Bodily Maps of Moral Concerns

Mohammad Atari

Department of Psychology, Brain and Creativity Institute,
University of Southern California

Aida Mostafazadeh Davani

Department of Computer Science, Brain and Creativity Institute,
University of Southern California

Morteza Dehghani

Department of Psychology, Department of Computer Science,
Brain and Creativity Institute, University of Southern California

Abstract

The somatosensory reaction to different social circumstances has been proposed to trigger conscious emotional experiences. Here, we present a pre-registered experiment in which we examine the topographical maps associated with violations of different moral concerns. Specifically, participants ($N = 596$) were randomly assigned to scenarios of moral violations, and then drew their subjective somatosensory experience on two 48,954-pixel silhouettes. We demonstrate that bodily representations of different moral violations are slightly different. Further, we demonstrate that violations of moral concerns are felt in different parts of the body, and arguably result in different somatosensory experiences for liberals and conservatives. We also investigate how individual differences in moral concerns relate to bodily maps of moral violations. Finally, we use natural language processing to predict activation in body parts based on the semantic representation of textual stimuli. The findings shed light on the complex relationships between moral violations and somatosensory experiences.

Keywords: emotion, feeling, morality, Moral Foundations Theory, natural language processing

Bodily Maps of Moral Concerns

Whether moral judgment is a product of reason or emotion has been an ongoing debate among philosophers and psychologists for decades. When moral psychology separated itself from moral philosophy, it almost exclusively focused on moral reasoning, rather than affective aspects of morality. The first empirical attempts in moral psychology started by examining cognitive-developmental components of understanding fairness and rules (Kohlberg, 1971; Piaget, 1948). Ever since, the field has been constantly growing and there has been an increasing interest in examining affective components of morality. Accumulating evidence suggests that emotion can ensue from, amplify, or directly cause moral judgement (Avramova & Inbar, 2013). Irrespective of the exact nature of the relationship between emotion and morality, distinct emotions are known to be associated with specific moral concerns as well as moral violations (Haidt, 2003).

Emotions are neural and somatic events whose evolutionary function is to prepare an organism to respond adaptively to a change in social or physical circumstances (Darwin, 1872). Once emotions are induced, individuals can consciously experience their emotions by mentally constructing a *feeling* (Damasio, 1999). Constructing feelings of emotions recruits brain systems that regulate and map body responses (Damasio & Carvalho, 2013). Both classic and modern theories of emotion postulate that interoception (i.e., sensing of physiological feedback from the body and its visceral organs) is essential for emotional experience (Damasio, 1999; James, 1994; Schachter & Singer, 1962). The link between interoception and emotion has been supported by several studies. For example, Barrett, Quigley, Bliss-Moreau, and Aronson (2004) found that arousal focus, the extent to which individuals emphasize the changes of feelings in their verbal reports of experienced emotion, is related to interoceptive sensitivity. Individuals who were sensitive to their heartbeat change due to emotion arousal images reported more intense emotional experiences compared to the less sensitive individuals (Barrett et al., 2004), hence supporting the association between bodily feedback and emotional states.

Even though emotions have been studied for a long time in psychology, the topographical distribution of the emotion-related bodily sensations have only recently been identified (Nummenmaa, Glerean, Hari, & Hietanen, 2014). Nummenmaa and colleagues (2014) mapped the “feeling space” for different emotions and focused on certain feedback from the body when certain emotions are felt. These authors examined the association between bodily sensations and different emotions, finding that consciously felt emotions are represented in human body by topographically distinct maps with partial overlaps. Particularly, some emotions are associated with “activations” in certain body parts, while other bodily regions might be “deactivated” in the same emotional experience. For example, feelings of fear are paired with activations in the chest and head area, whereas feelings of sadness are represented by exorbitant deactivations in lower limbs and slight activations in the chest.

