Moral Values Coding Guide

Joe Hoover^{1*}, Kate Johnson¹, Morteza Dehghani¹, Jesse, Graham¹

1. University of Southern California

Abstract

This guide is intended to be a manual for coding the moral content of natural language documents and a record of the USC Computational Social Science and Virtue Ideology and Morality labs' protocol for coding the moral content of social media and other texts. The general goal is for this guide to be a self-contained manual that can be used for training research assistants and establishing common practices for the identification and labeling of morally relevant sentiment in natural language. It currently includes an introduction to coding sentiment according to the Moral Foundations Theory framework.

Moral Values Coding Guide

Contents

Abstract	1
Moral Values Coding Guide	2
Introduction	3
Background	3
Moral Foundations Theory	4
Theoretical Background	4
Expressions of Moral Foundations in Text	5
Coding Moral Foundations Theory	6
General Text Coding for Sentiment Analysis	6
Introduction to Coding Moral Foundations Theory	6
Coding Moral Foundations Theory	9
CSSL & VIM lab Moral Foundations Theory Coding	11
References	14

Introduction

This guide is intended to be a manual for coding the moral content of natural language documents and a record of the USC CSSL and VIM labs' protocol for coding the moral content of social media and other texts. While it is currently under development, the general goal is for this guide to be a self-contained manual that can be used for training research assistants and establishing common practices for the identification and labeling of morally relevant sentiment in natural language.

Background

Moral values play many important roles in human social functioning. For example, they shape our behavior, our judgments, and the people and groups we affiliate with and distance ourselves from. However, despite the fact that moral values play such a large role in our lives, they are not always prominently visible in day-to-day life. For the most part, we cannot see someone's moral values simply by looking at them. While we often infer people's values from their behaviors, these inferences are often unreliable.

One way we make up for this invisibility is by expressing our moral values in language. For example, much of people's conversations revolve around describing and passing judgment on the behaviors and beliefs of others that we deem good or bad. Often, we use morally relevant language, or moral sentiment, to signal our values to others. Such signals can be ingenuous or disingenuous, but either way expressions of moral sentiment serve as informationally rich indicators of individuals' and groups' moral values, or at least the moral values they wish to display.

People's tendency to embed moral values in their daily language presents a valuable research opportunity for social scientists interested in moral values. For example, by measuring and tracking expressions of moral sentiment, we can gain insight into how moral values develop, transform, and spread. Moral sentiment analysis can also be a powerful tool for understanding social and cultural trends. Further, such research is particularly relevant and powerful today, given the rapid increase in digital content production and consumption that has been enabled by web-based technologies

and social media.

In order to take advantage of the massive amounts of data available online, it often is necessary to have sets of documents (e.g. Tweets or blog posts) that have been labeled for their moral content. Such sets of labeled documents can, for example, be used to test the accuracy of new methods or to train classifiers that can be used to predict the moral sentiment in new documents.

Moral Foundations Theory

Theoretical Background

Moral Foundation Theory (MFT) proposes that human moral values arise from a set of universal foundations with evolutionary origins (Graham, Haidt, & Nosek, 2009). While people's values vary substantially at multiple levels (e.g. between individuals, communities, cultures, etc.), MFT suggests that these values cluster into five basic bi-polar dimensions of moral concerns that represent the morally bad at one extreme and the morally good at the other. The descriptions below can be found at www.moralfoundations.org (*moralfoundations.org*, n.d.), as can suggestions for further reading.

1. *Care/Harm*: This foundation is related to our long evolution as mammals with attachment systems and an ability to feel (and dislike) the pain of others. It underlies virtues of kindness, gentleness, and nurturance.
2. *Fairness/Cheating*: This foundation is related to the evolutionary process of reciprocal altruism. It generates ideas of justice, rights, and autonomy.
3. *Loyalty/Betrayal*: This foundation is related to our long history as tribal creatures able to form shifting coalitions. It underlies virtues of patriotism and self-sacrifice for the group. It is active anytime people feel that it's "one for all, and all for one."
4. *Authority/Subversion*: This foundation was shaped by our long primate history of hierarchical social interactions. It underlies virtues of leadership and followership, including deference to legitimate authority and respect for traditions.

