Available online at www.sciencedirect.com

ScienceDirect

Journal of Interactive Marketing xx (xxxx) xxx

JOURNAL OF INTERACTIVE MARKETING

www.elsevier.com/locate/intmar

A Brand-New Look at You: Predicting Brand Personality in Social Media Networks with Machine Learning

Utku Pamuksuz [a] & Joseph T. Yun [b,*] & Ashlee Humphreys [c]

[a] Department of Business Administration, Gies College of Business, University of Illinois at Urbana Champaign, 446 Wohlers Hall, 1206 S 6th St., Champaign, IL 61820, United States of America

[b] Department of Accountancy, Gies College of Business, University of Illinois at Urbana Champaign, 360 Wohlers Hall, 1206 S 6th St., Champaign, IL 61820, United States of America

[c] Medill School of Journalism, Media, Integrated Marketing Communications, Northwestern University, 1845 Sheridan Road, Evanston, IL 60208, United States of America

Abstract

Tools for analyzing social media text data to gain marketing insight have recently emerged. While a wealth of research has focused on automated human personality assessment, little research has focused on advancing methods for obtaining brand personality from social media content. Brand personality is a nuanced aspect of brands that has a consistent set of traits aside from its functional benefits. In this study, we introduce a novel, automated, and generalizable data analytics approach to extract near real-time estimates of brand personalities in social media networks. This method can be used to track attempts to change brand personality over time, measure brand personality of competitors, and assess congruence in brand personality. Applied to consumer data, firms can assess how consumers perceive brand personality and study the effects of brand–consumer congruence in personality. Our approach develops a novel hybrid machine learning algorithmic design (LDA2Vec), which bypasses often extensive manual coding tasks, thus providing an adaptable and scalable tool that can be used for a range of management studies. Our approach enhances the theoretical understanding of channeled and perceived brand personality as it is represented in social media networks and provides practitioners with the ability to foster branding strategies by using big data resources.

Keywords: Brand personality; Social media networks; Data analytics; Machine learning; Word embeddings; Document embeddings; Transformers; LDA2Vec

Introduction

A brand is regarded as one of the most valuable assets owned by a firm. Strong and distinguished brands significantly enhance firm performance (Madden, 2006) and play a key role in building consumer perceptions about products and firms. In fact, scholars have demonstrated that brands exhibit personalities similar to human personalities (Aaker, 1997). As brands build these personalities, consumers sometimes interact with brands as if they were human (Levy, 1985) and form meaningful relationships with brands (Fournier, 1998). Naturally, consumers seek brands with personalities that are congruent with either their own or their aspirational (ideal) personalities (Batra, Ahuvia, & Bagozzi, 2012; Sirgy, 1982). Developing and measuring a strong brand personality is therefore key to many marketing efforts.

The growth of social media platforms has sparked an opportunity to understand how firms foster meaningful brands and use brand personality in these media. Due to the growing potential for social media networks to be used as efficient marketing and brand-building platforms, firms have increasingly expanded branding efforts to this digital interactive medium. As a result, branding in the digital medium has become an essential form of marketing communication to convey core brand personality. Having the ability to use an effective marketing communications strategy to distinguish a brand from competitors fosters customer relationships and can build brand equity. Thus, for a firm to understand their brand

* Corresponding author.
E-mail addresses: pamuksu2@illinois.edu (U. Pamuksuz), jtyun@illinois.edu (J.T. Yun), a-humphreys@northwestern.edu (A. Humphreys).

https://doi.org/10.1016/j.intmar.2021.05.001
1094-9968© 2021

personality in social media, they must develop the capability to assess both the intended and perceived brand personality through the generation of branded content and interactive dialog with consumers.

Since 1997, most of the marketing literature has embraced self-reporting tools (e.g., Likert scale surveys) based on Aaker's (1997) scale to assess brand personality. Such self-reporting tools are often expensive, labor-intensive, and time-consuming. They exhibit reporting bias issues, and the results can become outdated very quickly. In this age of data-driven analytics, brand personalities are also projected real-time on the brand's social media accounts, and traditional methods of surveying brand personality cannot cope with the speed of brand social media content creation.

Further, brands are increasingly embedded in a quickly changing social environment to which they must respond and adapt (Bhagwat, Warren, Beck, & Watson IV, 2020; Holt, 2016). Being able to measure brand personality in real time as firms both create messages and track response to those messages on social media is now an essential marketing function. For example, a brand like Walmart has a folksy, hometown brand personality that is high on competence and sincerity and low in sophistication (Arnold, Kozinets, & Handelman, 2001). If Walmart needs to issue a series of new messages about COVID, first about in-store safety measures and then about vaccine availability, what language best conveys these messages? Do consumers receive these messages in line with Walmart's existing values and personality? Are these values distinct from other competitors on social media? If Walmart creates other, more typical product and service messages, to what degree are those messages congruent with the brand personality reflected in previous messages? All of these questions can be answered by using the approach and tool proposed in this research.

In general, research in personality on social media networks is positioned at the intersection of individuals, organizations, and technology, and using advanced analytics to understand social data is an emerging research field across different academic disciplines including psychology, marketing, management, and information systems (Culotta & Cutler, 2016; Humphreys & Wang, 2018; Kern et al., 2016; Netzer, Feldman, Goldenberg, & Fresko, 2012). Typically, extant analytic methods require extensive content customization and static closed vocabulary approaches that show limitations in terms of comprehensiveness. Some recent work (Kern et al., 2016; Park et al., 2015) has conducted automated human personality assessments by using open vocabulary approaches that integrate unsupervised machine learning techniques with multiple feature selection methods to build robust language models on social media networks. Despite rigorous research efforts in human personality assessment in social media content, studies are limited in the brand personality domain. Thus, we were motivated to develop a data analytic approach to detect and analyze brand personality from social media data.

Text analysis has emerged as a key way to understand many classic marketing issues (Berger et al., 2020). While previous work has used social media data to map market structure (Lee & Bradlow, 2011; Netzer et al., 2012) and brand positions (Tirunillai & Tellis, 2014), there are no extant approaches, to our knowledge, developed to detect brand personality using machine learning from textual data. The closest and most recent work used a part-of-speech keyword extraction approach to detect perceptions of brand personality from an online fashion blog and various fashion company websites and Facebook pages (Ranfagni, Camiciottoli, & Faraoni, 2016). Other approaches have used linking to make inferences about brand personality (Culotta & Cutler, 2016). In this paper, we build upon this research by introducing a fully automated machine learning approach for marketing scholars and practitioners to analyze how personalities of brands are channeled and perceived via social media, and we offer a foundation for future advances in machine-learned approaches to examining brand–consumer relationships occurring in social media networks.

The automated measurement of brand personality from social media data has three important implications for marketing scholars and managers. First, firms can use the tool to perform self-auditing, either to ensure consistency in the communication of brand personality or to track strategic changes in brand personality over time. Secondly, the method can be used to measure brand personality in the context of consumer mentions and uses. As some scholars have noted, brand personality is not always under the sole discretion of the firm (Diamond et al., 2009; Thompson, Rindfleisch, & Arsel, 2006), as consumers interpret brand action and meaning in social context. This tool allows brand managers to assess consumer perceptions of brand personality by evaluating the language surrounding brand mentions online. Lastly, the tool can be used to examine the effects of self-congruency and brand personality. That is, what are the effects of “on-brand” and “off-brand” posts on consumer engagement with the post? Further, the tool can be used to examine hypothesized effects of personality matching between consumers and brands based on theories of cognitive dissonance.

We integrate closed vocabulary-based methods, supervised machine learning, and unsupervised open vocabulary methods into one refined model. At a high level, our algorithmic design takes the unstructured text data from social media accounts and returns scores for Aaker's (1997) five brand personality dimensions; Sincerity, Excitement, Competence, Ruggedness, and Sophistication in real-time, which provides a novel method for analyzing social media content that may considerably increase the scale and scope of brand research.

Prior Literature

Brand Personality

The term brand personality was first coined by Martineau (1958) who proposed that consumer behavior is dependent upon personality rather than objective reality by referring to a set of human characteristics related to a brand. For instance, users have characterized the brand personality of Mercedes Benz as *upscale* and *aspirational*, while Calvin Klein's brand personality has been characterized as *sexy* and *sophisticated*. There are product-related and non-product-related factors that drive the formation and perception of brand personality (Aaker, 1995). On one hand, the attributes of a product can signal a personality, such as a high-priced Burberry scarf that might portray signals of wealth, style, and perhaps a bit of arrogance. On the other hand, non-product-related signals include age, symbols, employees, CEOs, celebrity endorsers, and sponsorships. For instance, Red Bull's sponsorship of the International Ice-Skating Championship may reinforce the brand's *offbeat* and *youthful* personality. Considering all of these factors with brand personality formation, the growth of social media has sparked opportunities for firms to convey their branding efforts by performing integrated marketing activities with less effort and cost than previously. Traditionally, brands have only a few ways to communicate brand personality through mainstream media and service person interaction (Aaker, 1997). However, in social media, a brand has manifold opportunities for expressing its personality in tweets, posts, pictures, and by publicly observed interactions with consumers, all of which are potentially shared and circulated with other consumers.

