Forthcoming in *The Atlas of Language Analysis in Psychology* (Guilford Press)

Theory-driven Measurement of Emotion (Expressions) in Social Media Text

William J. Brady^{1*}, Killian McLoughlin¹ & M.J. Crockett¹

¹Yale University, Department of Psychology

Word count: 6,380

Theory-Driven Measurement of Emotion (Expressions) in Social Media Text

Sentiment analysis, or the measurement of opinions, attitudes and emotions in text, has exploded in the past 10 years. The number of academic articles published that mention “sentiment analysis” grew from 78 in 2008 to 1,876 in 2018 (Web of Science). This growth brought notable advances in sentiment classification whether it be measuring emotional valence (i.e., negative vs. positive sentiment) or the measurement of discrete, specific emotions (e.g., “anger”). In particular, there has been recent interest in measuring emotion expressions in social media text, as there are now over 3 billion users expressing emotions online on a regular basis. Although there is now a wealth of emotion expression data that exists on social media, only recently have psychological scientists begun to make use of it for the purposes of observational research and psychological hypothesis testing. In this chapter, we aim to provide psychologists with a practical guide for making use of advances in sentiment analysis for the purpose of measuring specific emotions in social media text, and ultimately testing hypotheses. We begin with an examination of the type of emotion-related information researchers can expect to measure in social media text. Next, we present a practical guide for the theory-driven application of supervised machine learning to measure specific emotion expressions on social media using the measurement of moral outrage expression as a case study. We end with considerations for researchers who wish to apply sentiment analysis of social media text for the purposes of psychological hypothesis testing.

What does sentiment analysis applied to social media measure?

Consider the following social media message:

"I am so. freaking. done. with our inept government!"

Intuitively, we might judge from this text that the message author is experiencing a high-arousal, negatively-valenced emotion such as moral outrage. To make this judgment, we rely on a shared cultural understanding of the concept of "moral outrage" to judge the message as such. For instance, American social media users understand that the message contains elements associated with the concept of 'moral outrage', or a high-arousal, negative emotional state that is elicited by a morally-relevant transgression (e.g., a violation of a prescriptive norm understood as beneficial for the good of society; Haidt, 2003; Salerno & Peter-Hagene, 2013; Tetlock, Kristel, Elson, Green, & Lerner, 2000).

More specifically, based on our concept of moral outrage, we rely on the detection of a combination of *linguistic features*: specific words ("freaking", "done", "inept), specific punctuation (repetitive use of periods, exclamation point), and the greater linguistic context (political opinion about the government). All told, our judgment that the message author is morally outraged depends on the identification of specific features that comprise the author's *emotion expression*.

Defining 'emotion expression' requires a theoretical understanding of what emotions are. Psychological constructionist theories of emotion are readily applicable to understanding emotions in the context of linguistic representation on social media, because they assume emotions are inherently based on conceptual knowledge (Barrett, 2013; Russell & Barrett, 1999). Specifically, emotions are defined as categorizations of

underlying feelings of affect and arousal ('core affect') with the use of conceptual knowledge. Following this definition of emotion, in the context of social media where emotions are expressed in text and images, emotion expression refers to the employment of culturally-defined concepts by the expresser that symbolically represent underlying feelings of valence and arousal, or 'affect' (see, e.g. Lindquist, MacCormack, & Shablack, 2015). By identifying specific conceptual and linguistic features that fall out of a cultural concept of an emotion, we can detect specific emotion expressions in social media text. In other words, the message author uses specific linguistic features to represent a culturally-defined concept of moral outrage, and as American social media users we perceive the expression and form a judgment about the expressed emotion using our shared concept of moral outrage.

The example above underscores how sentiment analysis measures the linguistic representation of emotion (emotion *expression*). It is worth noting a distinction between emotion expression in text and an underlying emotion *experience* of the message author at their computer or smartphone where they composed the message. For example, we may be tempted to make an additional inference based on the expression of moral outrage: that the message author is experiencing morally outraged. The validity of this inference is not something that can be directly measured by sentiment analysis of social media text. If the researcher wishes to measure emotion experience or infer it from expression, further steps would be required to validate the inference. We discuss this issue in the context of psychological hypothesis testing in the last section of this chapter.

