

VISUAL LISTENING IN: EXTRACTING BRAND IMAGE PORTRAYED ON SOCIAL MEDIA

Liu Liu
Daria Dzyabura
Natalie Mizik

Working Paper
No 258

NES Working
Paper series

August
2019

Visual Listening In: Extracting Brand Image Portrayed on Social Media

Liu Liu*

Daria Dzyabura†

Natalie Mizik‡

June 6, 2017

This Version: August 24, 2019

Abstract

We propose a “visual listening in” approach (i.e., mining visual content posted by users) to measure how brands are portrayed on social media. Using a deep-learning framework, we develop BrandImageNet, a multi-label convolutional neural network model, to predict the presence of perceptual brand attributes in the images that consumers post online. We validate model performance using human judges, and find a high degree of agreement between our model and human evaluations of images. We apply the BrandImageNet model to brand-related images posted on social media, and compute a brand-portrayal metric based on model predictions, for 56 national brands in the apparel and beverages categories. We find a strong link between brand portrayal in consumer-created images and consumer brand perceptions collected through survey tools. Images are close to surpassing text as the medium of choice for online conversations. They convey rich information about the consumption experience, attitudes, and feelings of the user. We show that valuable insights can be efficiently extracted from consumer-created images. Firms can use the BrandImageNet model to automatically monitor in real time their brand portrayal and better understand consumer brand perceptions and attitudes toward theirs and competitors’ brands.

Key words: social media, visual marketing, brand perceptions, computer vision, machine learning, deep learning, transfer learning, big data

* Leeds School of Business, University of Colorado Boulder, E-mail: liu.liu-1@colorado.edu.

† New Economics School, E-mail: ddzyabura@nes.ru.

‡ Foster School of Business, University of Washington, E-mail: nmizik@uw.edu.

1. Introduction

A profound shift has recently occurred in how people consume information and in the origins of the information itself. Corporations and traditional media have lost their monopoly on information creation and communication channels. Much of the brand-related content is now created and spread through social media platforms (e.g., Twitter and Instagram), user discussion forums, blogs, and so on. Firms now operate in an environment in which customers are actively participating in sharing their brand perceptions and experiences on social media and co-creating brands and brand identities. With this broader, more egalitarian distribution model, monitoring how a brand is portrayed on social media is essential for effective brand management.

In this paper, we focus on consumer-created visual content. With the proliferation of camera phones, generous data plans, and image-based social media platforms, photo-taking and sharing have become an important part of consumers' social lives (Diehl et al. 2016). Images are becoming an increasingly prevalent form of online conversations. Instagram has 1 billion monthly active users, and these users share an average of 95 million photos daily.¹ In these shared photos, consumers often tag brands, resulting in a large volume of images depicting brands. For example, a search of *#nike* on Instagram returns over 92 million photos tagged with Nike.² Through these images, consumers communicate about brands with each other.

Consumer-created brand images on social media are different from product images on retailers' websites. In the images posted on social media, consumers often link brands with usage context, feelings, and consumption experiences. For example, consider the two images in Figure 1. The first image is hashtagged with *eddiebauer* and the second, with *prada*. Image (a) shows a person wearing Eddie Bauer and resting on a rock, facing mountains in the distance. Image (b) shows the user taking a picture of herself wearing Prada sunglasses and bright red lipstick. The two images differ in terms of their content (mountainous landscape vs. head shot) and their visual properties (color palate, contrast, etc.) They also differ in terms of the mood and experience

¹ <https://instagram-press.com/our-story/>, accessed August 15th 2019; <https://www.socialpilot.co/blog/social-media-statistics>, accessed August 15th 2019.

² <https://www.instagram.com/explore/tags/nike/>, accessed August 15th, 2019.

consumers may associate with the two brands. The images created and shared by consumers offer a new avenue to study brand associations and consumers' brand perceptions. Firms need to monitor how consumers are portraying their brands on social media, or risk missing a large part of consumer brand conversations.

Figure 1. Sample Images from Instagram Hashtagged with Brand Names



(a) #eddiebauer



(b) #prada

Academic researchers and practitioners have developed a set of tools for monitoring social media and listening in on consumers' online brand conversations. The focus so far has been on text content. However, given that images are on their way to surpass text as the medium of choice for online conversations, and the rich information about the consumption experience and feelings conveyed through images, monitoring visual content is important to get a more complete understanding of brand-related online conversations.

We introduce a "visual listening in" approach for monitoring visual brand content created and shared by consumers on social media. We develop and validate the BrandImageNet model to allow firms to monitor their brand portrayal on social media and evaluate it relative to competitors' and to the firm's preferred brand positioning. The BrandImageNet model maps images to specific perceptual attributes (e.g., glamorous, rugged) and allows us to measure how brands are portrayed along these attributes. For example, is Prada portrayed as a glamorous brand and is it portrayed as more glamorous than Eddie Bauer? We focus on identifying perceptual brand attributes from images, which is different from identifying functional attributes

of the product (Gardner and Levy 1955, Park et al. 1986, Keller and Lehmann 2006). Positioning brands on intangible attributes is at the core of many marketing strategies because it allows firms to differentiate themselves from one another in categories with functionally similar products. In categories such as beverages and apparel (our focus categories in this paper) consisting of brands offering products with similar functionality, the feelings, perceptions, and attitudes consumers have toward the brand often make a big difference. The intangible brand attributes allow consumers to choose the brand that feels most appropriate to them and to express themselves through that brand.

We develop our BrandImageNet model using a deep-learning framework to predict whether a particular perceptual attribute (we study four: glamorous, rugged, healthy, and fun) is expressed in a given image. One challenge with studying images is that image data are unstructured. To draw any inferences, a first and important step is to create a meaningful representation of the images. Typically, doing so would require creating a feature vector summarizing the image data. An emerging area of machine learning, called deep learning, uses “deep” neural networks to simultaneously learn both the model parameters and the feature representation needed for prediction tasks. It works with raw image data directly, and automatically learns the representation of data with multiple levels of abstraction (LeCun et al. 2015). Our BrandImageNet is a multi-label convolutional neural network (ConvNet), fine-tuned from the Berkeley Vision and Learning Center (BVLC) Reference CaffeNet model on a set of Flickr images we collected for brand attribute identification. It achieves a high average out-of-sample accuracy of 90% and an average AUC of 0.85.