To our knowledge, topographical representations of moral emotions ensuing moral violations have not yet been studied. As mentioned, distinct emotions are known to be associated with various violations of moral norms (Haidt, 2003). Two decades ago, Rozin, Lowery, Imada, and Haidt (1999) proposed the Contempt, Anger, and Disgust (CAD) hypothesis indicating that the “other-condemning” moral emotions of contempt, anger, and disgust correspond to violations of three moral codes of Community, Autonomy, and Divinity. Consistent with the CAD hypothesis and based on the intuitionist perspective on moral judgment, the Moral Foundations Theory (MFT; Graham et al., 2013; Haidt & Joseph, 2004) was developed by searching for the best links between anthropological and evolutionary accounts of moral intuitions across cultures. This framework suggests that moral intuitions derive from innate psychological mechanisms that coevolved with cultural institutions. Each moral system produces fast, automatic gut-level reactions of like or dislike when certain phenomena are perceived in the social world, which in turn guide moral judgments of right and wrong. These systems, according to the MFT, have evolutionary adaptive underpinnings present in individuals across cultural norms: Care,

Fairness, Loyalty, Authority, and Purity (Graham et al., 2013).

Violating the norms of Care, Fairness, Loyalty, Authority, and Purity does not necessarily produce a uniform emotional response. Research suggests specific predictions regarding the types of response that violations of distinct moral norms may elicit. For instance, the specific correspondence model (Chapman & Anderson, 2013; Kemper & Newheiser, 2018) posits that specific moral emotions map onto specific violation types. Therefore, any action in which an entity was harmed should reliably elicit anger and any action in which a bodily norm was violated should elicit disgust (Rozin & Haidt, 2013; Russell & Giner-Sorolla, 2013). The specific correspondence model, thus, suggests that witnessing harmful actions should uniquely activate a desire to confront the violator, whereas witnessing impure actions should uniquely elicit avoidance. By contrast, constructionist models of moral emotion argue that there are no exclusive links between moral content domains and elicited emotions (Cameron, Lindquist, & Gray, 2015). Rather, contextual cues (e.g., framing language) and conceptual knowledge (e.g., who was harmed; Gray & Schein, 2012) inform the interpretation of moral violation, and no specific stimulus would predictably elicit the same emotion across all moral contexts (Cameron et al., 2015). Recent research suggests that individuals express distinctively high levels of desire to avoid (vs. confront) violators of purity norms. Violations of other moral norms, however, do not similarly elicit unique patterns of avoidance or confrontation. Thus, behavioral responses to moral violations depend in part on the norm that was violated, with impure acts eliciting a uniquely strong avoidance response (Kemper & Newheiser, 2018). Therefore, it stands to reason that topography of different moral violations would reveal different maps on the human body (see Nummenmaa, Hari, Hietanen, & Glerean, 2018).

Here, we aim to examine how emotions associated with violations of moral concerns are topographically represented in bodily sensations. Specifically, we address four pre-registered questions: (1) are bodily sensation maps distinct for the five moral concerns theorized by MFT? (2) can machine learning techniques be used to distinguish bodily

sensation maps associated with different moral violations? (3) can we differentiate the bodily sensation maps resulting from each moral violation for different political ideologies (e.g., do liberals feel “purity” in the same parts of their bodies as conservatives) and individual differences in moral concerns (e.g., do people who score high on purity concerns feel “purity” in the same parts of their body as those who score low on purity concerns)? (4) can the topographical bodily sensation maps be predicted from the textual descriptions of moral violation vignettes?

Methods

Participants

As mentioned in our pre-registration and based on Nummenmaa et al. (2014), we aimed to recruit 600 participants. Overall, 630 participants on Amazon’s Mechanical Turk (MTurk; see Buhrmester, Kwang, & Gosling, 2011; Litman, Robinson, & Abberbock, 2017) participated in this study. Consistent with our pre-registration, we excluded those participants who failed attention checks, resulting in a total of 596 individuals (age: $M = 36.5$ years, $SD = 11.8$ years; gender: *female* = 355, *male* = 237, *other* = 3, *unknown* = 1). Participants were randomly assigned to one of the five experimental conditions and completed the target task and several individual difference measures, explained below.

Measures

Moral Violation Scenarios. Participants were randomly assigned to one of five moral violation conditions based on the MFT, where they each read a vignette about violation of a particular moral foundation (Clifford, Iyengar, Cabeza, & Sinnott-Armstrong, 2015). For each foundation, we selected vignettes from a larger pool of stimuli provided by Clifford et al. (2015). For each foundation (Care, Fairness, Loyalty, Authority, and Purity), we included four vignettes, matched on average perceived wrongness, arousal, and frequency. A sample vignette is “You see a woman clearly avoiding

sitting next to an obese woman on the bus.” After presenting the scenario, participants were asked: “How morally wrong is the action depicted in this scenario?” (1 = *Not at all wrong*; 5 = *Very wrong*; $M = 3.30$, $SD = 1.18$) and “How strong was your emotional response to the behavior depicted in the scenario?” (1 = *Not at all strong*; 5 = *Very strong*; $M = 2.90$, $SD = 1.12$) on a 5-point Likert-type scale.