In summary, to measure specific emotion expressions on social media, we need to identify a specific social context (e.g., American culture), which comes with a shared understanding of concepts that represent various emotions that are recognized in the social context. We then identify various conceptual features that comprise the emotion concept, and corresponding linguistic features, or in other words the representation of the emotion (expression). **Fig. 1** outlines this process using the example of moral outrage.

Fig. 1. A conceptual work-flow for identifying specific emotions in social media text. By understanding that sentiment analysis on social media measures symbolic representations of contextually-defined concepts of emotion (emotion expressions), three steps become important to clarify in order to measure specific emotions precisely. First, identification of a specific social context such as a specific culture or social group. Second, targeting of an emotion concept that is a shared representation among individuals within the targeted social context. Third, identification of conceptual features that represent the specific emotion in the social context. Finally, identification of specific linguistic features that should appear in language associated with the shared concept of the targeted emotion.

Once a specific context, emotion concept, and features are identified, it is possible to precisely apply modern techniques in natural language processing and

machine learning to measure specific emotion expressions on social media. Below, we present a practical guide for psychologists focused on the application of supervised machine learning using the case study of measuring moral outrage expression on Twitter.

Applying supervised machine learning to measure emotion expressions

Supervised machine learning refers to the process by which a model learns a set of rules from labeled examples such that the algorithm can make predictions that assign labels to new examples (Kotsiantis, 2007). The model's learning is "supervised" because the labels are provided in the dataset for it to learn from. Here we focus on supervised *classification* models, or the prediction of a binary class (as opposed to regression models). For instance, a researcher might collect a dataset of tweets, hand-label them according to whether they think the tweets express moral outrage, and then provide the resulting binary labels to a model for it to learn.

In the case of measuring specific emotion expressions in social media text, there are several steps required to apply supervised machine learning. First, a "training dataset" from which the model will learn must be collected from social media. Second, the training dataset must be annotated in a theoretically-informed and consistent manner. Third, methods in natural language processing must be applied to process unstructured social media text into a format readable by machine learning methods. Fourth, a specific machine learning model must be selected and validated. Fifth, it is important to test how the model generalizes to novel datasets. Below, we consider each step in detail.

Collecting a training dataset of tweets

There are several considerations for choosing what social media data will comprise one's training dataset. The researcher at this point should have already chosen a specific social context (e.g., American Twitter users; see **Fig. 1**), so the key decision is how to focus data collection around user conversations that are likely to contain some signal of the emotion expression in question. This is important because supervised machine learning methods typically require the training dataset to have "class balance", i.e., having a roughly even number of positive and negative cases of the emotion expression. For instance, to measure moral outrage on Twitter, it is useful to collect tweets sent in the context of conversations about contentious topics within American politics (e.g., gun control, immigration, climate change), where moral outrage is likely to be expressed. On one hand, it is important to identify potential datasets where there is a high signal of the emotion expression in order to achieve class balance in the dataset. On the other hand, one also does not want to heavily overrepresent the base-rate of the emotion in the context of interest which could affect generalizability of the model (for a discussion of class balance in machine learning, see e.g., Batista, Prati, & Monard, 2004).

In order to collect Twitter data about specific topics, there are packages readily available for the *Python* (Tweepy, 2017) and *R* (rtweet; Kearney, 2018) programming languages that allow a researcher to mine Twitter's Standard API for specific keywords and #hashtags that can be used as an indication a user is tweeting about a topic. For example, collecting tweets about gun control could be achieved by using Twitter's API to scrape tweets that use the keyword "gun control" or for hashtags signaling specific

events that evoked national debates about gun control (e.g., a mass shooting event - #ParklandShooting).

It is important to recognize that collecting data with keywords and hashtags can also lead to “off-topic” data collection, such as a case where a popular music artist releases a song called “gun control” and the API returns tweets related to the song. The broader the keywords and hashtags used to search the API, the greater the chance of noisy data collection. For example, searching the API for keywords related to a specific event that evoked a conversation about gun control (e.g., #ParklandShooting) will likely yield more accurate data collection than searching for the keyword “guns”. To address the issue of noise in data collection, the researcher should make some attempts to estimate the validity of their data collection procedure to ensure the context of interest is being indeed being collected. There are several approaches to estimate data collection validity, such as manual coding or data-driven methods like topic modeling. As a manual coding approach, a researcher could take a small random sample of their data and manually code for whether each tweet is on topic. Repeating this procedure several times could form an estimate of the percentage of tweets that are likely to be off topic in the larger dataset. On the other hand, crowdsourcing could be employed to have every tweet in the dataset checked for its relevance to the target topic, where each tweet is rated by multiple participants on platforms such as Amazon’s *Mechanical Turk* or *Prolific.ac*. Regarding data-driven approaches, topic modeling methods such as Latent Dirichlet Allocation can be used to determine if the newly-collected data are comprised of similar “topics” compared to a known source of on-topic data (see Silge & Robinson,

2017, Chapter 6 and the 'methods' chapter of this book). The bottleneck for this approach is having a known on-topic dataset readily available.