We apply our BrandImageNet model to consumer-created and firm-created brand images shared on Instagram. We collect a set of images hashtagged with brand names for a total of 56 brands in two product categories: apparel and beverages. Using the BrandImageNet model, we compute the probability of attributes presence in each brand image. Then, for each brand and each of the brand attributes we study, we compute an image-based brand image (IBBI) metric as the average probability of attribute presence across all consumer-generated images of this brand. A higher score indicates more images portray the brand as having a given brand attribute. We

undertake validation analyses and show that our model predictions have a high degree of agreement with human judges evaluating images and that the IBBI metric reflects consumer brand perceptions. Specifically, we find strong positive correlation between IBBI and brand perception scores from survey data (Y&R BAV survey and our online brand perceptions survey of active Instagram users). These results show that consumers' brand portrayal on social media reflects and is indicative of consumer brand perceptions. Consumer-created brand images contain valuable brand-related information, and our model is able to pick it up.

Our study makes both a methodological and substantive contribution to the marketing literature. We propose a new approach to measure consumers' brand perceptions from the voluble stream of public consumer-generated brand imagery. We validate the model and demonstrate its performance through tests against human judgments and surveys of consumer brand perceptions. This study is one of the first to use a deep-learning approach in marketing. Deep-learning methods are becoming more popular and have significantly improved performance of many computer vision tasks (e.g., object detection); however, to the best of our knowledge, no extant work on predicting perceptual brand attributes from images exists.

Substantively, our study establishes a link between consumer-created images and brand perceptions. Past research has examined how firms use imagery, such as logos and image advertisements, to position brands, and how visual stimuli influence consumer behavior (Wedel and Pieters 2007). We show that visual content generated by consumers reflects brand perceptions. We present evidence showing that consumer-created brand images contain valuable brand information. Our approach goes beyond the practice of simply counting the frequency of brand logo presence in the images. Many firms have started to look into consumer-created visual content on social media, but most efforts are focused on brand logo detection and tracking brand mentions. We show that deeper insights can be extracted from consumer-created images. Firms can use our BrandImageNet model to monitor their brand portrayal and examine the effectiveness of their positioning strategies. Because our BrandImageNet model can be easily extended to include other attributes and is highly scalable, brands can use it to track relevant brand perceptions automatically and in real time from images posted on social media.

The rest of the paper is organized as follows. First, we discuss related literature in both marketing and computer science. Then, we introduce our BrandImageNet model and discuss its architecture, training process, and the data we use to train it. Next, we apply our model to brand-related images posted on Instagram to measure consumers' brand portrayal and demonstrate its performance using human judges and survey instruments. We conclude with a discussion of limitations and directions for future research.

2. Related Work

This paper is related to three streams of research in marketing and computer science: user-generated content (UGC), visual marketing, and computer vision and deep-learning applications in marketing. We review how we contribute to all three below.

Mining UGC for online marketing intelligence is becoming increasingly popular in both marketing and computer science research, as well as among practitioners. In the marketing literature, Netzer et al. (2012) propose an approach to help firms understand market structure and monitor the topics discussed in relation to their brand by mining brand associations from consumer-generated content on forums. Culotta and Cutler (2016) measure consumers' brand perceptions by mining the brands' social connections through social networks on Twitter. Several papers investigate the relationship between consumer-generated online content (e.g., product reviews and ratings) and sales (e.g., Chevalier and Mayzlin 2006, Liu 2006, Archak et al. 2011). The focus so far has been on text. Only recently have marketing researchers begun to consider visual content (Pavlov and Mizik 2019, Zhang and Luo 2018, Zhang et al. 2018). We propose leveraging the ever-increasing flow of consumer-generated images available online for marketing research. We contribute to this literature by advancing a new approach to measuring consumers' brand perceptions, using images consumers post on social media.

Images constitute a large part of consumer online conversations. With the exclusive focus on text content, a significant portion of online brand conversations is not being "heard." Indeed, many consumers, particularly millennials, prefer visual-based online communication, such as Snapchat or Instagram, to text. Further, for some consumer categories (e.g., beverages, apparel), visual content is more easily available than text content. Consider, for example, the beverages

category. Few people would write a review of bottled water or orange juice, but many people tag Fiji Water, Tropicana, and Minute Maid on Instagram. Finally, text mining of UGC tends to recover mostly functional product attributes. Netzer et al. (2012), for example, recover topics in a car discussion forum related to functional attributes such as powertrain warranty, good mileage, suspension noise, and emergency brakes. Similarly, Timoshenko and Hauser (2019) recover functional needs for oral care products, such as “An oral-care product that does not affect my sense of taste” and “An oral care product that is quiet.” Tirunillai and Tellis (2014) recover attributes such as ease of use, reliability, portability, safety, and comfort. We demonstrate that recovering non-functional dimensions of brand image from visual content is possible. Images may convey information about the usage situation, the setting, mood, and feelings associated with the product. Indeed, although a consumer is unlikely to write that he/she is feeling rugged wearing Levi’s jeans, a photograph tagged with Levi’s may convey ruggedness.

Our work also contributes to the literature on visual marketing, which has studied how consumers perceive different visual stimuli. Firms have been using visual stimuli to shape consumer brand perceptions through brand logos, advertisements, retail-store decorations, and, more recently, social media. Raghubir and Greenleaf (2006) examine how geometry, specifically, the ratio of the sides of a rectangular product or package, affect consumers’ purchase intentions and preferences. Zhang et al. (2009) model the effects of feature-ad design characteristics (e.g., ads’ surface size and number of supporting colors) on sales of the feature product, and investigate the mediating role of attention. Wedel and Pieters (2014) investigate how color schemes affect rapid gist perception of an ad, product category, and brand recognition under brief and blurred exposure conditions. They find that color of the central diagnostic object in the ad plays a key role in protecting gist perception. Xiao and Ding (2014) find that faces in print ads affect ad effectiveness. Our analysis contributes to this literature by demonstrating that visual content can convey perceptual brand attributes. We find that brand images created and shared by consumers reflect their brand perceptions.