5. *Purity/Degradation*: This foundation was shaped by the psychology of disgust and contamination. It underlies religious notions of striving to live in an elevated, less carnal, more noble way. It underlies the widespread idea that the body is a temple which can be desecrated by immoral activities and contaminants (an idea not unique to religious traditions).

From the MFT perspective, the variation in values between people and cultures is thus a difference in emphasis – while we all have the capacity to moralize issues associated with each foundation, the extent to which we moralize different foundations varies according to a range of cognitive, psychological, and cultural factors. Further, while these foundations are theoretically independent, they often overlap in daily life. Thus, a particular issue, such as a soldier going AWOL, can be associated with subversion and betrayal.

Expressions of Moral Foundations in Text

When people write or speak about phenomena that they moralize, they tend to use specific kinds of words that signal moral relevance. While some of the most direct of such signals are general terms, such ‘right’, ‘wrong’, ‘good’ or ‘bad’, much of the language that we use indicates not only moral relevance, but also specific categories of moral relevance that align with MFT. For example, when a person describes an event that is perceived as a fairness violation, they might say something like ‘I can’t believe that happened. So unfair!’ or, slightly less explicit, ‘Ridiculous. They had NO RIGHT to do that.’

Accordingly, to the extent that we can identify linguistic cues that signal associations with different foundations, we can detect and quantify moral values by analyzing natural language artifacts (e.g. written documents and speech). In early work, Graham et al. (2009) demonstrated this by using frequencies of foundation-relevant words to measure differences in moral values sentiment between conservative and liberal sermons. Further, by applying more sophisticated text analysis methods, researchers have recently shown that moral sentiment analysis can be applied

to a range of applications and domains (Dehghani et al., 2016; Sagi & Dehghani, 2014)

Coding Moral Foundations Theory

General Text Coding for Sentiment Analysis. While text analysis methods have become advanced rapidly, the fact that they fall far short of human comprehension necessitates the implementation of performance evaluation. While any text analysis method will yield results, the reliability of these results can vary widely as a function of many different factors (e.g. the specific method algorithm, the domain of the text, the construct being measured, the validity of the chosen construct representation). Accordingly, it is essential that empirical work employing these methods demonstrate sufficient measurement validity, especially because psychological text analysis is a relatively new field with few methodological standards.

Currently, the gold standard for demonstrating sentiment analysis measurement validity involves comparing human perceptions of sentiment to the predictions generated by the method being used. For example, if a researcher is interested in detecting positive and negative sentiment in a set of documents, they might have human coders label a set of documents as either positive, negative, or neutral. They could then use this set of coded documents to evaluate the performance of a text analysis method by using the method (sometimes in combination with a classifier) to generate predictions of the sentiment expressed in a subset of the coded documents.

Introduction to Coding Moral Foundations Theory. While intuiting which foundation a particular moral concern is associated with can be quite easy in daily life, it can be surprisingly difficult when coding texts. Whether a text expresses loyalty or authority concerns, for example, can be highly ambiguous. We refer to this as general sentiment ambiguity.

Further, general sentiment ambiguity can be exacerbated by a different kind of ambiguity, which is generated by the difficulty of distinguishing between the moral sentiment expressed in a text and the likely moral sentiment intended by the author. While social scientists who rely on text analysis methods to study morality are

primarily interested in the latter, in reality, they only have access to the former. Of course, there is on average a strong correspondence between the values seemingly expressed in a text and the values endorsed by an author – indeed, computational sentiment analysis is founded on this assumption. However, inferring this correspondence often requires contextual cues that are not available to human coders or algorithms. For example, a social media message might simply state that the author thinks ‘Everything that is going on with abortion these days is reprehensible.’ It is clear that this is likely a morally relevant statement, but it is less clear what foundation this statement is relevant to. If we knew that the author was a secular liberal, we might assume that the author is concerned about violations of abortion rights and therefore infer with reasonable confidence that the foundation at play in this document is fairness/cheating; in contrast, if we knew that the author was a conservative Christian, we might assume that the author was expressing an anti-abortion sentiment, perhaps associated with purity/degradation. Unfortunately, it is often the case that little to no contextual information is available, so such informed inferences are generally not possible. The point is, the same expression can be meant to convey profoundly divergent moral sentiments and it is not infrequent that competing interpretations of a particular document are difficult if not impossible to resolve systematically.