Bodily Sensation Task. Bodily topography of feelings was mapped using the emBODY tool (Nummenmaa et al., 2014). The online data acquisition package is publicly available at <https://version.aalto.fi/gitlab/eglerean/embody>. Yet, some parts of our analytic framework is different from that of Nummenmaa et al. (2014). Before engaging in the task, participants were shown a brief tutorial video to make sure they fully understood the task. Participants were then asked to color whereabouts of activations and deactivations felt in their body. Specifically, participants viewed two silhouettes placed on the right and left sides of the screen. After reading the moral violation scenario, participants were asked to draw on both silhouettes the portions of the body where they felt *activation* (left silhouette) and *deactivation* (right silhouette). Individual-level bodily topographies were then computed based on the difference between the left and right silhouettes. Final bodily sensation maps were represented by 48,954 pixels.

Political Orientation. All participants rated their affiliation with the republican political party as opposed to the democratic political party along a 7-point scale ranging from 1 (*Strong Democrat*) to 7 (*Strong Republican*). Another item asked participants to rate their political conservatism on a scale ranging from 1 (*Very Liberal*) to 7 (*Very Conservative*). We averaged these two items in order to achieve a political orientation score where higher scores indicate more conservative political orientation. Previous work has used a similar method for assessment of political ideology (Jost & Thompson, 2000). The internal consistency of these two items was relatively high in the current sample ($\alpha = .94$). We labeled those who scored one standard deviation ($SD = 1.86$) lower than the average political conservatism ($M = 3.83$) as liberal ($N = 91$), and those who scored one standard

deviation higher as conservative ($N = 126$).

Moral Concerns. Individual differences in moral concerns were measured using the 30-item Moral Foundations Questionnaire (MFQ; Graham et al., 2011). MFQ assesses the degree to which participants deem different considerations as relevant (1 = *Not at all relevant*; 5 = *Extremely relevant*) when making moral judgments and their agreement (1 = *Strongly disagree*; 5 = *Strongly agree*) with statements germane to morality. These items were used to create foundation-level scores for Care ($\alpha = .69$), Fairness ($\alpha = .63$), Loyalty ($\alpha = .76$), Authority ($\alpha = .76$), and Purity ($\alpha = .83$). Those who scored at least one standard deviation higher than the average of the respective MFQ subscale were considered “high” on that moral concern, and those who scored at least one standard deviation lower were considered “low” on that moral concern.

Procedure

This study was approved by the Institutional Review Board (IRB) at University of Southern California (UP-16-00695-AM003). Potential participants were invited to take part in a psychological study on Amazon’s Mechanical Turk (MTurk) for monetary compensation. Participation was on a voluntary basis and participants were compensated \$0.50 for their participation. Each participant completed the bodily sensation task after having read a vignette about a particular moral violation and then completed a set of self-administered measures of political ideology and moral concerns. This study’s hypotheses, predictions, and analyses were pre-registered on Open Science Framework (OSF; <https://osf.io/zbv6e/>) and all data and materials will be made publicly available upon acceptance of the article.

Analytic Strategy

As discussed in our pre-registration, we examined the topographic representation of each moral scenario across conditions, using the method described in Nummenmaa et al. (2014). Specifically, for each participant, a single map comprising 48,954 pixels was

obtained by subtracting the activation and deactivation maps to obtain 48,954 variables per person in a given condition. The differences between activations and deactivations were then assessed using 48,954 univariate t -tests: a one-sample t -test against zero was performed for each pixel within a condition, resulting in a statistical t -map. Then we produced effect-size-maps based on the t -maps and sample sizes in each condition.