Mining visual UGC requires different tools than mining text UGC. We draw on image-classification methods developed in the computer vision and deep learning literature. Until

recently, much of this literature has focused on object identification and recognition. The field of computer vision was originally developed for engineering applications, such as automatic product inspection, autonomous vehicle navigation, surveillance, and medical image processing to aid diagnostics. More recently, researchers have begun developing tools for predicting image aesthetics, image style or genre, image interestingness (e.g., Karayev et al. 2014, Jiang et al. 2013), and visual sentiment (Giannakopoulos et al. 2015). Our work is closely related to this stream of research on predicting abstract constructs from images. We study a new problem: measuring perceptual brand attributes from images. We solve the problem using deep convolutional neural network architecture, which is a type of deep neural networks (LeCun et al. 2015). Deep-learning methods are well suited for marketing problems, because of their ability to handle unstructured data, scalability, and superior predictive performance. As marketing researchers begin to adopt machine-learning methods (e.g., Liu et al. 2019, Zhang and Luo 2018, Gabel et al. 2019), ours is among the first applications of deep-learning methods to marketing problems.

3. Machine-learning-based Identification of Brand Attributes in Images

To measure how brands are portrayed along particular attributes on social media, we need a scalable method to determine whether a particular brand attribute is expressed in a given image: Does the image look rugged? Fun? Glamorous? Healthy? In fact, an image can depict multiple brand attributes. An image can, for example, look both rugged and fun. This attribute identification problem can be formulated as a multi-label image-classification task—that is, classifying images into multiple brand attributes—where each image may depict more than one attribute simultaneously. To solve this problem, we need (1) an annotated data set of images labeled with brand attributes to use for training, and (2) an algorithm that learns the underlying function mapping images to brand attributes. In this section, we describe how we created a data set of labeled images and developed our multi-label convolutional neural network model, BrandImageNet, for identifying perceptual attributes in an image.

To illustrate our brand attribute identification method, we explore four brand attributes that are highly relevant in the apparel and beverages brand categories we study: glamorous,

rugged, healthy, and fun. We selected these from the Young and Rubicam's BAV consumer-brand-perception survey (Mizik and Jacobson 2008) to allow subsequent comparisons of our model-based findings with external survey-based findings. We selected these specific four brand attributes because they are meaningful differentiators in the apparel and beverage sectors (they have a higher range of scores across brands in these sectors in the BAV survey).

3.1. Data

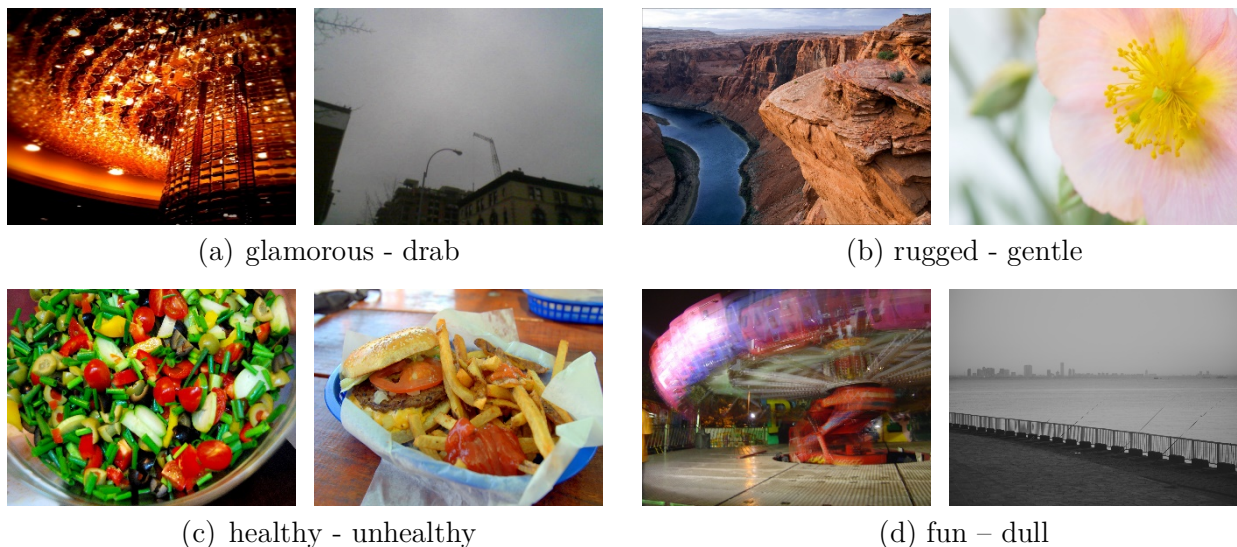
Training a multi-label image-classification model of brand attributes requires an annotated data set consisting of images labeled with respect to whether or not they express each attribute. Unlike object detection tasks, which are typically trained on large public data sets of labeled images, no existing data set is annotated with perceptual brand attributes. Thus, we created one.