Once properly formed, brand personality can be a longstanding asset for firms. The personification of brands may provide an important point of differentiation from competitors and assist corporations in developing brand equity (Ross, 2008). Marketers, therefore, need to ensure that a brand's personality is channeled consistently to the consumers and interpreted in ways intended by the firm. When a brand consistently nurtures its brand personality, the relationships between the brand and its consumers evolve in a way that is characterized by the values inherent in the brand's personality (Fournier, 1998). Corporate brands exhibit a brand personality that represents various characteristics of the brand, and this personality evolves largely from the brand's fundamental values and positioning (Harris & de Chernatony, 2001). The goal of corporate branding efforts is to develop a brand that is perceived as unique and positively valanced (Keller & Lehmann, 2006). Consumers' perceptions and behaviors are influenced by the brand personality that is channeled by the focal firm. However, channeled brand personality may also be partly a product of social consensus (Diamond et al., 2009; Thompson et al., 2006) which can be measured through social media data—what consumers and other thought leaders say about the brand. Thus, this work enables managers to measure and track social consensus towards their brand personality, and thus attempt to shape it. We illustrate the potential of using our tool to measure these perceptions of brand personality (perceived brand personality) by measuring alignment between brand communications and other social media content referring to a brand.

To date, Aaker's (1997) brand personality scale is the most widely employed brand personality measure for a theoretically understanding of the brand personality construct. Aaker analyzed the individual ratings of 37 brands on 114 personality traits by 613 respondents from the United States and developed a reliable, valid, and generalizable scale to measure brand personality. As a result, brand personality scales are composed of 42 traits. These traits are defined into five dimensions: Sincerity, Excitement, Competence, Sophistication, and Ruggedness. Sincerity captures traits such as down-to-earth, cheerful, sincere, and friendly. Excitement indicates traits including daring, young, trendy, imaginative, unique, and independent. Competence is represented by traits such as intelligent, reliable, secure, and successful. Sophistication is characterized by traits including upper-class, glamorous, charming, and good-looking. Finally, ruggedness encapsulates traits such as masculine, tough, and outdoorsy. Previous computational methods for analyzing brand personality from social media text have used dictionary-based approaches (Opoku, Abratt, & Pitt, 2006; Xu et al., 2016), but there are many methods that have been developed to predict human personality from social media text that might be usefully employed to more accurately measure brand personality.

Computational Detection of Brand Personality

In the brand literature, conventional empirical methods including self-reported surveys and standard personality questionnaires have been widely used for data collection and hypothesis testing efforts (Aaker, 1997; Carr, 1996). Yet the emergence of social media, such as Facebook and Twitter, have created novel online platforms for brands to interact with consumers. Such platforms have already transformed consumer behavior in terms of the creation of large amounts of user-generated content and mass consumption of this content, this transformation has generated vast data sources for marketing scholars and practitioners to unlock new consumer insights by using modern data analytic techniques (Zhang, Bhattacharyya, & Ram, 2016). As a result, the emergence of social media networks not only provides unbounded data sources to empirically test propositions for various disciplines but also enables the implementation of advanced analytical methods that considerably enhance the scope and scale of personality research (Golbeck, Robles, Edmondson, & Turner, 2011). We note that there is a need for such analytical advancements to be applied to the realm of brand personality, to assess how brand personalities are being channeled and to measure how they are being perceived by consumers (Aaker, 1995). There have been other attempts to measure brand personality. For example, Xu et al. (2016) conducted a predictive analysis of the drivers of brand personality embodied in social media. The authors focused on the factors that drive brand personality instead of direct brand personality detection from social media content. They used questionnaires and a closed-vocabulary approach (LIWC) as an illustration of the consumer-perceived brand personality without employing machine learning and advanced analytic implementations such as open-vocabulary based approaches (e.g. unsupervised cluster detection) and social network analytics (e.g. link prediction in a social network).

In another related work, Culotta and Cutler (2016) developed an automated data analytics tool to predict brand perceptions from Twitter. The first notable difference in our work is that we focus on detecting brand personality, which is a different construct than brand perception. The second major difference is in methodologies. While previous work follows a network similarity approach, we use a machine-learned text-analysis approach. Specifically, prior work has first categorized exemplar social media accounts by a certain perception category (e.g. eco-friendly) and then used Twitter Lists to identify Twitter accounts that exemplify eco-friendliness (e.g. @smartcarusa). Then, a brand perception score is calculated for a category for the brand Twitter account being analyzed by calculating the similarity of followers that followed that brand Twitter account and the exemplar Twitter account. Thus, if you were testing how eco-friendly @Coca-Cola was, their model would compare the followers of @smartcarusa and @Coca-Cola. Our method, in contrast, detects brand personality directly from social media posts. This can be advantageous because it allows a widely known brand like Coke to present as, for example, eco-friendly despite not being categorized that way a priori.

Ranfagni et al. (2016) used part-of-speech tagging keyword extraction methods to compare brand personality-related adjectives (e.g., daring, affordable, romantic) across brands and between brands and consumers. Using their method, researchers can measure the degree of alignment between company-defined and consumer-perceived brand personality, similarities between brands, and consumer-perceived similarity across brands.

With regard to brand research using machine-learning, Huang and Luo (2016) apply supervised machine learning to elicit consumer preferences. In addition, Jacobs, Donkers, and Fok (2016) integrate unsupervised learning for better identification of items purchased together. To our knowledge, there are no extant machine learning approaches for detecting brand personality. Thus, we believe that the introduction of a novel, automated, and generalizable method to extract near real-time estimates of brand personalities from the generated text content may have strong implications for the academic community in enhancing the theoretical understanding of the brand personality construct and for managers seeking to measure brand communication.

Methodology

Our methodology was guided by design science research principles to report relevance and enhance the rigor of our research process and results (Peffers, Tuunanen, Rothenberger, & Chatterjee, 2008). According to Peffers et al. (2008), we first introduce the design and development process in the following subsections: First, we start with a computational design in reporting on our sample selection, data sources, and machine learning implementation phases. Second, we demonstrate the results. Third, we evaluate the results and test robustness. Finally, we conclude the paper with future research considerations utilizing our tool and the implications of our work for researchers and practitioners.

Sample Selection and Data Sources

To test the generalizability of our approach across brands, we use a wide range of brands from a variety of sectors. To select brands, we used the website millwardbrown.com, which maintains a large selection of brands categorized by sector including apparel, cars, luxury, personal care, food drink, financial institutions, technology, telecommunication, insurance, and airlines. We chose five well-known brands from each category that totaled up to 100 brand accounts to train the learning model. In this sense, we used a similar concept to Culotta and Cutler (2016) of choosing "exemplar brand accounts" that span multiple, different categories. We then used an additional 20 separate brands to test the results of our framework. These 20 brands for testing were specifically chosen based on their publicly perceived visible personalities from different industries. For example, Patagonia is a brand that signals sincere and authentic personality, whereas technology firms such as Google signal strong and competent personalities. Thus, we assumed a demonstration of our results would be more interpretable if we focused on strong and publicly visible brands in the first stage of testing our algorithm.

For each of these brands, we used a dataset that was previously collected using the Crimson Hexagon cloud platform, in which Twitter and Facebook posts from official brand accounts were collected between June 1, 2014 until Dec 31, 2017 (Pamuksuz, 2017). We had originally considered only Twitter posts, but Facebook posts have more potential for descriptive markers due to the lack of a character limit, thus adding Facebook posts would invariably assist our machine-learned training. Crimson Hexagon warehouses all public Twitter and Facebook data stretching back to 2009. Consequently, we retrieved 266,105 posts in total for the training set and 53,221 posts for the test set. The average number of posts per brand in the training set was 2,274 with a standard deviation of 98 posts across the brands. The range across the brands was a minimum of 387 posts to a maximum of 3,921 posts. We did not have a minimum character limit to the posts, and on average each tweet had 38.4 characters, whereas each Facebook post had 74.1 characters.

Summary Overview of Our Machine-Learned Implementation

Before we present a detailed explanation of each step of our machine-learning process, we provide a summary outline of our machine-learning steps (see Fig. 1 for a visual workflow):

1. We collected raw data with no brand personality labels. To generate labels, which is required to train a learning model to predict any brand's personality in the future, we had two options as outlined below. We chose the second option due to the cost inefficiencies of the first option.
   a. Use human coders to label thousands of posts to obtain a labeled training set.
   b. Develop a sophisticated learning model (LDA2Vec) to label some part of the raw data by using a combination of an unsupervised and self-supervised approach.
2. We used LDA2Vec with previously predefined brand personality word dictionaries to label approximately 30% of the posts with one of the brand personalities dimensions. In order to design a robust automated personality detection tool, this number was not sufficient.