In order to classify the condition each participant was assigned to, we used Support Vector Machines (SVM; Cortes & Vapnik, 1995; Hearst, Dumais, Osuna, Platt, & Scholkopf, 1998) which are a class of supervised machine-learning algorithms considered to be robust-to-overfitting in classification settings. In the training phase of SVM, a hyperplane is chosen which maximizes the margin of separation between the classes, while also allowing for some data points to be misclassified. Based on the model built from training data, a new data point can be classified based on its position relative to the hyperplane. Further, non-linear classification can be performed using SVMs by projecting the data in to high-dimensional space using the so called *kernel trick*. In all of our analyses, we balanced the datapoints between the classes and performed 10-fold cross-validations 100 times to obtain estimates of variance in the performance of the models. We, then, used Fisher’s one-sample permutation test, with 5,000 permutations, to compare the average accuracy against chance to assess whether the classification algorithm performed significantly above the chance level. Averaged accuracies (across the 100 folds) and p -values associated with permutation tests are reported for each model.

To examine the role of political orientation in bodily emotional sensations, we divided the data by political orientation and investigated whether the maps in each condition are significantly different for liberals and conservatives (e.g., do liberals feel ‘purity’ in the same part of their body as conservatives do?). Similar to above, for each condition, we balanced the datapoints between the classes (i.e., liberals and conservatives) and, for each concern, performed 10-fold cross-validations 100 times to obtain estimates of variance in the performance of the models. We then compared the average accuracy against chance

(50%) using Fisher’s one-sample permutation test, with 5,000 permutations to investigate the reliability of the differences between the maps generated by liberals and conservatives.

We also broke down the MFQ responses to examine whether being high (vs. low) on a particular moral concern has a similar effect on bodily sensations. For each condition, we balanced the datapoints between the classes and performed 10-fold cross-validations 100 times to obtain estimates of variance in the performance of the models to examine whether high-scorers and low-scorers feel moral violations in different body regions. We then compared the average accuracy in each condition against chance (50%) using Fisher’s one-sample permutation test, with 5,000 permutations.

As a pre-registered exploratory analysis, we explored whether the topographical maps for different moral concerns can be predicted from the textual description of that moral violation as stated in the vignette in natural language. Whereas typical approaches for the analysis of moral language (e.g., Garten et al., 2018) estimate the “moral loadings” of a piece of text based on the moral words (as measured by a pre-specified dictionary), here the vignettes describe events which are potentially, but not necessarily, moral (Clifford et al., 2015). As such, we applied *InferSent*¹ (Conneau, Kiela, Schwenk, Barrault, & Bordes, 2017) to generate vector representations of each vignette. InferSent learns sentence-level representations by training a bi-directional Recurrent Neural Network (RNN) with Long Short-Term Memory (LSTM) on the Stanford Natural Language Inference dataset (Bowman, Angeli, Potts, & Manning, 2015). Using the best performing predictive model, we encoded each moral vignette into a vector of length 4,096. We then ran 48,954 ridge regressions using InferSent vectors as predictors of activation of each pixel, and then calculated R^2 for each model to build a R^2 -map to visualize which parts of the body’s (de)activation can be better explained by the semantic representation of textual stimuli.

We ran all analyses in R (version 3.4.1), Python (version 3.6) and Octave

¹ <https://github.com/facebookresearch/InferSent>

(version 4.4.1) programming languages. We used `Caret` (version 6.0)², `e1071` (version 1.7)³, `Psych` (version 1.8)⁴, and `EnvStats` (version 2.3)⁵ packages in R.

Results

Bodily sensation of moral violations

Figure 1 (top) displays the bodily sensation maps associated with each moral violation. For each condition, we normalized the difference between activation pixels and deactivation pixels subject-wise. Then, we performed 48,954 one-sample t -test against zero to get the resulting t -maps for each moral violation condition. Next, we transformed these t -maps into effect-size-maps by dividing each pixel's t value by the square root of sample size in that condition (Cohen's d). Effect-size-maps are visualized in Figure 1. As can be seen, in each condition, the moral violations are associated with slightly distinct distributions of body areas where activation was felt to either increase (or speed up) or decrease (or slow down).