We gathered an annotated set of images from Flickr, an online photo-sharing website. Users share photos on Flickr, label their uploaded photos with titles and descriptions, and provide free-form tags. Flickr has been used as a data source in previous visual and social network research (e.g., Zhang et al. 2012, Dhar et al. 2011, McAuley and Leskovec 2012). Flickr lends itself well to gathering an annotated set, because unlike many other social media platforms (e.g., Instagram), it provides a search engine that returns the most relevant images for a keyword. The search is based on text labels provided by users, image content, and clickstream data (Stadlen 2015). An image ranked at the top for a particular query has often been validated by tens of thousands of users who clicked on the image, reflecting a large population consensus regarding a strong association between the image and the query term. Because recognizing brand attributes from images is a more subjective task than recognizing objects from images, using annotations based on a large population consensus is important to decrease noise in the labels data.³

³ We considered collecting a training set of Instagram images based on hashtags as an alternative. However, a hashtag is too noisy to serve as an effective supervisory label, because users apply hashtags that are irrelevant or leave out relevant hashtags. As a result, training models using such weakly supervised data usually require extremely large data sizes (Mahajan et al. 2018). Using MTurk to label the images is also not practical for our purposes. The costs of securing a sufficiently large sample of images and human judges on Mturk to get reliable labels is prohibitive. Flickr, on the other hand, is cost effective and is also more reliable. In section 4, we present evidence validating the appropriateness of our BrandImageNet model (trained on Flickr) for application to Instagram data.

For each attribute, we queried the attribute term on Flickr search engine and collected about 2,000 images in the top search results. We use these images as our positive instances (i.e., images that express the brand attribute). We also needed to collect negative instances for each attribute consisting of images that do not express that brand attribute (images that are not rugged, not glamorous, etc.) We need both the positive and negative images for the algorithm to learn how to separate them. To gather negative images, we queried the antonyms of our attributes in Flickr search engine (drab for glamorous, gentle for rugged, unhealthy for healthy, dull for fun) and again collected about 2,000 images in the top search results. To capture a wider spectrum of the negative cases, we also used the images collected for the other attributes and their antonyms as negative examples of a given brand attribute, if they were not already included in the positive image set of that brand attribute. For example, for the attribute healthy, we used “healthy” as the query to collect positive instances. We used images that were returned for the query “unhealthy” and images that were returned for all other terms (glamorous, rugged, fun, drab, gentle, dull) but were not in the “healthy” set, as negative instances. The entire annotated set has a total of 16,368 images. Each image is associated with four labels for the four brand attributes we study. Figure 2 shows one sample image returned by Flickr for each queried brand attribute and its antonym.

Figure 2. Sample Images of Brand Attributes and Their Antonyms in Flickr Data Set



3.2. BrandImageNet Model

Next, we built and trained a multi-label image-classification model that maps images to perceptual brand attributes.

Image classification is challenging because the raw input data are difficult to capture quantitatively, more difficult than other unstructured data, such as text. Text consists of words, which can be grouped by topic or classified into positive or negative valence. Images, on the other hand, are composed of pixels, and a single pixel in isolation does not lend itself to meaningful interpretation. Therefore, any modeling of image data requires creation of a meaningful representation of the images. It requires creating a set of features, or predictor variables, through some complex transformations of the raw pixel data. The process of defining features is a critical step in achieving good predictive performance. Manual feature engineering is hard, time-consuming, and requires great domain knowledge, and there is no systematic approach that would guarantee finding features that best characterize the underlying problem. Thus, we adopt a deep-learning approach (LeCun et al. 2015).

Deep learning does not require human input to manually engineer image features to represent raw image data. It is a representation learning method, meaning it automatically optimizes the feature-extraction step for the prediction problem. It learns the features and trains the classifier over these features simultaneously, by minimizing classification error. A deep neural network does so by applying several “layers” of simple non-linear transformations of the raw data. The result, or output, of each transformation serves as the input into the next transformation layer. Through multiple layers of such transformations, the network extracts higher- and higher-level representations of the data, allowing the final layer to easily classify the data. For example, working with image data, lower layers of a deep-learning model may extract edges and textures, whereas higher layers detect motifs, object parts, and complete objects (Goodfellow et al. 2016). The final layer maps the resulting features onto the target variables with a classification function.

3.2.1. A Multi-label Convolutional Neural Network

We set up a multi-label convolutional neural network to identify multiple brand attributes from images. ConvNet lends itself well to this problem. It is a state-of-the-art method most commonly

used for analyzing visual imagery. It has demonstrated an outstanding performance in the computer vision field for tasks such as object detection and classification, as well as in other fields (e.g., medical image analysis). ConvNets vary in number of layers and each layer's type, parameters, and hyperparameters. A typical ConvNet involves four types of layers: convolutional, activation, pooling, and fully-connected layers.

The architecture of BrandImageNet, our multi-label ConvNet, is depicted in Figure 3. The input into the network is an image, which is 227 pixels wide and 227 pixels tall. Each pixel is represented in the RGB system by three integers ranging from 0 (black) to 255 (white), reflecting the contribution of each of the three (red, green, blue) color channels.⁴ The image is represented by a three-dimensional matrix and a total of 154,587 (227 x 227 x 3) numbers. The output of the network is an array of k numbers, with each number representing the probability of the image expressing a particular brand attribute.⁵ In our case, $k=4$ and the output contains the probabilities of the image being glamorous, rugged, healthy, and fun. The brand-attribute labels are not mutually exclusive (i.e., an image can be classified as positive on multiple attributes).

Our multi-label ConvNet model maps the input, an image (the original matrix of pixel values), to the output, the probability scores, through eight trainable layers (five convolutional layers and three fully-connected layers) and 15 non-trainable layers (ReLU activation, max pooling, normalization, dropout, and sigmoid layers). A trainable layer contains model parameters to be learned and hyperparameters to be set. The non-trainable layers only contain hyperparameters to be set. We describe each type of layer in detail below.

The architecture of our network is based on and is modified from the Berkeley Vision and Learning Center (BVLC) Reference CaffeNet model, which is an implementation, with minor variations, of the widely used AlexNet (Krizhevsky et al. 2012).⁶ Both AlexNet and BVLC Reference CaffeNet models are designed to classify a given image as one of 1,000 different objects (e.g., a tree, a ship, etc.) and have been trained on the ImageNet data set of 1.2 million images

⁴ RGB constructs all possible colors as a combination of the red, green, and blue channels.

⁵ Regular lower-case letters denote scalars, bold lower-case letters denote vectors, and bold upper-case letters denote matrices and sets.