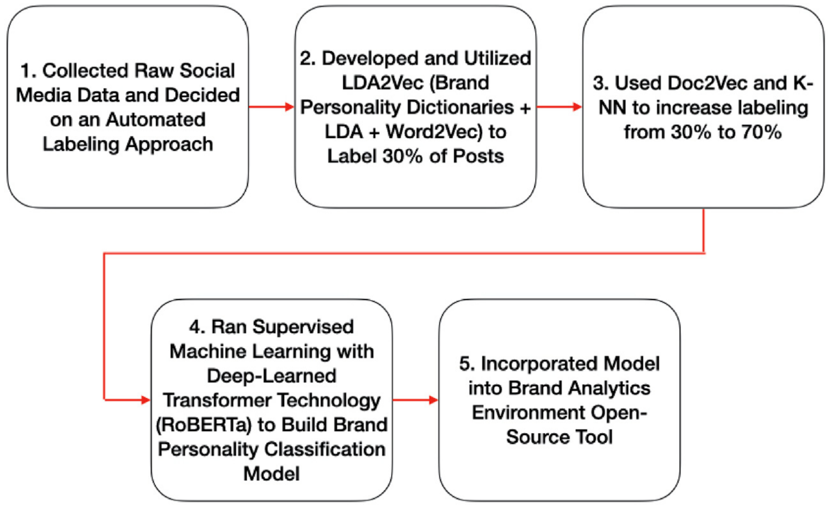

Fig. 1. High-level computational workflow of brand personality modeling.

3. We used Doc2Vec and K-Nearest Neighbors methods in order to increase the size of our automatically labeled training set to address posts that exhibited multi-personality features. This increased our labeled set from 30% to 70% at the end of this phase. This was a large enough training set to build a deep learned model.
4. We developed and tested several machine learning and deep learning methods, and eventually we implemented the highest performing transformer algorithm (RoBERTa) to build our brand personality detection prediction model.
5. Our resulting tool can classify and predict multi-label brand personality traits for each single post. A researcher can combine scores of individual posts to calculate an overall personality result (calculating an average) based on brand analysis needs.

Detailed Explanation of Our Machine-Learned Implementation

Social Media Post Labeling

We designed a novel unsupervised approach to generate labels for the training set using self-supervision or weak learning. Self-supervised learning formulates a supervised learning task where the labeling process is completely automated. This technique is widely used in recent language modeling and machine learning research (Baevski, Schneider, & Auli, 2019; Wang et al., 2019). We adopted this labeling approach for two main reasons. First, we needed the highest possible accuracy on identifying brand personalities from the collected social media corpus which should be close enough to error-free human manual coder performance. Unsupervised methods alleviate the complexity of the data and provide a simpler view for decision-makers, but its performance will be only supportive when it comes to the exact identification of brand personalities within the large unstructured text data. Second, since the cost of human coder knowledge for labeling a relatively large subsample in the training phase is expensive and not practical to train very large datasets, we focused on utilizing LDA (Blei, Ng, & Jordan, 2003), Word2Vec (Mikolov, Chen, Corrado, & Dean, 2013), and Doc2Vec (Le & Mikolov, 2014) unsupervised models along with previously identified brand personality word libraries (Aaker, 1997; Opoku et al., 2006) to implement our supervised methods. As far as we know, this is the first marketing study that has combined a dictionary-based approach, a novel unsupervised machine-learned approach, and self-supervised learning using dimensional feature distance to label social media posts for supervised learning of a brand construct such as brand personality. A simple visualization of the Phase 1 labeling process is shown in Fig. 2.

Mixing LDA Topic Clusters and Word2Vec Word Representations (Phase 1.1)

In phase 1.1 (see Fig. 2), our objective was to take all 266,105 brand posts within a hundred brand documents and find clusters of social media posts that are related to one another across all the brands. Therefore, our first step was to create refined clusters of posts utilizing LDA and Word2Vec (see Moody, 2016, for further technical details). LDA is a Bayesian version of pLSA (probability latent semantic analysis). It employs Dirichlet priors for the document-topic and word-topic distributions, lending itself to better generalization. The Dirichlet distribution provides a way of sampling probability distribution of a specific type. Although LDA is illustrative enough to generate multiple topics per document, it is not sufficient for multi-labeled corpora because, as an unsupervised bag-of-words model, it offers no obvious way of incorporating a supervised label set into its learning procedure. Thus, we incorporate Word2Vec (Mikolov et al., 2013) to leverage both global and local presentations of terms among clusters. Word2Vec is a predictive algorithm for learning word embeddings using a deep neural network model. Embeddings are vector representations of words represented by a set of hidden variables, and each word is represented by a specific embodiment of these variables. Word2Vec directly tries to predict a word from its local neighbors in terms of learned small, dense embedding vectors. For example, the vector for the word cat might be [1,2,3] (real word embedding vectors are

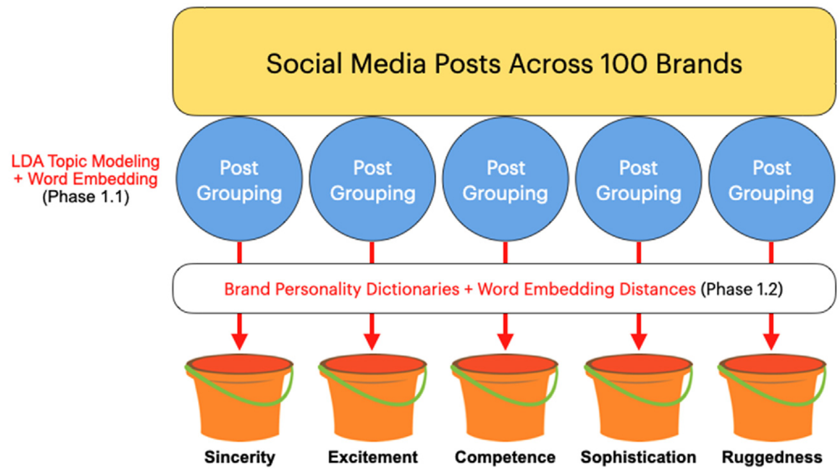

Fig. 2. Dictionary-based + unsupervised machine learning labeling process - LDA2Vec (Phase 1).

far larger in size) and the vector for cats might be [1,2,4]. If we were to plot these vectors in three-dimensional space as [x,y,z], we can imagine that these two points would be very close to one another. With regard to topic modeling, using word embeddings can refine topic models by drawing together more relevant and salient terms within topics (Das, Zaheer, & Dyer, 2015; Moody, 2016). Topic modelings and word embeddings used in two different research communities. Word embeddings come from the neural net research tradition, while topic modelings come from Bayesian model research tradition. Word embedding can be used to improve topic models like LDA2Vec (Das et al., 2015; Moody, 2016).

Although there are several different techniques for topic classification and text clustering tasks used in Marketing Science including the usage of LDA and other baseline topic modeling - LSA, pLSA or clustering algorithms such as K-Means and DBSCAN, we chose to develop this hybrid method, LDA2Vec, to achieve the best labeling performance after monitoring the low performance results of LSA, K-means and DBSCAN when tackling hundreds of dimensions where sparsity occurs upon the integration of word embeddings on distance matrices. For the choice of Word2Vec over other popular pretrained embeddings, including GloVe (Pennington, Socher, & Manning, 2014) and Numberbatch (Speer, 2017), we considered the congruence of language usage of brands highest with the documents used to train Word2Vec. Thus, we decided to apply Word2Vec embeddings to integrate into our hybrid LDA2Vec model.

Thus, as shown in Fig. 2, we used LDA2Vec to create refined clusters of words that describe the various topics being discussed across a hundred brands. We then took each of the social media posts and used feature vectors' distance similarities between the posts and each of the refined topics. We did this to see which posts most accurately represent each of the refined topic clusters since we wanted only the most relevant posts for our machine learning training data. Each social media post participated in only one refined topic cluster, and only posts that were 50% or more similar in vector distance to a refined topic cluster were included to move onto our next step.

Combining Dictionaries and LDA2Vec for Initial Labeling (Phase 1.2)

In the next step, Phase 1.2 (see Fig. 2), we take the brand social media posts that were associated with our refined topic clusters and associate them with existing brand personality dictionaries. First, we pulled the trait norms from Aaker's (1997) brand personality dictionary, and we combined these traits norms with the synonyms from the Opoku et al. (2006) brand personality dictionary. After this, we combined the terms from this dictionary with our method of analyzing refined topic clusters to label posts with one of the brand personalities classes. We then compared the vector distances for the words in each cluster and the words in each of the brand personality dimensions from the two dictionaries. Our last step was to take the social media posts that were in the topics closest to the brand personality dictionary words via vector distances and classified all those posts as the brand dimension that was matched.