Such ambiguities present a core challenge for text analysis, but their consequences are generally relegated to the error term without much need for concern as long as a certain threshold of performance is met. A more immediate problem is the effect of these ambiguities on human coder performance. Our personal experience with coding moral sentiment and in training research assistants to code moral sentiment has shown that it is often quite difficult to strike an acceptable balance between inferential coding and explicit coding. Because we typically do not have access to the authors of the texts we are analyzing – and sometimes not even to the discourse they are embedded in – we must constrain the degree to which we base document codes on what we think the author meant to express.

While this kind of theoretical/methodological responsibility might seem overly

conservative, excessive liberalness with context-based inference risks contaminating the data. Such inferences are prone to a range cognitive biases and can lead to artificially reduced variance of coded values. If each coder on a team relies too extensively on their own biases to infer the moral sentiment expressed in a document, coder reliability will almost certainly suffer. On the other hand, constraining interpretations of sentiment to a too literal plane can be equally problematic, because the subtleties of human language and morality will likely be lost. For example, an explicit coding of the text, "OMG I am going to MURDER THAT DOG, it is just TOO CUTE," references an act that is definitively relevant to the harm. However, taken in context, the profession of murderous intent is clearly idiomatic and it is not at all clear that this statement should be coded as equivalent to a statement like, "I hate him. He deserves to be murdered." Thus, a careful balance must be struck between inferential coding – coding that relies on context based assumptions – and explicit coding – coding that relies exclusively on literal interpretations of textual content. While this is difficult to do well and impossible to do perfectly, maintaining an awareness of these two extremes can hopefully limit coder biases in either direction.

Another important issue relevant to moral sentiment coding is the risk of artificially inflating coder agreement through over-training or excessively stringent coding schemes. When coding for MFT content, there are often disagreements about which foundation is relevant to a given statement and in many of these cases, even among high-level expert coders, it is not clear which perspective is correct. Of course, such disagreements can be resolved through discussion, and to an extent resolution is appropriate. However, it is our view that at some point, resolution of coder disagreement begins to artificially inflate inter-coder reliability. Moral values are inherently subjective phenomena and the true accuracy of a code cannot really be determined because there is no objective criterion to which it can be compared. The closest we can come to an objective criterion is consensus among some constituency. As consensus is approached, the certainty that a given phenomena has a strong subjective association with a specific moral foundation increases. This means that low consensus –

for example, among trained coders – is not simply a problem that must be resolved; it is a potential indication that the association between a foundation and a phenomenon might be subject to important boundary conditions, weak, or even illusory. Training coders so that they disagree minimally does not change this fact, but rather hides it. Consequently, while coders obviously need training, it is our view that training should focus on establishing a shared network of concepts and a few heuristics that can be used for generating codes, but also that this training should stop short of fabricating agreement.

Coding Moral Foundations Theory. When coding documents for Moral Foundations Theory (MFT) content, several initial decisions must be made. Often, the first decision regards what dimensions to code for. If a foundation-specific hypothesis will be tested, it might make sense to code only for that foundation. Frequently, though, it is necessary to code for multiple foundations. In such cases, the most obvious approach is to code for the presence of each of the five foundations. However, some research programs might require more fine-grained labels. While the poles of each dimension are related, they also express quite different sentiments that might be psychologically relevant. For example, the statement "We must end suffering," is probably not equivalent to the statement "We must provide kindness and compassion." It can thus also be useful to code for the presence of each pole of each foundation, which yields 10 individual codes. Additionally, it is important to code non-moral texts as such.

Another decision to make is whether to allow overlapping labels. In early coding sessions, we trained coders to label what seemed to be the primary sentiment in a document and also label any secondary moral sentiments that were expressed. However, during reliability analyses we found that while coders sufficiently agreed about whether a given moral sentiment was present, they agreed far less about primacy. Accordingly, we tend to just code for the presence/absence of each foundation.