Classification of moral violations

As mentioned previously, we ran binary SVM classifiers with 10-fold cross-validation 100 times for each condition where in each fold the target concern, and a randomly selected subset of the other concerns (equal in size) were classified. For Care, the average classification accuracy was 49.0%, which was not significantly higher than chance ($p = .991$). For Fairness, the average classification accuracy was 51.5%, which was significantly greater than chance ($p = .002$). For Loyalty, the average accuracy was 49.0%, which was not significantly higher than chance ($p = .983$). For Authority, the average accuracy was

² <https://cran.r-project.org/web/packages/caret/index.html>

³ <https://cran.r-project.org/web/packages/e1071/index.html>

⁴ <https://cran.r-project.org/web/packages/psych/index.html>

⁵ <https://cran.r-project.org/web/packages/EnvStats/index.html>

Figure 1. Bodily sensation maps of moral violations. Top: The bodily sensation maps show regions where activation increased (warm colors) or decreased (cool colors) in each moral violation condition. Bottom: The bodily sensation maps for liberals and conservatives across moral concerns. The color bar indicates the effect sizes.

49.0%, which was not significantly higher than chance ($p = .993$). Finally, the mean classification accuracy for Purity was 51.0%, which was significantly higher than chance ($p = .009$).

Individual differences and moral violations

We ran binary SVM classifiers with 10-fold cross-validation 100 times for each condition where in each fold the political orientation of the left-out participants was predicted. The average classification accuracies for Care ($n/\text{group} = 45$, $M = 66.7\%$, $p < .001$), Fairness ($n/\text{group} = 36$, $M = 62.3\%$, $p < .001$), Loyalty ($n/\text{group} = 49$, $M = 74.0\%$, $p < .001$), Authority ($n/\text{group} = 48$, $M = 66.7\%$, $p < .001$), and Purity ($n/\text{group} = 39$, $M = 72.0\%$, $p < .001$) were significantly greater than chance. Therefore, for all moral

concerns we could classify political ideology based on bodily sensation maps, indicating that liberals and conservatives *feel* moral violations, especially feelings of Loyalty and Purity, in different parts of their body. Bodily sensation maps of moral violations for liberals and conservatives are presented in Figure 1 (bottom). Although we pre-registered our analysis, these findings should be interpreted with caution because the number of liberals and conservatives included is small ($37 \leq n \leq 50$ in each group), and middle-of-the-road individuals were not included.

We also examined classification of high (vs. low) moral concerns based on bodily sensation maps. We ran binary SVM classifiers with 10-fold cross-validation 100 times for each condition where in each fold high-scorers and low-scorers on MFQ subscales were classified. The average classification accuracies for Care ($n/\text{group} = 39$, $M = 66.3\%$, $p < .001$), Fairness ($n/\text{group} = 37$, $M = 61.0\%$, $p < .001$), Loyalty ($n/\text{group} = 37$, $M = 50.5\%$, $p = .504$), Authority ($n/\text{group} = 42$, $M = 49.8\%$, $p = .999$), and Purity ($n/\text{group} = 43$, $M = 61.3\%$, $p < .001$) were relatively high except for Loyalty and Authority. Therefore, individual differences in moral concerns are associated with where in body individuals *feel* moral violations. Bodily sensation maps of moral violations for low-scorers and high-scorers are presented in the Supplementary Materials. Of note, high-scorers and low-scorers formed small groups ($37 \leq n \leq 43$ in each group) and these findings should be treated with caution until further replicated.

Semantic representation of moral violations

As mentioned in our analytic strategy, we represented each vignette onto a 4096-dimensional sentence embedding. We used these vectors to predict (de)activation of each pixel using a ridge regression with 10-fold cross-validation. Then we computed a R^2 -map, indicating the explained variance in each pixel based on the semantic representation of the textual stimuli. Results of the these ridge regressions are presented in Figure 2. It can be seen that semantic representation of the texts used as experimental

stimuli can predict larger proportions of variance in activation of the gut area.

Figure 2. The percentage of variance explained in the activation of different body parts based on the semantic representation of textual stimuli. The color bar indicates the explained variance.

Discussion

The social intuitionist model of moral cognition (Haidt, 2001; Haidt & Bjorklund, 2008) suggests that moral judgments are caused by emotional responses to a person, action, or violation. Drawing upon intuitionist models of human morality (Haidt, 2001; Haidt & Joseph, 2004) and recent research on bodily maps of emotions (Hietanen, Glerean, Hari, & Nummenmaa, 2016; Nummenmaa et al., 2014, 2018), we conducted a pre-registered examination of bodily sensations associated violations of different moral concerns.