⁶ https://github.com/BVLC/caffe/tree/master/models/bvlc_reference_caffenet

map represents a certain type of local feature extracted. For example, the first convolutional layer in our network has 96 kernels and outputs 96 two-dimensional feature maps.

A convolutional layer l is defined by several hyperparameters: number of kernels d^l , size of each kernel's weight matrix (k^l, k^l, d^{l-1}) , padding p^l , and stride s^l .⁷ The padding p^l determines the number of zeros to pad around the border of the input before the convolution operation. The stride s^l determines the number of pixels the kernel slides, horizontally or vertically, between applications of the kernel to the input image during convolution operation. For example, $s^l=2$ means the kernel will be applied to every other pixel. We set these hyperparameters to values specified in the BVLC Reference CaffeNet model. The weight matrices $\{\mathbf{K}_1^l, \dots, \mathbf{K}_{d^l}^l\}$ and biases $\{b_1^l, \dots, b_{d^l}^l\}$ of the d^l kernels are the parameters of the convolutional layer l we estimate.

The convolutional layer l outputs \mathbf{Y}^l , consisting of d^l two-dimensional feature maps of size $H^l \times W^l$, where $H^l = (H^{l-1} - k^l + 2p^l)/s^l + 1$ and $W^l = (W^{l-1} - k^l + 2p^l)/s^l + 1$. The i^{th} feature map, denoted \mathbf{Y}_i^l , is the result of the convolution of the i^{th} kernel and input from previous layer \mathbf{Y}^{l-1} , defined as

$$\mathbf{Y}_i^l = \sum_{j=1}^{d^{l-1}} \mathbf{Y}_j^{l-1} * \mathbf{K}_{ij}^l + b_i^l,$$

where \mathbf{K}_i^l is the weight matrix and b_i^l is the bias of the i^{th} kernel, j indexes the depths of the input data and kernel matrix \mathbf{K}_i^l . $*$ denotes convolution. The unit at position (m, n) of the i^{th} feature map, denoted $\mathbf{Y}_i^l(m, n)$, is computed as

$$\mathbf{Y}_i^l(m, n) = \sum_{j=1}^{d^{l-1}} \sum_{q=1}^{k^l} \sum_{r=1}^{k^l} \mathbf{Y}_j^{l-1}(q + (m-1) \times s^l - p^l, r + (n-1) \times s^l - p^l) \cdot \mathbf{K}_{ij}^l(q, r) + b_i^l.$$

For example, the size of the kernels in the first convolutional layer is (11, 11, 3). A unit in a feature map is a dot product of a kernel and a $11 \times 11 \times 3$ region in the input image it connects to. Convolutional layers utilize parameter sharing (i.e., units in the same output feature map share the same kernel weights and biases).

⁷ d^{l-1} is equal to the depth (the third dimension) of the input data from the previous layer.

ReLU Activation Layer. An activation layer takes the input from the previous layer and outputs activation maps by applying an element-wise non-linear activation function over the input volume. A ReLU activation layer uses rectified linear unit as its activation function, defined as

$$ReLU(x) = \max(0, x).$$

The activation function is an element-wise operation. Thus, the dimensions of the output and the input are the same. In our model, all five convolutional layers and the first two fully-connected layers are followed by a ReLU layer. A ReLU layer does not change data size.

Max Pooling Layer. A pooling layer, sometimes called a subsampling layer, is often inserted between convolutional layers to progressively reduce the spatial size (height and width) of the data representation. It helps reduce the number of model parameters to be learned and control overfitting. It also helps make the extracted feature maps robust to small image distortions.

A max pooling layer applies a max pooling function to down-sample the input data. Similar to convolutional layers, it has hyperparameters defined by the size of the max filter (k^l, k^l) and stride s^l . For each depth slice of the input data, it slides spatially and computes a max over local neighborhoods of size (k^l, k^l) and disposes of the rest, thus down-sampling the input data. For example, a max pooling layer with a max filter of size (2, 2) and a stride of 2 will down-sample an input of size (H, W, D) to an output of size (H/2, W/2, D). Note that as it operates independently on every depth slice of the input, the depth of output is the same as that of the input. In our model, the output of the first, second, and fifth ReLU layer passes a max pooling layer.

Fully-connected Layer. Following the convolutional layers are three fully-connected layers. Unlike convolutional layers where each unit is connected to a local region of data from the previous layer, units in a fully-connected layer are connected to the entire previous layer. The unit i in a fully-connected layer l is defined as the dot product of the entire output from previous layer \mathbf{Y}^{l-1} and a unique weight matrix \mathbf{W}_i , plus a bias term b_i :

$$y_i^l = \mathbf{Y}^{l-1} \cdot \mathbf{W}_i + b_i.$$

Each unit in a fully-connected layer is associated with a different weight matrix and bias offset. For example, the final fully-connected layer FC8 contains four units. Each of these four units has

a different bias term and weight vector (it is a vector because the size of the input from the previous layer is (4096, 1)).

The resulting dot products from fully-connected layers FC6 and FC7 are passed through a ReLU layer to add non-linearity. The output of fully-connected layer FC8 is fed into the final Sigmoid8 layer for final multi-label classification.

Sigmoid Layer. The final Sigmoid layer is an activation layer. It takes the input from the previous layer and applies a sigmoid function to each element of the input to produce an output between 0 and 1:

$$\text{Sigmoid}(x) = \frac{1}{1 + e^{-x}}.$$

In our network, it takes the four output units from FC8 and produces final predicted probabilities for our four brand attributes.

Other Layers. We have also added normalization layers (to limit the unbounded activation from ReLU activation layers) and dropout layers (to reduce the number of model parameters) after fully-connected layers FC6 and FC7 for regularization purposes to reduce overfitting, as in Krizhevsky et al. (2012). Table A in the appendix offers full details of our network architecture and hyperparameters of each layer.