When determining an adequate sample size to form a strong training dataset, there are both theoretical and practical issues. Theoretically, supervised machine learning methods produce more accurate predictions given a greater amount of training data. For instance, simpler models that perform worse than more complex models can achieve the same accuracy when the training dataset size is increased, but there is also a point at which increased training dataset size only improves accuracy of the model trivially (Banko & Brill, 2001). Practically, there are two potential bottlenecks for training dataset size. First, the availability of tweets might be constrained depending on how specific of a topic the researcher is interested in. If a researcher is interested in measure the moral emotion “awe” (Keltner & Haidt, 2003) in the context of tweets about a niche religious festival, it may be difficult to find more than a few hundred tweets. Furthermore, the greater the specificity in data collection, the greater the potential for low generalizability (see “testing model generalizability” below). Second, the researcher must consider what strategy they will use to annotate their training dataset ahead of time (see “annotating a training dataset” below). There may be limits on the amount of data that can be annotated in a timely manner depending on how intensive the annotation scheme will be. An estimate of how many tweets can be reasonably annotated given the chosen strategy can serve as an upper-bound of the training data sample size.

Annotating a training dataset of tweets for sentiment analysis

In order to provide labels from which a model will learn, the training dataset must be annotated based on criteria determined by a researcher. Here, we focus on a theory-driven approach, where researchers create criteria for annotation based on theories in emotion science and ask human annotators to apply those criteria when labeling tweets (see e.g., Brady, McLoughlin, Doan, & Crockett, 2020; Hoover et al., 2020). This approach is more resource-intensive but can produce high-quality labels that will hopefully increase discriminant validity between two classes (emotion expression present vs. not present in a given tweet).

Human annotation can be achieved by various means, but two common approaches are using an “expert” approach vs. a “crowd-sourcing” approach. Training “expert” annotators relies on a smaller number of individuals (for instance, a dedicated group of 4 research assistants) who can establish high group reliability by training together (see below) and then label a large number of tweets on their own. The strength of the expert annotator approach is that each individual annotator should demonstrate high precision due to extensive training and comparison with other expert annotators before they label tweets on their own. Using a small number of annotators could come with reduced cost since the researcher does not have to pay many annotators, however a small number of researchers could also spend months carefully training and applying annotators. One weakness is that there is a ceiling on how many tweets any one annotator can be expected to label on their own (yielding a lower sample size for the training dataset).

In the crowd-sourcing approach, hundreds or even thousands of annotators are recruited to label many tweets, such that every tweet in the training dataset is annotated by many individuals. Annotating emotion stimuli with crowd-sourcing has also been called a “normative rating” approach (Lang, Bradley, & Cuthbert, 1997). Applied to binary classification of tweets, the crowd-sourcing approach typically relies on the accuracy of consensus judgments determined by the majority of all raters for a given tweet. For example, if each tweet in a training dataset is labelled by 10 annotators, then the tweet will be labeled based on the label applied by 6 or more annotators. In cases where there are 5 annotators for each possible label (no consensus), the tweet could be dropped from the training set. However, if the criteria for labelling the emotion expression is suitably clear, tweets with no consensus should be very rare given a large enough number of annotators per tweet (with the exception that groups of annotators vary systematically in their labeling, see discussion of annotator political ideology below). While any one crowd-sourced annotator may not be completely consistent in their ratings, the majority rule should ultimately yield high precision. Crowd-sourced annotation can also be used to rate a large number of tweets since any one annotator does not need to rate a large amount of tweets.