We demonstrated that moral violation scenarios are associated with changes in activation and deactivation of specific body regions comparable with subjective feelings (Nummenmaa et al., 2018). The topographic maps associated with moral violations manifested more commonalities than differences (see Figure 1). These similarities are consistent with Kemper and Newheiser (2018) who did not find evidence for a specific correspondence mapping of behavioral response to moral violations based on MFT. It can be seen that across moral violations, the head/face area is highly activated, paired with varying levels of activation in the chest. The consistent activation observed in the head

area suggests that people subjectively associate moral violations with high-level cognitive processing. This is not surprising as the moral violation scenarios require a high level of social cognitive processing as well as an evaluation of personal states in relation to standards of right and wrong. The patterns of activation and deactivation observed for the head area consistently occurred across moral violations, and they may reflect subjective changes not only in the brain, but also in the face. Based on the current data we cannot infer whether these activations reflect perceived mental processing, facial blushing, or a combination of both. Deactivation of the limbs represented a consistent pattern across all moral scenarios. Of note, the chest area is activated in Care, Fairness, Loyalty, and Authority, but less so in Purity. Instead, people reported higher activations in the abdomen area in Purity.

We trained classifiers that could reliably predict political ideology based on bodily maps. Over the past decade, a body of research has shown that liberals and conservatives rely on different moral foundations and react differently to different moral violations (for a review see Graham et al., 2013). This is the first work indicating that political orientation influences where and how moral violations are felt in the body. These findings contrast null effects of political orientation on the link between moral transgressions and emotion using self-report measures (Landmann & Hess, 2018). Therefore, it is possible that, as we show, liberals and conservatives feel moral violations in different body regions, interpret them as distinct complex feelings (Landmann & Hess, 2018), and subsequently make different moral and political judgements.

We also trained classifiers that were able to reliably predict people's moral concerns (high vs. low in moral foundations). It stands to reason that those who score high on, for example, Purity may have different reactions to impure actions compared to their counterparts who do not endorse Purity values as strongly (e.g., Heerdink, Koning, Van Doorn, & Van Kleef, 2018). Those who are more concerned with Purity are more sensitive to cues of degradation and Purity violations; thus such individuals are more likely

to report stronger subjective activations and/or deactivations when primed with stimuli that violate relevant norms. As mentioned, these results are promising, but should be treated with caution until replicated in future studies.

Semantic representation of textual stimuli (Clifford et al., 2015) predicted (de)activation of different body parts, especially in the abdomen/gut area. Research in cognitive linguistics and social cognition suggests that individuals construe the world in large part through conceptual metaphors, which enable them to understand abstract concepts using knowledge of superficially dissimilar, but more concrete phenomena (Lakoff, 2016). Interestingly, the vignettes that we used did not include highly moral words, rather they simply described a social scenario in which a particular moral norm was violated. We found that the semantic representations can locate bodily representation of moral violations, suggesting that the semantic space of natural language describing moral violations can be coupled with emotional states and their bodily sensation maps.

This study has limitations worth noting. First, we collected self-report data regarding where activations and deactivations are felt in the body because our main objective was to examine representation of moral transgressions in the body, so we did not collect data with regard to change in physiological states. Second, we collected data on two 2D silhouettes to represent activation and deactivation of different body parts. Future research can use more accurate silhouettes, including 3D ones to better disentangle different body parts (e.g., occipital and frontal parts of the head). Third, we used vignettes that were matched on frequency and wrongness, only representing the moral foundations suggested by MFT. A good next step is to include other morally controversial scenarios (e.g., transgressions from norms of honesty or humility) with varying levels of wrongness, frequency, and weirdness.

Conclusion

Morality is a central feature of human life. Moral violations occur frequently across time and cultures, and humans are evolutionarily equipped to *feel* these transgressions and

act accordingly. Here, we show that violations of standards of Care, Fairness, Loyalty, Authority, and Purity are associated with slightly different bodily sensation maps. Further we tested how these maps differ based on political ideology and individual differences in moral concerns, finding that liberals and conservatives feel moral violations in different parts of their body. Finally, we show how semantic representation of textual stimuli can predict bodily sensations of moral violations in specific body parts, particularly in the abdomen/gut area. These results extend the current models of the role of somatosensation and embodiment in moral judgment. Topographical representation of moral violations in the body could, therefore, provide a somatosensory biomarker for moral intuitions.