Therefore, we were able to label social media posts in specific clusters that matched with word embeddings drawn from previously published brand personality dictionaries. As a note, we only classified posts that had 50% or more alignment with any one of the dimensions. For those that did not achieve this threshold, we classified them as NULL posts. This is common within machine learning workflows so that the model can learn from training data points that are good examples of "none of the above." Upon completion of this phase, we achieved the labeling of all the social media posts from different brand accounts into one of the brand personalities dimensions or the NULL category. The percentage of posts that were classified as NULL was 71.1%, which was not within an acceptable threshold for class imbalance problems (Lematre, Nogueira, & Aridas, 2017). We, therefore, needed an additional self-learning technique to increase the training set label size.

Validation of the LDA2Vec Labeling Approach

Since our hybrid approach of combining a dictionary-based approach with an unsupervised machine learning with word embeddings approach is seemingly the first of its kind, we first tested the reliability of the method. To test the generalizability of our approach across different personality dimensions, we considered three alternative labeling techniques and monitored the detection results. To do this, we compared the LDA2Vec results with Aaker's (1997) brand personality dictionary and human coder evaluations for the sample posts of each brand as described below.

We randomly picked 50 posts for each personality dimension based on LDA2Vec results, which totaled 250 posts to be validated. In addition, we selected 50 more posts that seemed likely to be classified with one of the personality dimensions. Note that these additional 50 posts were not initially labeled by the automated process.

The second step was the development of a training document for human coders and a coding scheme to classify tweets into personality dimensions. We followed Morris' (1994) methodology to classify content based on a coding scheme to ensure replicable results. We defined single posts as the unit of analysis because they could be objectively recognized by coders without losing contextual information (Harwood & Garry, 2003).

Next, two research assistants from the first and second authors' institution were trained based on the theoretical foundations (Aaker, 1997) and the comprehensive trait norms dictionary (Opoku et al., 2006). The research assistants coded the posts for each of the five personality dimensions. Several iterative practice sessions were conducted with Twitter and Facebook data sub-samples to train the coders with the content. These sub-sample posts were only used for the training of human coders and were eventually excluded from the final dataset. When coding this subsample, we observed an inter-coder reliability score of 0.89, which is greater than the threshold recommended by Krippendorff (2012). We then had each research assistant code 300 social media posts according to the brand personality dimension that they felt each social media post most exemplified. Finally, we took only the agreed-upon coded posts by both research assistants and compared them against the dictionary-based approach and our hybrid approach. A snapshot of the Spearman correlations for this reliability check is presented in Table 1.

As we expected, when we use human coders' ratings as a benchmark, Aaker's keyword-based dictionary (closed-vocabulary approach) was not comprehensive enough to assess the personality from the social media posts. Our LDA2Vec method was almost twice as more accurate than the dictionary-based approach compared to human coders. In essence, we captured every brand personality dimension that a dictionary-based approach would have detected, but we also detected brand personality in posts that a dictionary-based approach would not have detected.

Table 1
Spearman correlations of three labeling methods (n = 300).

	LDA2Vec	Aaker's dictionary (1997)	Human coder
LDA2Vec	1		
Aaker's dictionary (1997)	0.51	1	
Human Coder	0.84	0.46	1

Multi-Label Self-Supervised Learning

Phase 1's methodology would have been sufficient if our aim was to apply only one label to each social media post (e.g., most exemplifying sincerity or most exemplifying competence). However, our documents are multi-labeled meaning that a document can have multiple labels. This task is common in genre classification algorithms where a movie can have multiple genres (e.g., one movie being in a romantic comedy genre as well as in an action genre). Similarly, a brand post can have multiple labels (e.g., sincerity, excitement, and sophistication with different weights). Thus, our LDA + Word Embeddings model is used to initially label our highest confidence posts in the first phase that are weighted heavily towards one dimension of brand personality. The remaining 71.1% NULL documents still contained personality traits but had a more complex and nuanced multi-label make-up. In order to label these NULL documents, we used document embeddings (Doc2Vec) and a nearest neighbors approach to further label them.

Doc2Vec was proposed by Le and Mikolov (2014) as a simple extension to Word2Vec to extend the learning of embeddings from just words (word embeddings) to word sequences (document embeddings). In document embeddings, the relationships between words within a document are retained in the subsequent document embeddings vectors that are generated. In our case, every single tweet or Facebook post served as a document. We adopted a Doc2Vec and Nearest Neighbors (NN) approach to assign the most accurate labels for our complex social media post. NN's primary objective is to join our NULL documents with their previously labeled nearest neighbors within n-dimensional Euclidean proximity space. This n-dimensional space is created by the 100-dimensional vectors using Doc2Vec on our social media posts. We used average proximity distance within each topic cluster obtained in the first phase and used that metric to identify nearest neighbors of our previously labeled documents.

Fig. 3 illustrates a two-dimensional approximation (since we cannot easily visualize 100 dimensions) of the positionings of some documents. Blue dots are previously labeled social media posts and red dots are NULL labeled posts. Our method detects the nearest labeled posts to the NULL post and applies labels from the nearest labeled posts to the NULL post; thus, this takes a previous unlabeled NULL post and converts it to a multi-labeled post. Through this process we were able to train 40.2% more social media posts, thus resulting in 69.1% of the posts being labeled with brand personality dimensions and 30.9% remaining as null. This resulted in an even more acceptable threshold than phase 1's results with regard to class imbalance problems (Lematre et al., 2017). As a clarification, the fact that we have 30.9% of the posts labeled as null does not mean that our prediction model will perform in a way where 30.9% of the posts will be labeled as null. Rather, it is an indication that our automated labeling methodology was not able to label 30.9% of

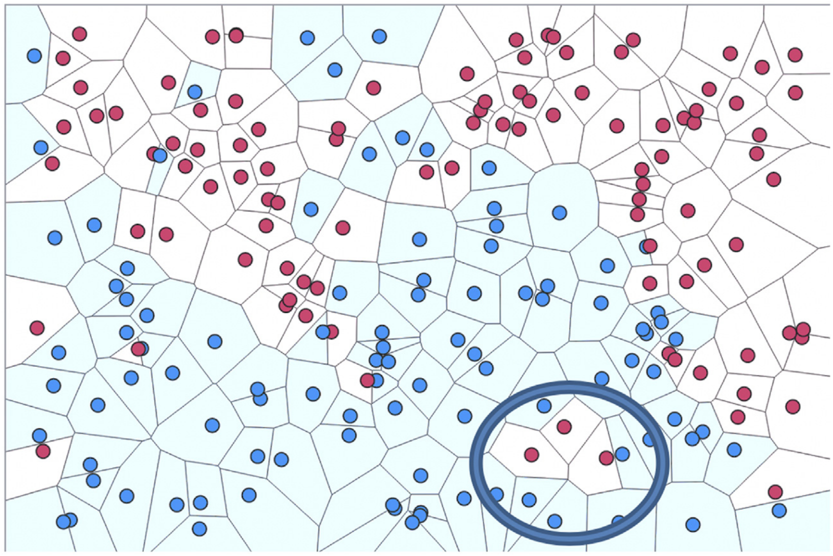

Fig. 3. Visual representation of labeling NULL documents with Doc2Vec and k-NN.

the posts, thus they are used by machine-learned modeling to indicate weak brand personality signaling.

Machine-Learned Model Training

Once we had our automatically labeled training data, we focused on testing various machine learning and deep learning algorithms to build a brand personality machine-learned prediction model.

First, we cleaned the labeled posts by applying pre-processing to remove stop words, stemming, and punctuation. We then transformed the labeled posts into a computational format by using the scikit-learn machine learning package for the Python programming language (Han, Kamber, & Pei, 2012; Pedregosa et al., 2012). We then conducted feature extraction to transform unstructured text data into numerical vectors for computational processing. This process takes these sets of terms and transforms them into numerical feature vectors.

Our next step was to train the actual machine learning classification model. The goal of this step is to select the best classification algorithm(s) for our analyses, keeping in mind that our main priority is to minimize classification error and that our context is one where there are multiple classes for detection (five brand personality classes). We examined four major classification approaches and used the best performing algorithm in each type of method. The evaluation metric we selected for our task was precision instead of recall or accuracy in order to compare binary class values for each personality trait. The precision levels of each classification method are shown in Table 2. Our method for testing the accuracy of each classifier was by using 10-fold cross-validation with the pre-labeled training set.

Although these results were somewhat promising, we suspected that our models were underperforming due to the fact that our machine learning models were missing the contextual understanding of words. To remedy this, we used a recent open-source collection of Transformer-based models (Wolf et al., 2019) that were developed from the transformer work on BERT (Devlin, Chang, Lee, & Toutanova, 2018). Transformers were originally introduced via researchers at Google Brain (Vaswani et al., 2017) and they help computers understand words in the context of other words within sentences, such as with the sentence, “John went to the store because he knew it would be open at this time.” Transformers help the computer understand that “it” is referring to the store rather than John. Subsequent work in transformer technology includes models such as BERT (Devlin et al., 2018), RoBERTa (Liu et al., 2019), ALBERT (Lan et al., 2019), and DistilBERT (Sanh, Debut, Chaumond, & Wolf, 2019). For our brand personality task, we tried all three BERT derived models and found that the best performing model was RoBERTa (Robustly optimized BERT approach). Substantially improved precision results are shown in Table 2.5.