3.2.2. Model Training with Fine-Tuning

Convolutional layers and fully-connected layers are trainable layers with model parameters to be learned. These parameters are the kernel matrices and biases $\{\mathbf{K}, \mathbf{B}^c\}$ in the convolutional layers and the weight matrices and biases $\{\mathbf{W}, \mathbf{B}^f\}$ in the fully-connected layers. We train the model by minimizing a cross-entropy loss function using stochastic gradient descent:

$$L(\mathbf{K}, \mathbf{B}^c, \mathbf{W}, \mathbf{B}^f) = -\frac{1}{N} \sum_n \sum_{a \in \mathbf{A}} [y_n^a \log(\text{Prob}(y_n^a = 1)) + (1 - y_n^a) \log(1 - \text{Prob}(y_n^a = 1))],$$

where N is the number of training images, $\mathbf{A} = \{\text{glamorous, rugged, healthy, fun}\}$, y_n^a is the label of the image n for attribute a (1 for positive training examples and 0 for negative training examples), and $\text{Prob}(y_n^a = 1)$ is the probability of the n^{th} image expressing attribute a .

A challenge with training a ConvNet model is that it requires very large sets of training data, because both the kernel and the weight matrices have high dimensionality and can have

millions of parameters. Training the model from scratch (i.e., initializing all the parameters with random numbers) with a small dataset like ours can result in overfitting.

We adopt a fine-tuning approach to avoid the overfitting problem. That is, instead of initializing model parameters with random numbers, we initialize model parameters using learned parameters from the BVLC Reference CaffeNet model.⁸ The fine-tuning approach is a case of transfer learning (Bengio et al. 2011, Bengio 2012): knowledge learned in one domain is transferred to help model learning in another domain. It is similar to using an informed prior in Bayesian estimation. Previous work has shown that fine-tuning a pre-trained network can significantly increase model performance and avoid overfitting (Donahue et al. 2014, Yosinski et al. 2014, Girshick et al. 2014).

We fine-tune our BrandImageNet model on our data as follows. We initialize parameters of the five convolutional layers and the first two fully-connected layers FC6 and FC7 with parameters learned from the BVLC Reference CaffeNet model. We fine-tune the parameters with a small learning rate of 0.0001 over 10,000 iterations. We initialize parameters of the final fully-connected layer (FC8) with random values and train it from scratch on our data set with a larger learning rate of 0.001. We use a higher learning rate for the final layer to allow the parameters in the final layer to change quickly with our data. We use a smaller learning rate for earlier layers to preserve the parameters learned from the BVLC Reference CaffeNet model and to transfer their extracted features to our application.

We use the Caffe deep-learning framework (Jia et al. 2014) to fine-tune our BrandImageNet model on a single K80 GPU node in a university's high-performance cluster. We use 80% of the images for model training, 10% of the images as a validation set to choose the best model snapshot, and the remaining 10% as a hold-out sample for evaluation of model performance.

⁸ As we later discuss in the Sensitivity Analyses section, we undertook multiple tests of model fine-tuning (e.g., varying the fine-tuning learning rate and using different values to initialize parameters) to train our BrandImageNet model for best performance. We found that initializing with BVLC Reference CaffeNet parameters and using a 0.0001 learning rate in Conv1-Conv5 and FC6-FC7 and with a 0.001 learning rate in FC8 generates the best model according to model performance on the validation set. Thus, we focus on this specification in our main discussion. This specification also generates best performance in Flickr out-of-sample tests and in Instagram tests against human judges and consumer-perceptions surveys.

We train the model for 10,000 iterations on the training set and choose the iteration at which the model has the best performance on the validation set to select our best model snapshot. We get the best model at iteration 2,800.

3.3. Model Evaluation

We evaluate the performance of our multi-label BrandImageNet model using its per-attribute prediction accuracy and AUC (Area Under the ROC Curve, Bradley 1997). The prediction accuracy is computed as the number of correctly predicted instances divided by the total number of instances. We set the predicted image label for the attribute a presence equal to one, if the predicted probability of attribute a presence is greater than 0.5, and set it to zero otherwise.

The AUC is the probability that the classifier will rank a randomly chosen positive instance higher than a randomly chosen negative instance (i.e., assign a higher probability of being positive to the positive than to the negative). It evaluates how well a classifier ranks its predictions and separates two classes. The higher the AUC, the better the classifier at separating positive from negative instances.

Table 1 presents our model performance for the four perceptual brand attributes we study. Our multi-label ConvNet model achieves an average accuracy of 90% and an average AUC of 85%, demonstrating good predictive performance.

Table 1. Out-of-Sample Predictive Performance for Multi-label BrandImageNet Model

	Accuracy	AUC
glamorous	88.6%	0.846
rugged	91.3%	0.853
healthy	89.9%	0.859
fun	89.4%	0.827

4. Brand Portrayal on Social Media

We have trained BrandImageNet, a multi-label ConvNet model, to ascertain the presence of four perceptual attributes in an image. We wish to apply this model to brand images posted on social media to understand consumer brand portrayal. In this section, we first show that our model trained on Flickr images can be applied to and performs well on Instagram data. Then, we show

that brand portrayal in the consumer posts on social media reflects consumers' brand perceptions. Finally, we present an illustration: a case study of underwear brands.

4.1. Instagram Brand Image Data

Instagram is an image-based social media platform and a popular communication medium. It was launched in 2010 and has 1 billion active monthly users who, on daily basis, share an average of 95 million photos. Instagram users often hashtag brands, creating a collection of brand-related images.

We collected data for 56 large national brands covered in the Y&R BAV survey (Mizik and Jacobson 2008) in two product categories for which consumers frequently post photos: apparel and beverages. For each brand, we used Instagram's API to obtain consumer-created photograph posts hashtagged with the brand name. We filtered out spam, resale, and photos posted by the official account of the brand, and retained about 2,000 photographs for each brand. The final consumer-generated-image data set contains 114,367 photographs. All data were collected between May and October 2016.

We also collected firm-created brand images from the brands' official Instagram account pages, giving us a data set containing 72,089 images. Three beverage brands included in our consumer image data are not represented in this set because they do not have an official account on Instagram.