References

- Avramova, Y. R., & Inbar, Y. (2013). Emotion and moral judgment. *Wiley Interdisciplinary Reviews: Cognitive Science*, 4, 169–178. doi: 10.1002/wcs.1216
- Barrett, L. F., Quigley, K. S., Bliss-Moreau, E., & Aronson, K. R. (2004). Interoceptive sensitivity and self-reports of emotional experience. *Journal of Personality and Social Psychology*, 87, 684–697. doi: 10.1037/0022-3514.87.5.684
- Bowman, S. R., Angeli, G., Potts, C., & Manning, C. D. (2015). A large annotated corpus for learning natural language inference. *arXiv preprint arXiv:1508.05326*.
- Buhrmester, M., Kwang, T., & Gosling, S. D. (2011). Amazon’s mechanical turk: A new source of inexpensive, yet high-quality, data? *Perspectives on Psychological Science*, 6, 3–5. doi: 10.1037/e527772014-223
- Cameron, C. D., Lindquist, K. A., & Gray, K. (2015). A constructionist review of morality and emotions: No evidence for specific links between moral content and discrete emotions. *Personality and Social Psychology Review*, 19, 371–394. doi: 10.1177/1088868314566683
- Chapman, H. A., & Anderson, A. K. (2013). Things rank and gross in nature: A review and synthesis of moral disgust. *Psychological Bulletin*, 139, 300–327. doi: 10.1037/a0030964
- Clifford, S., Iyengar, V., Cabeza, R., & Sinnott-Armstrong, W. (2015). Moral foundations vignettes: A standardized stimulus database of scenarios based on moral foundations theory. *Behavior Research Methods*, 47, 1178–1198. doi: 10.3758/s13428-014-0551-2
- Conneau, A., Kiela, D., Schwenk, H., Barrault, L., & Bordes, A. (2017). Supervised learning of universal sentence representations from natural language inference data. *arXiv preprint arXiv:1705.02364*.
- Cortes, C., & Vapnik, V. (1995). Support-vector networks. *Machine Learning*, 20, 273–297. doi: 10.1007/bf00994018
- Damasio, A. (1999). *The feeling of what happens: Body and emotion in the making of*

- consciousness*. New York, NY: The New York Times Book Review.
- Damasio, A., & Carvalho, G. B. (2013). The nature of feelings: Evolutionary and neurobiological origins. *Nature Reviews Neuroscience*, *14*, 143–153. doi: 10.1038/nrn3403
- Darwin, C. (1872). *The expression of the emotions in man and animals*. London: John Murray.
- Garten, J., Hoover, J., Johnson, K. M., Boghrati, R., Iskiwitch, C., & Dehghani, M. (2018). Dictionaries and distributions: Combining expert knowledge and large scale textual data content analysis. *Behavior Research Methods*, *50*, 344–361. doi: 10.3758/s13428-017-0875-9
- Graham, J., Haidt, J., Koleva, S., Motyl, M., Iyer, R., Wojcik, S. P., & Ditto, P. H. (2013). Moral foundations theory: The pragmatic validity of moral pluralism. *Advances in Experimental Social Psychology*, *47*, 55–130. doi: 10.1016/B978-0-12-407236-7.00002-4
- Graham, J., Nosek, B. A., Haidt, J., Iyer, R., Koleva, S., & Ditto, P. H. (2011). Mapping the moral domain. *Journal of Personality and Social Psychology*, *101*, 366–385. doi: 10.1037/a0021847
- Gray, K., & Schein, C. (2012). Two minds vs. two philosophies: Mind perception defines morality and dissolves the debate between deontology and utilitarianism. *Review of Philosophy and Psychology*, *3*, 405–423. doi: 10.1007/s13164-012-0112-5
- Haidt, J. (2001). The emotional dog and its rational tail: A social intuitionist approach to moral judgment. *Psychological Review*, *108*, 814–834. doi: 10.1037/0033-295x.108.4.814
- Haidt, J. (2003). The moral emotions. In R. J. Davidson, K. R. Scherer, & H. H. Goldsmith (Eds.), *Handbook of affective sciences* (pp. 852–870). Oxford: Oxford University Press.
- Haidt, J., & Bjorklund, F. (2008). Social intuitionists answer six questions about moral