There are, of course, many coding schemes that can be used for the above mentioned dimensions and the software that is used for displaying documents and storing codes can introduce idiosyncratic constraints. After several rounds of

experimentation, we have found that spread sheets provide an acceptable balance between efficiency, ease of use, and data management. To simplify the coding process, we represent foundations by the first letters of their poles, if we are coding them as holistic dimensions, or by the first letter of the virtues followed by a ‘p’ or ‘n’, which correspond to positive and negative, to indicate valence:

5 Foundations and non-moral:

- *Care/Harm*: CH
- *Fairness/Cheating*: FC
- *Loyalty/Betrayal*: LB
- *Authority/Subversion*: AS
- *Purity/Degradation*: PD
- *Non-moral*: NM

10 Foundations x Poles and non-moral:

- *Care*: CP
- *Harm*: CN
- *Fairness*: FP
- *Cheating*: FN
- *Loyalty*: LP
- *Betrayal*: LN
- *Authority*: AP
- *Subversion*: AN
- *Purity*: PP
- *Degradation*: PN
- *Non-moral*: NM

CSSL & VIM lab Moral Foundations Theory Coding. This section is intended to be a basic hand-guide for coding MFT sentiment as it is currently done in our labs.

Each moral virtue and vice should be coded as the capitalized first letter of the foundation and a capitalized P if the foundation is a virtue and a N if the foundation is a vice. P and N correspond to Positive and Negative. If a document does not have any moral content it should be coded as NM, which corresponds to non-moral.

The entire scheme is:

- *Care*: HN
- *Harm*: HP
- *Fairness*: FP
- *Cheating*: FN
- *Loyalty*: LP
- *Betrayal*: BN
- *Authority*: AP
- *Subversion*: AN
- *Purity*: PP
- *Degradation*: PN
- *Non-moral*: NM

The documents that you will code will be displayed in a spreadsheet. Each document will be contained in a row. Your job is to add the above described labels to the column called 'mft_sentiment'. See 1 for an example.

When coding a document, try to follow this order of operations:

1. Does the document (or sentence/phrase) seem to have moral content? If no, enter 'NM' in the 'mft_sentiment' column. If yes, continue to 2.
2. Which foundation does the document seem most associated with? Label the document with this code in the '1mft_sentiment' column proceed to 3.
3. Does the document seem to be associated with any other foundations? If no, proceed to 5. If yes, add these to the 'mft_sentiment' column
4. Is there anything important to note about the document? Is it in a different language? Does it seem in any way like it should be excluded (e.g. because it is fake, because it has been repeated multiple times, etc.) If so, add a note describing these issues in the 'note' column.

While the language that is used to signal moral relevance can be highly variable and abstract, it can be useful to have a core set of terms or concepts to use as anchors for identifying the presence or absence of a signal. Below such terms are listed in association with each foundation pole:

Harm:

Positive (HP): save, defend, protect

Negative (HN): harm, war, kill

Fairness:

Positive (FP): fair, equal, justice, rights, honesty

Negative (FN): unfair, unequal, unjust, dishonest, cheating

Loyalty:

Positive (LP): together, nation, family, solidarity, friends, support

Negative (LN): betray (in-group) , abandon, enemy, against, rebel (against group)

Authority:

Positive (AP): duty, law, honor, obligation, job, order

Negative (AN): rebelling (against authority), chaos, disorder, betray (your role)

Purity:

Positive (PP): clean, sacred, preserve, pure, serenity, noble

Negative (PN): dirty, repulsive, disgusting, revolting, animalistic

In addition to coding documents, please also keep a log of difficult cases that you encounter. We will discuss these with the other coders and, when it is appropriate, we will try to resolve these cases.

For example, if prayer is mentioned in a document, this could be relevant to purity, however if someone is praying for another person's wellbeing, it might also be relevant to care (HP). While it is certainly acceptable to code both as present, some degree of conservatism should be maintained because we do not want to inflate or suppress the frequency of a given moral sentiment.

References

- Dehghani, M., Johnson, K., Hoover, J., Sagi, E., Garten, J., Parmar, N. J., . . . Graham, J. (2016, 4 January). Purity homophily in social networks. *Journal of experimental psychology. General*.
- Graham, J., Haidt, J., & Nosek, B. a. (2009). Liberals and conservatives rely on different sets of moral foundations. *Journal of personality and social psychology*, 96(5), 1029–1046.
- moralfoundations.org*. (n.d.). <http://moralfoundations.org/>.
- Sagi, E., & Dehghani, M. (2014, 1 April). Measuring moral rhetoric in text. *Social science computer review*, 32(2), 132–144.

Table 1

Spreadsheet Example

ID	text	mft_sentiment	notes
123	Cause I want to be anarchy, Its the only way to be	AN	