4.2. Model Application to Instagram Image Data

We use our multi-label BrandImageNet model to compute the probabilities that a given brand image expresses the four brand attributes we study. Then, we compute the average probability of attribute presence across all consumer-generated images for each brand and each attribute. We use these average probabilities as our metric of how a brand is portrayed by consumers. The higher this score, the more visual content portrays the brand as having a given attribute.

That is, let $\{\mathbf{X}_1^{(b)}, \dots, \mathbf{X}_{N^{(b)}}^{(b)}\}$ be $N^{(b)}$ images hash-tagged with brand name b . We apply our BrandImageNet model to each image n of brand b to compute the probability that a perceptual attribute a is expressed in this image n : $Pr(y_n^{(b)}(a) = 1 | \mathbf{X}_n^{(b)})$, $n = 1, \dots, N^{(b)}$, $a \in$

{glamorous, rugged, healthy, fun}. We compute the average of these probabilities and use it as our IBBI metric. IBBI for brand b , attribute a is defined as follows:

$$IBBI_{ba} = \frac{\sum_{i=1}^{N^{(b)}} \Pr(y_n^{(b)}(a) = 1 | \mathbf{X}_n^{(b)})}{N^{(b)}}.$$

We compute IBBI for each of the four attributes, for the entire set of 56 brands. We also compute IBBI for each brand, using firm-created brand images to measure firms' brand portrayal. The full list of our estimates of both consumers' and firms' brand portrayal is provided in Tables OA1-OA4 in the Online Appendix. Consider, for example, Prada and Eddie Bauer brands. Based on consumer-created brand images, $IBBI_{\{\text{Prada, glamorous}\}}$ is 0.24, whereas $IBBI_{\{\text{Eddie Bauer, glamorous}\}}$ is 0.18. This difference is significant ($p < 0.001$): consumers portray Prada as more glamorous than Eddie Bauer on Instagram. On the ruggedness attribute, $IBBI_{\{\text{Eddie Bauer, rugged}\}}$ is 0.17, which is significantly higher ($p < 0.001$) than Prada's $IBBI_{\{\text{Prada, rugged}\}}$ of 0.07. The estimates we obtain match our intuition and align with the data pattern in the BAV survey. In the next section, we present a more formal evaluation of model and IBBI metric performance.

4.3 Model Validation on Instagram Image Data

We trained our model on Flickr data because doing so is cost effective and provides reliable labels for brand attributes. However, to study consumer portrayal of brands on social media, we intend to use Instagram data, which might potentially be different from the Flickr data. Thus, we undertake a thorough validation of our BrandImageNet model application to Instagram images. First, we conduct small-sample validation tests of model performance using human judges, and show our model is applicable and performs well on Instagram data. Then, we consider our IBBI metric against two surveys of consumer perceptions (BAV, a large-scale survey of a nationally representative sample of US adults, and the survey we conducted with active Instagram users) and show the IBBI metric of brand portrayal on social media reflects consumer brand perceptions.

4.3.1 Assessing Model Performance against Human Judges

First, we examine the degree of agreement between our model and human judges' identification of brand attributes in Instagram images. For each of our four brand attributes (glamorous, rugged, healthy, fun), we select a set of Instagram images classified by our multi-label BrandImageNet

model as (1) “attribute is present” (50 images with the highest predicted probability), (2) “attribute is not present” (50 images with the lowest predicted probability), and (3) “ambiguous” (50 images with around 50% predicted probability of reflecting a given attribute). We present these images (total number of images=600, 150 images per attribute) to US-based Level 3 judges on Figure Eight network and ask them to indicate whether the image is glamorous? Rugged? Healthy? Fun? Figure Eight is an on-demand data label and annotation platform specifically designed for collecting training data for machine-learning tasks.⁹ We use Level 3 judges—described by the platform as the “smallest group of most experienced, highest accuracy contributors”—to assure quality. We collect 20 independent judgments for each attribute and each image. Each judge is asked to evaluate up to a maximum of 50 images and is compensated based on the number of total evaluations they submitted.

Table 2, column 1 presents the AUC of our model using the majority opinion of human judges as the “ground truth.” Column 2 presents a measure of agreement between the model and human judges’ labeling of images. We compute agreement as the percentage of images for which the majority of human judges evaluating an image assign this image the same label as our model (1=attribute is present, 0=attribute not present). To offer a relevant benchmark for the agreement measure, we also calculate and report in column 3 this measure for each individual human judge versus the majority vote of human judges. The average percentage of images for which an individual judge agrees with the majority of judges is very close and, on average, slightly lower than the BrandImageNet model versus humans’ agreement numbers in column 2. Importantly, the Instagram data reveal the difficulty of identifying brand attributes from images and the heterogeneity in human evaluations of images: we have only seven (of 600 studied) images with 100% agreement across 20 human judgments.

⁹ <https://www.figure-eight.com/>.

Table 2. Aggregate BrandImageNet Model Performance According to Human-based Image Labels

	AUC: Model vs. the Majority Vote of Human Judges	Agreement: Model vs. the Majority Vote of Human Judges	Agreement: A Single Human Judge vs. the Majority Vote of Human Judges, Average
glamorous	0.93	83%	85%
rugged	0.96	85%	83%
healthy	0.91	78%	80%
fun	0.94	84%	80%
Average:	0.94	82.5%	82%

Note: 600 images (150 images for each attribute), 12,000 total judgments (20 judgments for each image and attribute). The model-based label is equal to 1 if the model-based probability estimate for attribute presence is greater than 50%, and zero otherwise. The human-based label for an image is equal to 1 if the majority of the judges indicate attribute presence, and zero otherwise. Agreement is the percentage of images for which the majority of human judges evaluating an image assign this image the same label as our model. The total cost of data collection is \$288.96, with an average cost per judgment of \$0.024.

Table 3 presents our measure of agreement between the model and human judges for the three types of images we selected (attribute present/ ambiguous/ attribute not present). We find high levels of agreement for unambiguous images: agreement between the model and human judgments is notably higher in Groups 1 and 3, where the presence or absence of an attribute is more pronounced. For the ambiguous images in Group 2, where the model-based probability score for an attribute presence is close to 50%, we see lower agreement between model-based and human-based labels.