RoBERTa was developed at Facebook and is a retraining of BERT with improved training methodology. Compared to BERT, they used 1,000% more data (160 GB of text for pre-training, including the 16 GB document corpus used in BERT)

Table 2
Comparison of classification algorithms (precision with 10-fold cross validation).

Method	Sincerity	Competence	Excitement	Sophistication	Ruggedness
Random forest	0.613	0.581	0.618	0.641	0.495
Support vector machines	0.582	0.543	0.591	0.594	0.447
K-nearest neighbors	0.512	0.498	0.505	0.56	0.423
Naïve Bayes	0.504	0.467	0.461	0.501	0.425
Seq2Seq LSTM	0.496	0.440	0.452	0.487	0.483

Table 2.5
Comparison of classification algorithms including RoBERTa (precision with 10-fold cross validation).

Method	Sincerity	Competence	Excitement	Sophistication	Ruggedness
Random forest	0.613	0.581	0.618	0.641	0.495
Support vector machines	0.582	0.543	0.591	0.594	0.447
K-nearest neighbors	0.512	0.498	0.505	0.56	0.423
Naïve Bayes	0.504	0.467	0.461	0.501	0.425
Seq2Seq LSTM	0.496	0.440	0.452	0.487	0.483
RoBERTa-based multi-label classifier	0.796	0.793	0.824	0.857	0.723

and compute power in training RoBERTa. As an additional improvement to BERT's training methodology, RoBERTa eliminates the Next Sentence Prediction (NSP) task and offers dynamic masking so that the masked token changes during the training process. Since our aim is to conduct multi-label classification within our brand personality detection tool, RoBERTa better serves our needs (Further information and statistical results of RoBERTa can be found at Liu et al., 2019).

To clarify as to why we incorporate a mixture of LDA2Vec with RoBERTa is to set a foundation for the reality that social data will continue to reflect the evolution of consumer perception. These constant changes within new social data will create the need for future model re-training. If we chose to solely use LDA2Vec with human integration to group the personality clusters of incoming posts, our LDA2Vec parameters would remain static from our first initial training rather than adjusting to the new social data. We would need to recruit additional human effort to redefine the boundaries of each cluster. On the other hand, by drawing on deep neural network architecture such as RoBERTa, we can just fine-tune the last layer of the machine-learned neural network whenever we receive new data and update training parameters without any human intervention. BERT-based learning models, such as RoBERTa, are built to adjust with dynamic context as compared to static embedding vectors created by LDA2Vec.

Our resulting machine-learned brand personality prediction model has the ability to take in a social media post and calculate scores for each of the five brand personality dimensions. Researchers can also then take the individual results of each social media post (e.g., each post of a brand's Twitter timeline) and take the average (mean) of the aggregated brand personality scores to come up with a total brand personality score for the grouping of posts (e.g., the brand personality score for a brand's Twitter timeline).

Although posts can be analyzed individually by our model, we recommend employing aggregated results to ensure context congruency. It is still possible to use our model to analyze individual posts, but as an example, we have found that individual post scores for competence show to be more accurate than individual post scores for sincerity. One reason for this could be due to the fact that sincerity is more difficult to detect without a greater degree of context, especially in the case of sarcasm. Furthermore, even aggregated textual posts will need more context from other brand medium (e.g., images, videos, audio), but text is the current starting point due to computing infrastructure and power required to build and run models with images, video, and audio. We do recognize though that analyzing individual posts may have the broadest use case for researchers and practitioners, thus our model does allow for this in its current implementation.

Findings

Since we achieved relatively sufficient precision for our brand personality model, we applied our trained model to predict brand personalities of 20 separate brands' Twitter timelines (brands that were not in the set of 50 brands we used to train the model). These 20 brands for testing were chosen based on their publicly perceived visible personalities (Dvornechuck, 2020). Our analyses were run in April 2020; thus, one concern was that many Twitter feeds of brands may have been full of COVID-19 messages, thus potentially preventing an “apples to apples” comparison between brands. Thus, we made sure that all the brand Twitter feeds that we were analyzing did not have more than a few tweets regarding COVID-19 and we pulled the maximum 3,200 tweets per brand that the Twitter public API allows for us to pull. The overall findings of the brands' channeled personality results are presented in Table 3. The numbers corresponding to brands represent the multi-label density percentage (probabilities) of each dimension. As a note, the dimensions of sophistication and ruggedness show low confidence scores overall, which suggests that the brands that we sampled use words associated with those two dimensions far less than with the other three dimensions. This trend is likely to hold as true across other brands that future researchers will run through our brand personality prediction model.

We can see that our model results correspond quite well with intuition. For example, major technology companies rank high in competence, Red Bull ranks high in excitement, outdoor brands rank high in ruggedness, Timberland and Patagonia rank high in sincerity, and luxury brands rank high in sophistication. Some less intuitive, but explainable, examples are Google ranking high in ruggedness and McDonald's ranking high in sincerity. Google tweets quite frequently about their impact all around the world, especially in more remote areas, which could be why they rank so highly in ruggedness. Additionally, according to a report from Sprinklr, a major social media analytics company, McDonalds makes honesty a major priority in their social media postings (Walter, 2014).

Perceived vs. Channeled Brand Personality

As discussed in the theoretical background, our brand personality model also allows for the measurement of

Table 3
Channeled brand personality results for the test set (20 brands).

Brands	Competence	Excitement	Ruggedness	Sincerity	Sophistication
Google	**0.242**	0.077	0.063	0.041	0.004
Patagonia	**0.161**	0.131	**0.078**	**0.325**	0.019
Microsoft	**0.142**	0.218	0.042	0.080	0.017
Intel	**0.142**	0.215	0.048	0.143	0.022
Nike	**0.142**	0.148	0.030	0.118	0.008
Volvo Cars	0.133	0.141	0.023	0.240	0.026
Dove	0.119	0.093	0.036	0.208	**0.066**
Louis Vuitton	0.100	**0.385**	0.037	0.090	**0.066**
Cabelas	0.098	0.117	**0.085**	0.272	0.026
Bass Pro Shops	0.086	0.111	0.062	0.218	0.026
Disney	0.081	0.200	0.030	0.104	0.040
Coca Cola	0.079	0.059	0.008	0.164	0.006
Harley Davidson	0.076	0.157	0.051	0.124	0.019
Toms	0.066	0.123	0.049	0.132	0.019
T-Mobile	0.060	**0.339**	0.009	0.062	0.005
Timberland	0.060	0.16	0.033	**0.398**	0.016
Burberry	0.057	**0.372**	0.017	0.087	**0.100**
Victoria's Secret	0.056	0.103	0.016	0.217	0.040
McDonalds	0.036	0.124	0.015	**0.321**	0.009
Red Bull	0.029	0.106	**0.063**	0.066	0.008

Top performers in each dimension are bolded.

consumers' perceived personalities of brands. On one hand, previous theoretical work on brand personality formation suggests that consumer-perceived and employee-perceived brand personality have more predictive power than intended personality (e.g. official social media account announcements) in brand personality formation (Xu et al., 2016). On the other hand, properly and consistently channeled brand personality has been shown to have a significant effect on audience perception (Parker, 2009). Parker (2009) considers this congruence between intended and perceived brand personality through the lens of congruity theory. Table 4 provides examples of intended personality dimensions from official brand Twitter accounts and perceived personality from Twitter user accounts.

In Table 5, we display the full results of the perceived brand personality scores for all of our brands in which we ran our model against conversation on Twitter mentioning each brand (using the @brandaccountname keyword; Intel is missing due to a data ingest issue with @intel and the Twitter API at the time of writing this paper). As a note, we have substantially redacted random portions of the perceived personality posts to respect the privacy of these individuals and to prevent easy searching of these posts on various social media analytics platforms (Humphreys & Wang, 2018). Interestingly, technology companies do not rate very high in competence, but rather brands such as Patagonia, Dove, and Volvo. This could signal that although tech companies may speak with competence, consumers do not necessarily perceive them in terms of competence. Victoria's Secret ranks highest in excitement, which seems in line with how one would speak about lingerie. Outdoor brands, such as Bass Pro Shops and Patagonia, rank high in ruggedness, which follow the trend of our channeled personality results. Regarding sophistication, Dove ranks high in perceived personality, which matches its results for channeled personality. A notable difference in sincerity is with McDonald's of which they rank high in sincerity for their channeled personality, but low in sincerity with regard to perceived personality.

To illustrate one way that researchers and practitioners can easily assess the congruence (or incongruence) of their channeled to perceived personalities, we illustrate calculations using cosine similarity, a metric used in previous personality

Table 4
Channeled and perceived brand personality examples.