- psychology. In W. Sinnott-Armstrong (Ed.), *Moral psychology, vol. 2. the cognitive science of morality: Intuition and diversity* (pp. 181–217). Cambridge, MA, US: MIT Press.
- Haidt, J., & Joseph, C. (2004). Intuitive ethics: How innately prepared intuitions generate culturally variable virtues. *Daedalus*, *133*, 55–66. doi: 10.1162/0011526042365555
- Hearst, M. A., Dumais, S. T., Osuna, E., Platt, J., & Scholkopf, B. (1998). Support vector machines. *IEEE Intelligent Systems and Their Applications*, *13*, 18–28. doi: 10.1109/5254.708428
- Heerdink, M. W., Koning, L. F., Van Doorn, E. A., & Van Kleef, G. A. (2018). Emotions as guardians of group norms: Expressions of anger and disgust drive inferences about autonomy and purity violations. *Cognition and Emotion*, 1–16. doi: 10.1080/02699931.2018.1476324
- Hietanen, J. K., Glerean, E., Hari, R., & Nummenmaa, L. (2016). Bodily maps of emotions across child development. *Developmental Science*, *19*, 1111–1118. doi: 10.1111/desc.12389
- James, W. (1994). The physical basis of emotion. *Psychological Review*, *101*, 205–210. doi: 10.1037/0033-295X.101.2.205
- Jost, J. T., & Thompson, E. P. (2000). Group-based dominance and opposition to equality as independent predictors of self-esteem, ethnocentrism, and social policy attitudes among African Americans and European Americans. *Journal of Experimental Social Psychology*, *36*, 209–232. doi: 10.1006/jesp.1999.1403
- Kemper, N. S., & Newheiser, A.-K. (2018). To confront or to avoid: How do people respond to violations of moral norms? *Social Psychological and Personality Science*, *9*, 734–743. doi: 10.1177/1948550617722831
- Kohlberg, L. (1971). Stages of moral development. *Moral Education*, *1*, 23–92.
- Lakoff, G. (2016). Language and emotion. *Emotion Review*, *8*, 269–273. doi: 10.1177/1754073915595097

- Landmann, H., & Hess, U. (2018). Testing moral foundation theory: Are specific moral emotions elicited by specific moral transgressions? *Journal of Moral Education*, *47*, 34–47. doi: 10.1080/03057240.2017.1350569
- Litman, L., Robinson, J., & Abberbock, T. (2017). Turkprime.com: A versatile crowdsourcing data acquisition platform for the behavioral sciences. *Behavior Research Methods*, *49*, 433–442. doi: 10.3758/s13428-016-0727-z
- Nummenmaa, L., Glerean, E., Hari, R., & Hietanen, J. K. (2014). Bodily maps of emotions. *Proceedings of the National Academy of Sciences*, *111*, 646–651. doi: 10.1073/pnas.1321664111
- Nummenmaa, L., Hari, R., Hietanen, J. K., & Glerean, E. (2018). Maps of subjective feelings. *Proceedings of the National Academy of Sciences*, *115*, 9198–9203. doi: 10.1073/pnas.1807390115
- Piaget, J. (1948). *The moral judgment of the child*. New York, NY, US: Free Press.
- Rozin, P., & Haidt, J. (2013). The domains of disgust and their origins: Contrasting biological and cultural evolutionary accounts. *Trends in Cognitive Sciences*, *17*, 367–368. doi: 10.1016/j.tics.2013.06.001
- Rozin, P., Lowery, L., Imada, S., & Haidt, J. (1999). The CAD triad hypothesis: A mapping between three moral emotions (contempt, anger, disgust) and three moral codes (community, autonomy, divinity). *Journal of Personality and Social Psychology*, *76*, 574–586. doi: 10.1037/0022-3514.76.4.574
- Russell, P. S., & Giner-Sorolla, R. (2013). Bodily moral disgust: What it is, how it is different from anger, and why it is an unreasoned emotion. *Psychological Bulletin*, *139*, 328–351. doi: 10.1037/a0029319
- Schachter, S., & Singer, J. (1962). Cognitive, social, and physiological determinants of emotional state. *Psychological Review*, *69*, 379–399. doi: 10.1037/h0046234