Once an annotation method has been chosen, the next step following **Fig. 1** is to identify the emotion of interest and the conceptual features theorized to be associated with the concept of the emotion in the chosen context. This step is perhaps one of the most important steps since it will determine the criteria set by the researcher and taught to the annotators during training. Developing well-defined and precise criteria ultimately determines the success of the sentiment classification project, regardless of how perfect

the computational model is. Even the most accurate machine learning model will make poor classifications if the labels it learns from are poor, also known as the “garbage in, garbage out” principle. In order to ensure the criteria are well-defined, researchers can leverage theories in emotion science to identify conceptual features associated with various emotions.

In measuring moral outrage, Brady and colleagues (Brady, McLoughlin, et al., 2020) leveraged research in moral psychology and emotion science to determine three key conceptual criteria for determining whether a tweet expressed moral outrage (Gintis, Bowles, Boyd, & Fehr, 2005; Salerno & Peter-Hagene, 2013; Tetlock, Kristel, Elson, Green, & Lerner, 2000): a tweet expresses moral outrage if (1) it expresses feelings in response to a violation of the user’s sense of morality, (2) it contains expressions of negative, high-arousal feelings such as anger, disgust and/or contempt, and (3) there is an implication or explicit mention of blaming or punishment.

It can also be useful to think of related conceptual features that could easily be confused for the emotion expression in question to reduce conceptual confusion among the annotators. For instance, often on social media people “troll” others in moral or political contexts where they seemingly express anger or disgust that might make the expression appear similar to moral outrage. However, there is usually also some perceptible element of “poking fun” such as the use of crying-laughing emoji at the end of the trolling tweet. By identifying positive conceptual features as well as related but negative conceptual features of the emotion in question, annotators can have a better idea of what they are looking for in the tweet text.

Once conceptual criteria have been set, annotators must be trained to detect the conceptual features in the text itself (recognizing conceptual features as linguistic features). Annotators should be given the criteria to study as well as clear examples of tweets that represent positive and negative classifications of the specific emotion expression. **Table 1** presents examples of positive and negative classifications for moral outrage in tweets, and each example should come with a detailed explanation of why the tweet is classified as containing outrage. For example, the first tweet in **Table 1** is considered a positive instance of outrage expression because we see evidence of high-arousal, negative feelings with the use of swearing and name-calling. The context is clearly morally-relevant in the sense that it is explicitly political, and there is also an explicit element of blame and desire for punishment, as the author calls for @realDonaldTrump to be removed from office.

Tweet	Classification
<i>@realDonaldTrump Fuck you and your goon cabinet members. You've gone way too far with the SHIT policies and should be kicked out of office</i>	Outrage
<i>Jeff sessions is completely unfit for his job and is making a series of horrible judgments. We must stop the damage he's causing!</i>	Outrage
<i>The NRA is just trying to help protect our second amendment rights, that's it</i>	Non-outrage
<i>The White house can't name one mass shooting that would have been stopped by new guns laws #2a</i>	Non-outrage

Table 1. Examples of outrage and non-outrage tweets used to train annotators.

For both the expert and crowd-sourced annotation methods, it is important to establish empirically that annotators are applying the researcher's criteria consistently. Toward this end, an interclass correlation (ICC) can be established (Shrout & Fleiss,

1979). The ICC assesses how consistently annotators are applying the label of moral outrage. The calculation of ICC depends on the specific annotation situation (see Shrout & Fleiss, 1979 for more details), and ICC values around .80 generally indicate that the labels were applied consistently across annotators (Xu & Shrout, 2014).

Another consideration, if the researcher aims to measure more complex emotion expressions on social media such as moral outrage, is the demographic characteristics of the annotators and how they might impact their labels systematically. In the case of moral outrage, one obvious issue is whether annotators of different political ideologies will perceive moral outrage expression differently depending on what moral viewpoint is being expressed by the tweet they are labeling. Brady et al., (2020) take several steps to address this issue such as establishing that annotator consistency does not systematically vary in ideologically homogenous vs. heterogenous groups (see Brady et al., 2020, SI Appendix).

Training and evaluating a machine learning model for sentiment analysis

Once the training dataset is labeled, there are several steps required to develop a model that can validly classify specific emotion expressions. Here, we focus on the use of neural network architectures which have been shown to produce state-of-the-art performance in sentiment analysis tasks (Zhang, Wang, & Liu, 2018). A workflow for developing an emotion expression classifier based on a neural network architecture is presented in **Fig. 2** and discussed in detail below.

Fig. 2. Workflow for developing an emotion expression classifier based on a neural network architecture for use with Twitter data.

