., Fleischmann, K.R., Cheng, A.-S., & Templeton, T.C. (2010). Investigating multi-label classification for human values. *Proceedings of the 73rd Annual Meeting of the American Society for Information Science and Technology*, Pittsburgh, PA.
16. Kim, D., Pan, Y., & Park, H.S. (1998). High- versus low-context culture: A comparison of Chinese, Korean, and American Cultures. *Psychology and Marketing 15*, 507-521.
17. Kirkpatrick, S., Gelatt, C.D., & Vecchi, M.P. (1983). Optimization by simulated annealing. *Science, 220*, 671-680.
18. Koepfler, J.A. (2014). *Values and self-presentation in online communication by stakeholders related to homelessness*. Unpublished Doctoral Dissertation, University of Maryland.
19. Koepfler, J.A., & Fleischmann, K.R. (2012). Studying the values of hard-to-reach populations: Content analysis of tweets by the 21st Century homeless. *Proceedings of the 7th Annual iConference*, Toronto, ON.
20. Koepfler, J.A., Shilton, K., & Fleischmann, K.R. (2013). A stake in the issue of homelessness: Identifying values of interest for design in online communities. *Proceedings of the 6th International Conference on Communities and Technologies*, Munich, Germany.
21. Koepfler, J.A., Templeton, T.C., & Fleischmann, K.R. (2012). Exploration of values and frames in social media texts related to the Homeless Hotspots debate. *Proceedings of the 75th Annual Meeting of the American Society for Information Science and Technology*, Baltimore, MD.
22. Le Dantec, C.A., & Edwards, W.K. (2008). Designs on dignity: Perceptions of technology among the homeless. *Proceedings of the ACM SIGCHI Conference on Human Factors in Computing Systems (CHI)*, Florence, Italy.
23. Pang, B., & Lee, L. (2008). Opinion mining and sentiment analysis. *Foundations and trends in information retrieval, 2*, 1-135.
24. Rokeach, M. (1973). *The nature of human values*. New York: Free Press.
25. Schwartz, S.H. (1994). Are there universal aspects in the structure and contents of human values? *Journal of Social Issues, 50*(4), 19-45.
26. Schwartz, S.H. (2007). Value orientations: Measurement, antecedents, and consequences across nations. In R. Jowell, C. Roberts, R. Fitzgerald, & G. Eva (Eds.), *Measuring attitudes cross-nationally: Lessons from the European Social Survey*. London, England: Sage.
27. Shilton, K, Koepfler, J., & Fleischmann, K.R. (2013). Charting sociotechnical dimensions of values for design research. *The Information Society*, 29(5), 259-271.
28. Takayama, Y., Tomiura, Y., Ishita, E., Oard, D.W., Fleischmann, K.R., & Cheng, A.-S. (2014). A word-scale probabilistic latent variable model for detecting human values. *Proceedings of the ACM Conference on Information and Knowledge Management (CIKM)*. Shanghai, China.
29. Takayama, Y., Tomiura, Y., Ishita, E., Wang, Z., Oard, D. W, Fleischmann K. R, and Cheng, A.-S. (2013), Improving automatic sentence-level annotation of human values using augmented feature vectors. *Conference of the Pacific Association for Computational Linguistics (PACLING 2013)*, Tokyo, Japan.
30. Templeton, T.C., & Fleischmann, K.R. (2011). The relationship between human values and attitudes toward the Park51 and nuclear power controversies. *Proceedings of the 74th Annual Meeting of the American Society for Information Science and Technology*, New Orleans, LA.
31. Templeton, T.C., Fleischmann, K.R., & Boyd-Graber, J. (2011). Simulating audiences: Automating analysis of values, attitudes, and sentiment. *Proceedings of the 3rd IEEE International Conference on Social Computing*, Boston, MA.
32. Wilson, T., & Wiebe, J. (2005). Annotating attributions and private states. *Proceedings of the Workshop on Frontiers in Corpus Annotation II: Pie in the Sky*, Ann Arbor, MI.
33. Woelfer, J.P., Iverson, A., Hendry, D.G., Friedman, B., & Gill, B.T. (2011). Improving the safety of homeless young people with mobile phones: Values, form, and function. *Proceedings of the ACM SIGCHI Conference on Human Factors in Computing Systems (CHI)*, Vancouver, BC.