Brand personality dimensions	Channeled personality	Perceived personality
Competence	“In partnership with journalists and fact checkers, the Google News Initiative supported company, http://Debunk.eu uses AI to help identify disinformation and reduce its harmful impact on society.”	“[redacted names] @GoogleAI All IC designers may hope so, but this is @Google. TPU v3 Pod has 100+ petaflops of compute power. Trained BERT models in just over an hour. [redacted portion] reinforcement learning approaches that were successful in solving games, like AlphaGo”
Sophistication	“Elegant and seductive lips with sheer understated eyes – the @Burberry siren red runway look:”	“@Burberry [redacted portion, talking about a product] is so stylish and classy”

Table 5
Perceived brand personality results for the test set.

Brands	Competence	Excitement	Ruggedness	Sincerity	Sophistication
Patagonia	**0.555**	0.059	**0.065**	0.034	0.005
Dove	**0.210**	0.215	0.030	**0.145**	**0.182**
Volvo Cars	**0.204**	0.216	0.052	0.096	0.043
Nike	0.132	0.070	0.025	**0.126**	0.016
Burberry	0.129	0.147	**0.053**	0.081	0.048
Victoria's Secret	0.091	**0.380**	0.017	0.071	**0.149**
Google	0.070	0.113	0.052	0.074	0.010
Harley Davidson	0.067	0.157	0.030	0.100	0.022
Timberland	0.065	0.120	0.048	0.091	0.016
Toms	0.062	**0.237**	0.038	**0.116**	0.035
Coca Cola	0.062	0.058	0.024	0.067	0.010
T-Mobile	0.057	0.096	0.026	0.068	0.016
Microsoft	0.050	0.052	0.021	0.041	0.008
Bass Pro Shops	0.046	0.05	**0.070**	0.074	0.007
Red Bull	0.035	0.089	0.042	0.103	**0.052**
Louis Vuitton	0.024	**0.361**	0.008	0.056	0.032
McDonalds	0.022	0.033	0.009	0.047	0.003
Disney	0.020	0.055	0.009	0.085	0.016
Cabelas	0.012	0.005	0.009	0.016	0.001

Top performers in each dimension are bolded; Intel is excluded from this table due to errors in ingesting @intel conversation data via the Twitter API.

similarity studies (Netzer et al., 2012; Yun, Pamuksuz, & Duff, 2019). Cosine similarity takes two vectors and calculates the cosine of the angle between the two vectors as in the following equation where θ is the angle between vectors, the numerator is the dot product of the two vectors, and the denominator is the product of the vector lengths. The results are shown in Table 6.

Therefore, social media conversations that mention @HarleyDavidson exhibit a 99.2% similarity in the five brand personality dimension levels as the brand communication coming from Harley Davidson themselves. This may be an indication to Harley Davidson's brand management team that they are successfully conveying their brand personality, although they are not necessarily a top performer in any of the dimensions (see Table 3). Despite the incongruence with sincerity, McDonald's similarity score (0.934) indicates that it is relatively congruent with consumer perception. Patagonia's brand personality is notably incongruent with customer perception (0.511), which may merit further investigation and strategic intervention. Strategically, one should note that this provides managers an opportunity to either hone its communication of existing personality or take the personality where consumers may be leading it, depending on the strategic organization of the firm (see e.g., Carpenter & Humphreys, 2019). This demonstration shows just one way that this automated brand personality detection tool could be used by researchers and practitioners. Although not the focus of our research, the tool could be further used on an individual-post level to later examine the effects of congruence between customers and the firm.

Discussion

An enormous amount of textual data is now available online. It is both the outcome and input to consumer perceptions about brands. We show that this conversation can be measured and systematically used to provide real time measures of brand personality. For firms, this means that they can track alignment of marketing messages—which are now more numerous and varied than in previous mass media environments (Humphreys, 2016)—with existing brand personality. They can also track consumer response to brand personality and adapt to changes in the cultural and socio-political environment, finding ways to tailor the message in ways that maintain consistency and relate to consumer values. Finally, firms can create benchmarks to compare their brand personality to competitors to ensure that they remain distinct, particularly in an environment that may encourage convergence due to echo chambers and other social effects (Hewett, Rand, Rust, & Van Heerde, 2016; Moe & Schweidel, 2014).

Table 6
Congruence between perceived and channeled brand personality.

Brand	Similarity distance
Harley Davidson	0.992
Louis Vuitton	0.976
Microsoft	0.958
Nike	0.951
Cabelas	0.945
McDonald's	0.934
Toms	0.929
Bass Pro Shops	0.918
Coca Cola	0.918
Red Bull	0.912
Disney	0.872
Dove	0.864
T-Mobile	0.862
Burberry	0.838
Volvo Cars	0.836
Timberland	0.820
Google	0.750
Victoria's Secret	0.620
Patagonia	0.511

Implications for Researchers

For researchers, this enables greater exploration of brand personality. For instance, future studies may expand the brand personality model, as proposed by Aaker (1997), to other dimensions. There have already been discussions about the limitations and generalizability of Aaker's work (Austin, Siguaw, & Mattila, 2003). Since our methodology of building a brand personality detection model is not limited by how many dimensions or by specific keywords, scholars may test their future propositions using our methodology.

Our findings also allow for the consideration of human personality dimensions and brand personality dimensions within the realm of social media. Human personality dimensions and brand personality dimensions are treated distinctly in literature. Fournier's (1998) study on consumers and their relationships with brands provides evidence that there are indeed enduring relationships between brands and consumers and there may be beneficial “matches” between certain dimensions of human and brand personality. Clearly, the line between human personality and brand personality is challenged.

There are further implications for increasing online engagement based on matches between brand and consumer personality. Based on self-congruence theory, engaging with like-personality social media content may reduce cognitive dissonance for consumers who share personality traits with a brand. As some research has shown, consumers also expect different types of relationships from different types of brands (e.g., communal vs. transactional; Aggarwal, 2004). Based on these norms, brand personality can be measured with the expectation that congruencies will reduce cognitive dissonance and thereby increase engagement while incongruencies will increase cognitive dissonance and thereby reduce engagement. Researchers may want to examine interactions on the individual post level to further test predictions about brand-consumer self-congruence. We have already discussed that our model can allow for individual post analysis with the caveat that it may not be as accurate as aggregate post analysis across dimensions, but researchers could potentially engage in a time series analysis of individual post congruence with historical aggregated brand personality versus engagement over a period of time. In the case of individual post analysis, receiving output from the model of a weak or even null (zero) signal could still provide insight into why individual posts are deviating so strongly from the historical data that was used to train the machine learning model.

Finally, there are potential implications of intended and perceived personality congruence. Organizational behavior and marketing researchers interested in firm reputation may, for example, examine the congruence between a brand's channeled brand personality and its own CEO brand personality and analyze the impacts of possible incongruence on firm practices. Furthermore, scholars may look at longitudinal data across multiple years and investigate the effects of strategic changes within firms (e.g., CEO turnover) or crises on brand personality over time.

Implications for Practitioners

There is a potential for practitioners to benefit from using the implementation of our brand personality model. Most directly, marketing managers can monitor how efficiently their brand's personality is being channeled through social media. Since branding strategies can be improved through observed personality and consumer engagement, congruence between channeled and perceived brand personality can also serve as an important metric to evaluate branding strategies. Firms can equally measure shifts in brand personality overtime and measure their personality against actual or potential competitors in the marketplace.

As a tool for evaluating potential changes to brand personality, firms can use the tool to monitor consumer perceptions of the brand and make adjustments accordingly. For instance, the impact of dissident stakeholder perceptions on brand personality or shifts in brand perceptions due to changes in social or political context (e.g., COVID) can be investigated (e.g., employees, activists). Managers can even consider reaching out to previously unidentified customer segments by using the personality similarity of the brand and social media users.

Finally, the proposed method could be used to identify potential celebrity spokespersons or influencers for brand marketing. Since celebrities are considered to have their own personal brands, measuring the cosine similarity between two may lead practitioners to have greater insight in choosing the most suitable endorser for their brands. Yun et al. (2019) found evidence that people on Twitter follow brand Twitter accounts (e.g., Harley Davidson) that are closer to them in human personality (big five) than brand Twitter accounts that are further from them in human personality. With our tool, practitioners can use our model to measure the brand personality of a celebrity (e.g., Kim Kardashian) and use cosine similarity (see methodology in Yun et al., 2019) to see how close Kim Kardashian's brand personality is to the brand personality of the brand being analyzed. If the celebrity is very close to the brand personality of the brand being analyzed, this could be a good indication that this celebrity could be hired as an influencer for the brand.