Preprocessing of raw text is especially important for sentiment analysis on Twitter where tweets carry many features that are heuristically and empirically found to contribute minimally to successful sentiment classification (Kharde & Sonawane, 2016). Steps for preprocessing Twitter data for sentiment analysis are presented in **Table 2** below. These preprocessing steps help reduce data sparsity and substantially shrinks the feature space for models which can lead to better performance in Twitter sentiment analysis (Hassan, Abbasi, & Zeng, 2013).

Preprocessing Step	Description
Lowercasing	All words / letters are converted to lowercase
Remove URLs	Removal of URLs in text e.g., “https://...”
Remove #hashtags	Removal of hashtags
Remove @ symbol	Removal of references to usernames (e.g. @realDonaldTrump) such as mentions and replies
Remove “RT” and “cc”	Remove text that signals a retweet or someone using the “carbon copy” phrase to tag another user
Remove non-English words	Ensure document only contains tweets of the same language
Remove white space	Remove unnecessary spacing / padding
Remove special characters	Remove punctuation symbols that do not add substance
Remove stop words	Stop words are common words such as articles, prepositions and pronouns that are unlikely to explain variance in sentiment
Stemming / lemmatization	Reduce different grammatically, inflectional and derivative forms of a word to a common base form (e.g., “ran”, “runs”, “running” are all reduced to “run”)

Table 2. Standard preprocessing steps for preparing raw Twitter data for sentiment analysis.

Once tweets are preprocessed, the text is formatted in a manner that can enhance performance of machine learning models. For use of neural network classification models, this also entails converting individual words into numerical vectors, and here we focus on the use of dense word vectors, or *word embeddings*. Word embeddings are useful for capturing meaning of words in a given context because they represent different words that have similar meanings as closer in vector space. The use of word embeddings has been highly successful for the purposes of sentiment analysis on Twitter (Severyn & Moschitti, 2015). In theory, a researcher could train a model to learn word embeddings based on their own dataset of tweets alone but often the researcher does not have a large enough training dataset for the word embedding

model to capture the full semantic space of the language used in the given context (Chollet, 2018). However, there are pretrained word embedding models available that have been trained on billions of unlabeled tweets, such as the Twitter version of Global Vectors for Word Representation (GloVe; Pennington, Socher, & Manning, 2014). The GloVe model can then be fine-tuned with the researcher's existing dataset to improve accuracy in sentiment classification (see, e.g. Chollet, 2018).

The word embeddings represent the words in a vector space to capture their meaning, and then the neural network architecture is built on the word embeddings for classification of the emotion expression using the training data labels (see **Fig. 2**). Broadly speaking, the neural network model does the “feature extraction”, where salient features associated with positive instances of the emotion expression (in terms of the word embedding representations) are learned through an iterative process. There are several types of neural network architectures that can be employed (for a review, see Goldberg, 2017), and researchers can test different architectures to determine which produces the best results. For example, Brady et al. (2020) use a recurrent neural network architecture that showed considerable improvement in moral outrage classification of tweets compared to linear machine-learning methods such as an implementation with a random forest model. Neural network architectures are readily implemented using packages such as *keras* which is available for both *Python* and *R* programming languages.

In order to evaluate the model, one standard approach is to use *k-fold cross-validation*, which is a resampling procedure that performs test-training splits based on *k* groups of data. For example, Brady et al. (2020) test their model using a standard 5-fold

cross-validation approach, where 5 groups of randomly selected tweets are formed, and the model is repeatedly trained on 4 out of the 5 groups, and then tested on the hold-out group. This method has been shown to provide less biased estimates of accuracy for unseen data (James, Witten, Hastie, & Tibshirani, 2013). There are various metrics that can be used to assess model performance such as raw accuracy, precision, recall and F-1 scores.

How accurate should a researcher expect their models to be? Recent work has shown that state-of-the-art methods in sentiment classification achieve accuracies across various test sets in the range of 60-80% (Barnes, Klinger, & Schulte im Walde, 2017). However, raw accuracy may not be the only consideration, depending on how the researcher would like to apply the emotion expression classifier. Precision and recall metrics can help determine the type of errors that comprises the model's accuracy (e.g., false alarm rates), which can help the researcher assess their model on the basis of whether of the research question is more tolerant of type I vs. type II errors.