Implementation of Our Model for Testing

To encourage testing of our model, especially for those without computer programming backgrounds, we have integrated our model into the open-source project named The Social Media Macroscope (Yun et al., 2019). Our model implementation for testing channeled brand personality can be accessed via the Brand Analytics Environment (BAE) tool (Yun et al., 2019), and our model implementation for testing perceived brand personality can be accessed via the Social Media Intelligence and Learning Environment (SMILE) tool (Yun et al., 2019), all within the The Social Media Macroscope (Yun et al., 2019). BAE allows users to input Twitter handles so that our brand personality model can be run against them and the brand personality scores will be outputted when analyses

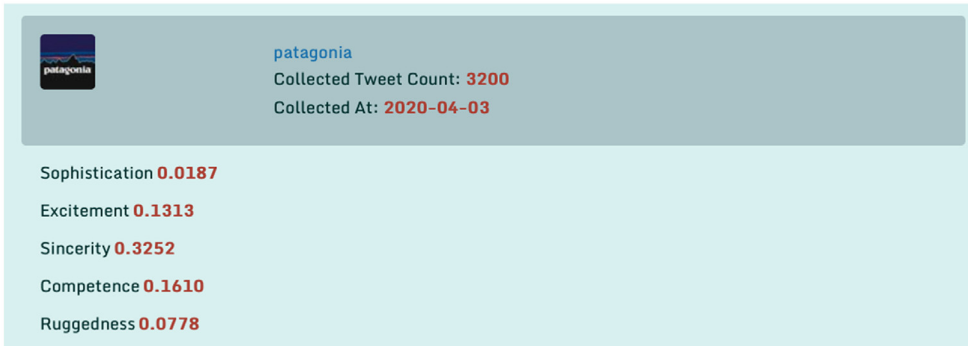

Fig. 4. Example of brand personality output within BAE.

have completed. Example screenshots for Patagonia within BAE and SMILE are shown in Figs. 4 and 5 respectively.

Both implementations of our model in the tool (channeled and perceived brand personality measurement) allow for the analysis of brand personality scores per individual post as well as in aggregate averaged across posts.

Limitations

We point out several limitations to this work. First, our work is limited to analyzing brands that maintain a Twitter and Facebook presence. Brands that are not active in their Twitter or Facebook postings would not be the best candidates for utilizing our automated method. Second, adding more data to our model would most likely bring additional improvement to brand personality detection. While we incorporate a wide range of brands from different industries, clusters may provide better results with more diverse contexts (more brands, more posts, and more social media platforms). Third, we rely only on text data provided by brands. Future research may also focus on images to identify the real context of the posts instead of solely relying on the usage of words. We believe a combination of text, image, and other media type processing would enable researchers to extract more robust and accurate information from both individual and aggregated posts, although access to computing power and infrastructure to accomplish this at scale

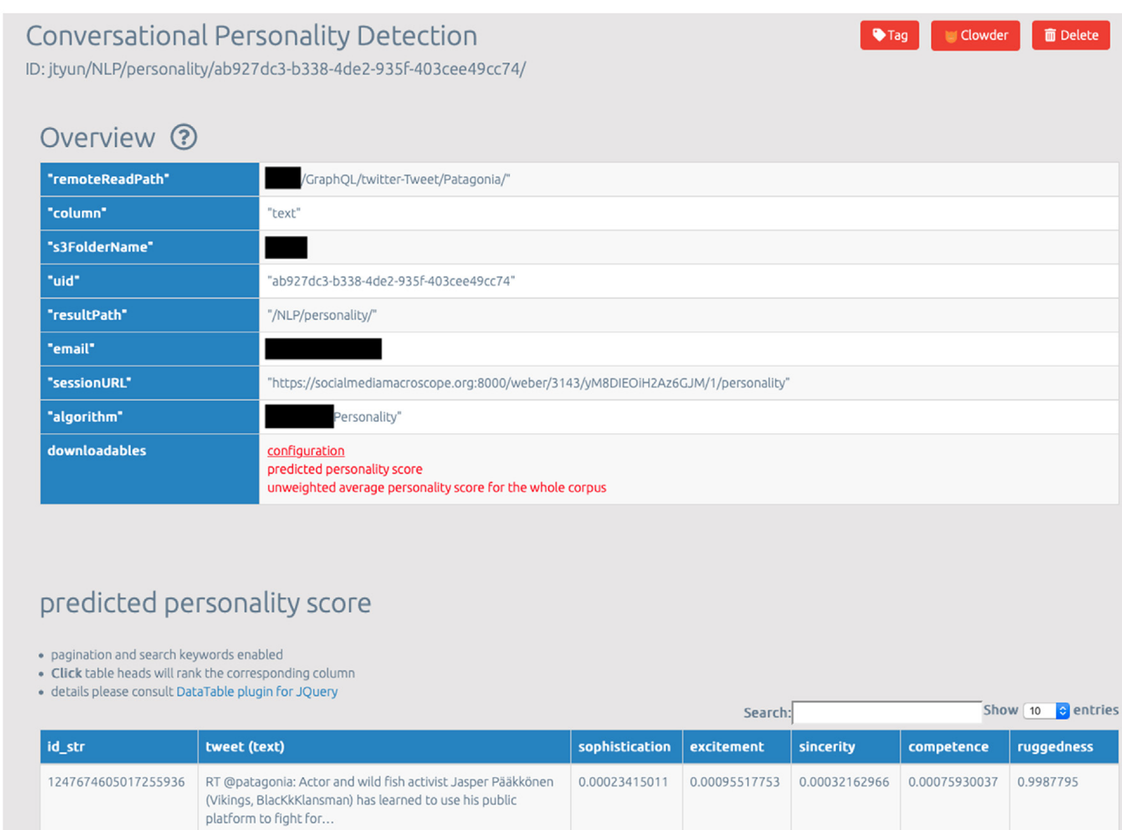

* Author's names blacked out

Fig. 5. Example of brand personality screenshot within SMILE.

is still cost prohibitive to many researchers. We believe that demonstrating the potential of brand personality detection—in our case initially through text—will provide a foundation and motivation for the analysis of other unstructured data types as data modeling and infrastructure costs go down over time. With this first of a kind implementation, we hope that the method introduced in this paper provides a useful tool for researchers and practitioners interested in automatically monitoring brand personalities.

Conclusion

Brand personality is a critical resource for the firm that binds together many different stakeholders—consumers, employees, shareholders, and management. It helps delineate the character of a firm and its products and facilitates the emotional connection between a brand and its target audience. Today's firms are challenged to efficiently define, manage, and control their own brand personality to succeed in a competitive advantage over competitors (Madden, 2006). Brand personality can be an influential tool to induce emotions, build trust and loyalty (Fournier, 1998), and enhance consumer preference (Aaker, 1997). Social media is a natural platform in which brand personality is discussed—both by consumers and by brands—and therefore a natural arena in which to measure, compare, and track brand perceptions. To the best of our knowledge, our proposed approach is pioneering. It provides several opportunities for both researchers and practitioners especially in generating personality assessments regularly, easily, and efficiently for brands of interest.

References

Aaker, D. (1995). Building strong brands. *Brandweek, 36*.

Aaker, J. L. (1997). Dimensions of brand personality. *Journal of Marketing Research, 34*, 347–356.

Aggarwal, P. (2004). The effects of brand relationship norms on consumer attitudes and behavior. *Journal of Consumer Research, 31*(1), 87–101.

Arnold, S. J., Kozinets, R. V., & Handelman, J. M. (2001). Hometown ideology and retailer legitimation: the institutional semiotics of Wal-Mart flyers. *Journal of Retailing, 77*(2), 243–271.

Austin, J. R., Siguaw, J. A., & Mattila, A. S. (2003). A re-examination of the generalizability of the Aaker brand personality measurement framework. *Journal of Strategic Marketing, 11*, 77–92.

Baevski, A., Schneider, S., & Auli, M. (2019). Vq-Wav2vec: self-supervised learning of discrete speech representations. *arXiv preprint*, 11–12 arXiv:1910.05453.

Batra, R., Ahuvia, A., & Bagozzi, R. P. (2012). Brand love. *Journal of Marketing, 76*, 1–16.

Berger, J., Humphreys, A., Ludwig, S., Moe, W. W., Netzer, O., & Schweidel, D. A. (2020). Uniting the tribes: using text for marketing insight. *Journal of Marketing, 84*(1), 1–25.

Bhagwat, Y., Warren, N. L., Beck, J. T., & Watson IV, G. F. (2020). Corporate sociopolitical activism and firm value. *Journal of Marketing, 84*(5), 1–21.

Blei, D. M., Ng, A. Y., & Jordan, M. I. (2003). Latent Dirichlet allocation. *Journal of Machine Learning Research, 3*, 993–1022.

Carpenter, G., & Humphreys, A. (2019). What the wine industry understands about connecting with consumers. *Harvard Business Review* https://hbr.org/2019/03/what-the-u-s-wine-industry-understands-about-connecting-with-customers.

Carr, S. D. (1996). The cult of brand personality. *Marketing News, 30*, 4–9.

Culotta, A., & Cutler, J. (2016). Mining brand perceptions from Twitter social networks. *Marketing Science, 35*, 343–362 mksc.2015.0968.