Another practical test to determine how well the model is performing is to substitute the human annotators in a group with the classification model and determine if the model applies emotion labels like the group of humans. For instance, Brady et al. (2020) had 8 extensively-trained annotators rate a dataset of 500 tweets for whether they express moral outrage. They computed the ICC for the group, which revealed that the annotators showed strong consistency with $ICC(2,8) = .88$, 95% CI = [.86,.89]. Next, they substituted 1 human with their outrage classifier in an iterative fashion, so the ICC was computed repeatedly when substituting the model with every human once (see Fig. 3). The mean ICC of this process was then computed at $ICC(2,8) = .87$, 95% CI =

[.86, .89]. This process showed that their outrage classifier was classifying tweets in a manner that was statistically indistinguishable for human annotators in terms of consistency.

Fig. 3. Method for accessing how consistent an emotion classification model is compared to trained human annotators, using an example with 3 annotators. An initial ICC estimate is determined based on a group of human annotators. Next, the model is substituted for each human once and an ICC is computed for the group after each substitution. The mean ICC produced from this procedure can be compared to the initial human annotator ICC to determine whether the model is consistent with human annotation.

While k-fold cross-validation can provide suitable estimates of a model's ability to accurately label unseen data, it is still purely based on the specific context represented by the tweets that were used as training data. For instance, if a researcher trains an emotion expression classifier on 10,000 tweets that were communications about a specific event (e.g., a mass shooting event), there is no guarantee that the model will accurately classify the emotion expression in out-of-sample data. In fact, much previous work in tweet sentiment classification suffers from poor generalizability to out-of-sample datasets (Barnes et al., 2017). In other words, we can draw a distinction between

internal generalizability (as accessed with cross-validation) and external generalizability. If a researcher intends to apply the classifier to a new Twitter data that might represent a different context (e.g., a Twitter conversation about pets when the original model was trained on political conversations), some effort should be made to empirically validate the expected accuracy of the out-of-sample classifications. One difficulty is that this may require considerable effort since new annotations would be required to test the model against labeled data.

Besides accessing generalizability across topics, the researcher may also worry about generalization across *time*. In particular on social media, norms of expression can change rapidly, making it possible that linguistic features associated with a specific emotion expression can also vary over time. Some effort to test the robustness of the model across time empirically can help address this issue. For instance, Brady et al. (2020) collected 1000 tweets about contentious political topics across 4 different years and had them labeled by trained annotators for whether or not they contained moral outrage. This test demonstrated that their outrage classifier showed small variation in accuracy but remained consistent across the different years of data collection. Indeed, one benefit of using psychological theory to identify conceptual features of an emotion (rather than measuring emotion expression in a purely data-driven manner) is that these features might represent a signal of the emotion over and above contextual variation in language, thus increasing generalizability.

Psychological hypothesis testing with emotion classifiers

We end with a few considerations for researchers who wish to apply an emotion classifier to testing a psychological hypothesis. First, it is important to keep in mind the distinction between emotion expression (measured by emotion classifiers) and emotion experience. A large body of work in emotion science has demonstrated with overwhelming evidence that expressions of emotions do not represent direct “read outs” of an underlying emotion experience (for a review, see Barrett, Adolphs, Marsella, Martinez, & Pollak, 2019). According to this research, which draws almost exclusively from data on facial expressions, the inference that a user on social media is experiencing anger at their computer because there is emotion language in their tweet might be a highly tenuous inference. However, there is some reason to be more enthusiastic about the inference from expression to experience when it comes to emotion expressions on social media that relies on symbolic representations in language. Language organizes highly variable perceptual instances into broader categories, and thus it can act as a filter that reduces the variability of instances of emotion expressions for a given emotion experience (Lindquist & Gendron, 2013). In other words, emotion expression in language is more likely to have a direct correspondence with the an individual’s experience of an emotion that is represented itself by a categorization of underlying valence and arousal.

Nonetheless, if a researcher is classifying tweets specifically as a proxy for a person’s emotion experience offline (such as using emotion expressions in tweets as evidence for state-level rates of mood disorders) it could be helpful to validate the emotion expression classification with an offline measure of emotion experience. On the

other hand, many research questions regarding how emotional tweets influence perceptions and behaviors do not necessarily require the inference that tweet expressions are associated with offline emotion experiences. Users on social media have direct access to the representation of other users' emotion based on the expressions in tweets, but no access to their actual emotional experience at the computer or smart phone. Thus, the expression present in the tweet itself is the information that influences perceivers. For example, in research investigating how the expression of moral emotions can draw increased engagement during the context of political conversations, the spread of the emotion expressions themselves and the consequences for online social networks is the question of interest (Brady, Gantman, & Van Bavel, 2020; Brady, Wills, Burkart, Jost, & Van Bavel, 2019; Brady, Wills, Jost, Tucker, & Van Bavel, 2017). Research questions such as those may not require any assumptions about the relationship between emotion expression and underlying emotional experiences.

Another important consideration for psychological hypothesis testing based on emotion expressions on social media is that communication norms change dynamically in computer mediated environments (Postmes, Spears, & Lea, 2000). This suggests that although the researcher can do their best to identify broad signals (features) of an emotion after selection of a social context, sub groups within the context such as a specific social network might evolve unique expressions over time. One extreme case is the rapid evolution of context-specific memes often used as tools to expression specific emotions on Twitter (for an interesting case study, see Ma, 2016). Researchers should take care to name the boundary conditions of their effects based on specific social

media context in which the emotion expressions were classified (Van Bavel, Mende-Siedlecki, Brady, & Reiner, 2016).

Conclusion

With over 3 billion users across the world, social media platforms provide researchers with an opportunity to study a massive volume of emotion expressions unfolding in real-time social interactions. As new computational tools are becoming available, social media data increasingly affords researchers the ability to study psychological questions about emotions at an unprecedented scale. This chapter outlined theoretical and practical considerations to guide researchers in this endeavor. We hope that the principles outlined here can apply broadly to the study and measurement of emotions as new social technologies unfold rapidly in the digital age.

References

- Banko, M., & Brill, E. (2001). Mitigating the paucity-of-data problem. *Computational Linguistics*, 1–5.
- Barnes, J., Klinger, R., & Schulte im Walde, S. (2017). Assessing State-of-the-Art Sentiment Models on State-of-the-Art Sentiment Datasets. In *Proceedings of the 8th Workshop on Computational Approaches to Subjectivity, Sentiment and Social Media Analysis* (pp. 2–12). Copenhagen, Denmark: Association for Computational Linguistics.
- Barrett, L. F. (2013). Psychological construction: The darwinian approach to the science

- of emotion. *Emotion Review*, 5(4), 379–389.
- Barrett, L. F., Adolphs, R., Marsella, S., Martinez, A. M., & Pollak, S. D. (2019). Emotional Expressions Reconsidered: Challenges to Inferring Emotion From Human Facial Movements. *Psychological Science in the Public Interest*, 20(1), 1–68.
- Batista, G. E. A. P. A., Prati, R. C., & Monard, M. C. (2004). A study of the behavior of several methods for balancing machine learning training data. *ACM SIGKDD Explorations Newsletter*, 6(1), 20.
- Brady, W. J., Gantman, A. P., & Van Bavel, J. J. (2020). Attentional Capture Helps Explain Why Moral and Emotional Content Go Viral. *Journal of Experimental Psychology: General*, 149(4), 746–756.
- Brady, W. J., McLoughlin, K., Doan, T., & Crockett, M. J. (2020). Social Learning of Moral Outrage in Online Social Networks. Retrieved from <https://osf.io/9he4n/>
- Brady, W. J., Wills, J. A., Burkart, D., Jost, J. T., & Van Bavel, J. J. (2019). An Ideological Asymmetry in the Diffusion of Moralized Content on Social Media Among Political Leaders. *Journal of Experimental Psychology: General*, 148(10), 1802–1813.
- Brady, W. J., Wills, J. A., Jost, J. T., Tucker, J. A., & Van Bavel, J. J. (2017). Emotion shapes the diffusion of moralized content in social networks. *Proceedings of the National Academy of Sciences*, 114(28), 7313–7318.
- Chollet, F. (2018). *Deep Learning with R*. Shelter Island, NY: Manning Publications.
- Gintis, H., Bowles, S., Boyd, R., & Fehr, E. (2005). *Moral sentiments and material interests: The foundations of cooperation in economic life*. Cambridge, MA: MIT

Press.

Goldberg, Y. (2017). *Neural network methods for natural language processing*.

Synthesis Lectures on Human Language Technologies (Vol. 10). Morgan & Claypool Publishers.

Haidt, J. (2003). The moral emotions. In R. J. Davidson, K. R. Scherer, & H. H.