Das, R., Zaheer, M., & Dyer, C. (2015). Gaussian Lda for topic models with word embeddings. *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, 795–804.

Devlin, J., Chang, M.-W., Lee, K., & Toutanova, K. (2018). Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint* arXiv:1810.04805.

Diamond, N., Sherry, J. F., Muñiz, A. M., McGrath, M. A., Kozinets, R. V., & Borghini, S. (2009). American girl and the brand gestalt: closing the loop on sociocultural branding research. *Journal of Marketing, 73*, 118–134.

Dvornechuck, A. (2020). Brand personality: Traits of top brands. *Ebaqdesign Blog*.

Fournier, S. (1998). Consumers and their brands: Developing relationship theory in consumer research. *Journal of Consumer Research, 24*, 343–353.

Golbeck, J., Robles, C., Edmondson, M., & Turner, K. (2011). Predicting personality from Twitter. *Proceedings - 2011 IEEE International Conference on Privacy, Security, Risk and Trust and IEEE International Conference on Social Computing, PASSAT/SocialCom 2011*, 149–156.

Han, J., Kamber, M., & Pei, J. (2012). *Data Mining: Concepts and Techniques*, , 745. Morgan Kaufmann Publishers.

Harris, F., & de Chernatony, L. (2001). Corporate branding and corporate brand performance. *European Journal of Marketing, 35*, 441–456.

Harwood, T. G., & Garry, T. (2003). An overview of content analysis. *The Marketing Review, 3*, 479–498.

Hewett, K., Rand, W., Rust, R. T., & Van Heerde, H. J. (2016). Brand buzz in the Echoverse. *Journal of Marketing, 80*(3), 1–24.

Holt, D. (2016). Branding in the age of social media. *Harvard Business Review, 94*(3), 40–50.

Huang, D., & Luo, L. (2016). Consumer preference elicitation of complex products using fuzzy support vector machine active learning. *Marketing Science, 35*, 445–464.

Humphreys, A. (2016). *Social Media: Enduring Principles*. Oxford University Press.

Humphreys, A., & Wang, R. J.-H. (2018). Automated text analysis for consumer research. *Journal of Consumer Research, 44*(6), 1274–1306.

Jacobs, B. J. D., Donkers, B., & Fok, D. (2016). Model-based purchase predictions for large assortments. *Marketing Science, 35*, 389–404.

Keller, K. L., & Lehmann, D. R. (2006). Brands and branding: Research findings and future priorities. *Marketing Science, 25*, 740–759.

Kern, M. L., Park, G., Eichstaedt, J. C., Schwartz, H. A., Sap, M., & Smith, L. K., et al (2016). Gaining insights from social media language: Methodologies and challenges. *Psychological Methods, 21*, 507–525.

Krippendorff, K. (2012). Content analysis: An introduction to its methodology. *Language Arts & Disciplines, 79*, 441.

Lan, Z., Chen, M., Goodman, S., Gimpel, K., Sharma, P., & Soricut, R. (2019). Albert: A lite bert for self-supervised learning of language representations. *arXiv preprint* arXiv:1909.11942.

Le, Q., & Mikolov, T. (2014). Distributed representations of sentences and documents. *31st International Conference on Machine Learning, Beijing, China*, 1188–1196.

Lee, T. Y., & Bradlow, T. (2011). Automated marketing research using online customer reviews. *Journal of Marketing Research, XLVIII*, 881–894.

Lematre, G., Nogueira, F., & Aridas, C. K. (2017). Imbalanced-learn: A Python toolbox to tackle the curse of imbalanced datasets in machine learning. *Journal of Machine Learning Research, 18*, 1–5.

Levy, S. J. (1985). Dreams, fairy tales, animals, and cars. *Psychology and Marketing*, 67–81.

Liu, Y., Ott, M., Goyal, N., Du, J., Joshi, M., & Chen, D., et al (2019). Roberta: A robustly optimized bert pretraining approach. *arXiv preprint* arXiv:1907.11692.

Madden, T. J. (2006). Brands matter: an empirical demonstration of the creation of shareholder value through branding. *Journal of the Academy of Marketing Science, 34*, 224–235.

Martineau, P. (1958). The personality of the retail store. *Harvard Business Review, 36*, 47–55.

Mikolov, T., Chen, K., Corrado, G., & Dean, J. (2013). Efficient estimation of word representations in vector space. *Arxiv*, 1–12.

Moe, W. W., & Schweidel, D. A. (2014). *Social Media Intelligence*. Cambridge University Press.

Moody, C. E. (2016). *Mixing Dirichlet Topic Models and Word Embeddings to Make Lda2vec*.

Morris, R. (1994). Computerized content analysis in management research: A demonstration of advantages & limitations. *Journal of Management, 20*, 903–931.

Netzer, O., Feldman, R., Goldenberg, J., & Fresko, M. (2012). Mine your own business: Market-structure surveillance through text mining. *Marketing Science, 31*, 521–543.

Opoku, R., Abratt, R., & Pitt, L. (2006). Communicating brand personality: Are the websites doing the talking for the top South African business schools? *Journal of Brand Management, 14*, 20–39.

Pamuksuz, U. (2017). *Unlocking the Insights: Data-Analytic Examinations of social signals in Organizational Practices*. Dissertation. Business Administration . University of Illinois at Urbana-Champaign.

Park, G., Schwartz, H. A., Eichsteadt, J. C., Kern, M. L., Kosinski, M., & Stillwell, D. J., et al (2015). Automatic personality assessment through social media language. *Journal of Personality and Social Psychology, 108*, 1–25.

Parker, B. T. (2009). A comparison of brand personality and brand user-imagery congruence. *Journal of Consumer Marketing, 26*, 175–184.

Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., & Grisel, O., et al (2012). Scikit-learn: machine learning in Python. *Journal of Machine Learning Research, 12*, 2825–2830.

Peffers, K., Tuunanen, T., Rothenberger, M. A., & Chatterjee, S. (2008). A design science research methodology for information systems research. *Journal of Management Information Systems, 24*, 45–78.

Pennington, J., Socher, R., & Manning, C. D. (2014). Glove: global vectors for word representation. *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing*, 1532–1543.

Ranfagni, S., Camiciottoli, B. C., & Faraoni, M. (2016). How to measure alignment in perceptions of brand personality within online communities: Interdisciplinary insights. *Journal of Interactive Marketing, 35*, 70–85.

Ross, S. D. (2008). Assessing the use of the brand personality scale in team sport. *International Journal of Sport Management and Marketing, 3*, 23.

Sanh, V., Debut, L., Chaumond, J., & Wolf, T. (2019). Distilbert, a distilled version of bert: Smaller, faster, cheaper and lighter. *arXiv preprint* arXiv:1910.01108.

Sirgy, M. J. (1982). Self-concept in critical review consumer behavior. *Journal of Consumer Research, 9*, 287–300.

Speer, R. (2017). Conceptnet Numberbatch 17.04: better, less-stereotyped word vectors. *ConceptNet blog, April24*.

Thompson, C. J., Rindfleisch, A., & Arsel, Z. (2006). Emotional branding and the strategic value of the Doppelgänger brand image. *Journal of Marketing, 70*(1), 50–64.

Tirunillai, S., & Tellis, G. J. (2014). Mining marketing meaning from online chatter: Strategic brand analysis of big data using latent Dirichlet allocation. *Journal of Marketing Research (JMR), 51*, 463–479.

Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., & Gomez, A. N., et al (2017). Attention is all you need. *Advances in Neural Information Processing Systems*, 5998–6008.

Walter, E. (2014). *How Mcdonald's Wins in Social: Ranking at Number One | Social Media Today.*.

Wang, A., Pruksachatkun, Y., Nangia, N., Singh, A., Michael, J., & Hill, F., et al (2019). Superglue: A stickier benchmark for general-purpose language understanding systems. *Advances in Neural Information Processing Systems*, 3261–3275.

Wolf, T., Debut, L., Sanh, V., Chaumond, J., Delangue, C., & Moi, A., et al (2019). Transformers: State-of-the-art natural language processing. *arXiv preprint* arXiv:1910.03771.

Xu, A., Liu, H., Gou, L., Akkiraju, R., Mahmud, J., & Sinha, V., et al (2016). Predicting perceived brand personality with social media. *ICWSM*, 436–445.

Yun, J. T., Pamuksuz, U., & Duff, B. R. L. (2019). Are we who we follow? Computationally analyzing human personality and brand following on Twitter. *International Journal of Advertising, 38*(5), 776–795.

Yun, J. T., Nickolas, V., Chen, W., Luigi, M., Joseph, T., & Curtis, D., et al (2019). The Social Media Macroscope: A Science Gateway for Research Using Social Media Data. *Future Generation Computer Systems.*

Zhang, K., Bhattacharyya, S., & Ram, S. (2016). Large-scale network analysis for online social brand advertising. *MIS Quarterly, 40*, 849–868.