Goldsmith (Eds.), *Handbook of affective sciences* (pp. 572–595). New York, NY: Oxford University Press.

Hassan, A., Abbasi, A., & Zeng, D. (2013). Twitter sentiment analysis: A bootstrap ensemble framework. In *Proceedings -*

SocialCom/PASSAT/BigData/EconCom/BioMedCom 2013 (pp. 357–364).

Hoover, J., Portillo-Wightman, G., Yeh, L., Havaladar, S., Davani, A. M., Lin, Y., ...

Dehghani, M. (2020). Moral Foundations Twitter Corpus: A Collection of 35k Tweets Annotated for Moral Sentiment. *Social Psychological and Personality Science*.

James, G., Witten, D., Hastie, T., & Tibshirani, R. (2013). *An introduction to statistical learning* (Vol. 112). Springer.

Kearney, M. W. (2018). rtweet: Collecting Twitter data.

Keltner, D., & Haidt, J. (2003). Approaching awe, a moral, spiritual, and aesthetic emotion. *Cognition & Emotion*.

Kharde, V. A., & Sonawane, S. S. (2016). Sentiment Analysis of Twitter Data: A Survey of Techniques. *International Journal of Computer Applications*, 139(11), 5–15.

Kotsiantis, S. B. (2007). Supervised machine learning: A review of classification techniques. *Informatica*, 31(3), 249–268.

- Lang, P. J., Bradley, M. M., & Cuthbert, B. N. (1997). International affective picture system (IAPS): Technical manual and affective ratings. *NIMH Center for the Study of Emotion and Attention*.
- Lindquist, K. A., & Gendron, M. (2013). What's in a word? Language constructs emotion perception. *Emotion Review*, 5(1), 66–71.
- Lindquist, K. A., MacCormack, J. K., & Shablack, H. (2015). The role of language in emotion: Predictions from psychological constructionism. *Frontiers in Psychology*, 6.
- Ma, X. (2016). From internet memes to emoticon engineering: Insights from the baozou comic phenomenon in China. In *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*.
- Pennington, J., Socher, R., & Manning, C. (2014). Glove: Global Vectors for Word Representation. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*.
- Postmes, T., Spears, R., & Lea, M. (2000). The formation of group norms in computer-mediated communication. *Human Communication Research*, 26(3), 341–371.
- Russell, J. A., & Barrett, L. F. (1999). Core affect, prototypical emotional episodes, and other things called emotion: Dissecting the elephant. *Journal of Personality and Social Psychology*.
- Salerno, J. M., & Peter-Hagene, L. C. (2013). The interactive effect of anger and disgust on moral outrage and judgments. *Psychological Science*, 24(10), 2069–2078.
- Severyn, A., & Moschitti, A. (2015). Training Deep Convolutional Neural Network for

- Twitter Sentiment Classification. In *UNITN* (pp. 464–469).
- Shrout, P. E., & Fleiss, J. L. (1979). Intraclass correlations: Uses in assessing rater reliability. *Psychological Bulletin*.
- Silge, J., & Robinson, D. (2017). *Text mining with R: A tidy approach*. “ O’Reilly Media, Inc.”
- Tetlock, P E, Kristel, O. V, Elson, S. B., Green, M. C., & Lerner, J. S. (2000). The psychology of the unthinkable: taboo trade-offs, forbidden base rates, and heretical counterfactuals. *Journal of Personality and Social Psychology*, 78(5), 853–870.
- Tetlock, Philip E., Kristel, O. V., Elson, S. B., Green, M. C., & Lerner, J. S. (2000). The psychology of the unthinkable: Taboo trade-offs, forbidden base rates, and heretical counterfactuals. *Journal of Personality and Social Psychology*, 78(5), 853–870.
- Tweepy. (2017).
- Van Bavel, J. J., Mende-Siedlecki, P., Brady, W. J., & Reinero, D. A. (2016). Contextual sensitivity in scientific reproducibility. *Proceedings of the National Academy of Sciences*, 113(23), 6454–6459.
- Xu, J. H., & Shrout, P. E. (2014). Intraclass correlation coefficients. *The Encyclopedia of Clinical Psychology*, 1–6.
- Zhang, L., Wang, S., & Liu, B. (2018). Deep learning for sentiment analysis: A survey. *Wiley Interdisciplinary Reviews: Data Mining and Knowledge Discovery*, 8(4), e1253.