Table 3. Agreement between BrandImageNet Model and Human Labels by Image Group

	Group 1: “Attribute Present” (high predicted probability)	Group 2: “Ambiguous images” (~50% predicted probability)	Group 3: “Attribute Absent” (low predicted probability)
glamorous	92%	52%	100%
rugged	100%	56%	100%
healthy	84%	52%	98%
fun	94%	58%	100%

Note: 600 images (150 images for each attribute, 50 images per attribute in each group), 12,000 total judgments (20 judgments for each image and attribute). Agreement is computed as the percentage of images for which the majority of human judges evaluating an image assign this image the same label as the

BrandImageNet model. The model-based label is equal to 1 if the model-based probability estimate for attribute presence is greater than 50%, and zero otherwise.

Table 4 presents an alternative metric to evaluate agreement between the model and human judgments. It shows the percentage of human judgments indicating the presence of an attribute in an image, by image group. As expected, we see a monotonic decline in the positive identification of attribute presence by human judges as we go from high to low model-predicted likelihood of attribute presence.

Table 4. Percentage of Human Judgments Indicating the Presence of an Attribute in Instagram Images by Group

	Group 1: “Attribute Present” (high predicted probability)	Group 2: “Ambiguous images” (~50% predicted probability)	Group 3: “Attribute Absent” (low predicted probability)
glamorous	83%	34%	8%
rugged	79%	39%	6%
healthy	76%	24%	23%
fun	80%	37%	14%

Note: 600 images (150 images for each attribute, 50 images per attribute in each group), 12,000 total judgments (20 judgments for each image and attribute).

4.3.2 Accessing IBBI Metric against Survey-based Brand Perceptions Data

To gain further insights, we compare the model-predicted IBBI scores from consumer-generated and firm-created brand images on Instagram with the data from a large-scale consumer survey of brand perception conducted by the Y&R BAV among US adults.

The BAV survey asks a large, nationally representative sample of respondents to indicate whether they perceive a given brand as possessing a particular attribute. The BAV data are summarized as a percentage of consumers who responded affirmatively to the question of whether a particular brand represents a given attribute. We use BAV data for glamorous, rugged, fun, and healthy from the first quarter of 2016 to match our image data collection timeframe.

Table 5 presents Pearson correlations for the brand metrics computed from consumer-created brand images, firm-created brand images, and BAV survey. We see a high degree of correspondence between the consumer and firm image-based IBBI measures. Most, but not all,

correlations are high and highly significant. We also find a high degree of correspondence between consumer image-based IBBI and BAV measures and firm image-based IBBI and BAV measures for the attributes that are more relevant for the category. For apparel, we observe higher correlations for the glamorous and rugged attributes. For beverages, we see high correlations for healthy.

Table 5. Correlation Analyses of Model-predicted IBBI from Consumer- and Firm-Created Images on Instagram and the BAV Survey-based Measures of Brand Perceptions

		Consumer IBBI vs. Firm IBBI	Consumer IBBI vs. BAV	Firm IBBI vs. BAV
Apparel	glamorous	0.7838***	0.5519***	0.6100***
	Rugged	0.9122***	0.5467**	0.5035**
	Healthy	0.4680**	0.1794	0.3225*
	fun	0.6061***	0.3583*	0.2883 [±]
Beverage	glamorous	0.5518**	0.4568**	0.6582***
	rugged	0.8259***	0.3596*	0.4708*
	healthy	0.7370***	0.6976***	0.4766**
	fun	0.3775*	0.1791	0.2584

Note: [±]p<.1, *p < 0.05, **p < 0.01, ***p < 0.001.

Table 5 shows high and significant correlations between IBBI measures and BAV survey data. However, a significant difference likely exists between the population of BAV respondents and the population of active Instagram users posting images online. To gain further insights, we replicate our correlation analyses using the data from a survey we conducted among active Instagram users. If our IBBI metrics of brand portrayal reflect consumer brand perceptions, we should see a higher degree of correspondence between IBBI and brand perceptions of Instagram users.

Indeed, the BAV survey data come from a nationally representative sample of US adults. Consumer-created images on Instagram are posted by active Instagram users, who tend to be relatively young and tech savvy. Moreover, many consumer brand images depict personal consumption experience. For example, at least one of three human raters we asked to review the images flagged 81.5% of the consumer-generated Prada photos as depicting personal consumption

experience, whereas only about 3.5% of the BAV survey respondents report using the brand regularly and 86.2% have never used it.

Thus, we conducted an online survey to collect brand perceptions data from a random sample of Instagram users. We followed the data-collection protocol similar to that used by BAV. We have a total of 176 Instagram users responding to the survey of apparel brands and 175 Instagram users responding to the survey of beverage brands. Similar to BAV, we removed respondents who indicated no familiarity with a specific brand. As such, the number of respondents differs by brand and is lower for lesser-known brands. We used Figure Eight to collect survey data at a total cost of \$547.20, and have an average of 148 and a minimum of 87 responses per brand.

Table 6 presents correlations between our IBBI metrics and the survey of Instagram users. We find notably higher correlations of firm- and consumer-image-based IBBI with the Instagram users' survey data than with the BAV survey data (shaded cells in Table 6 mark notable improvement over respective Table 5 results), and the average increase in correlations of Instagram users' survey with the consumer-based IBBI is somewhat greater than the increase with the firm-based IBBI. We observe higher correlations for glamorous, rugged, and fun in the apparel category. For beverages, we observe higher and more significant correlation for glamorous, rugged, and healthy. These results based on a survey of a more relevant population of Instagram users suggest our IBBI measure indeed reflects consumer brand perceptions. Although we cannot fully eliminate the population differences (a random sample of Instagram users we obtained may be different from the Instagram users who hashtag brands, and people who hashtag Victoria's Secret may be different from people who hashtag Eddie Bauer), these results further support the point that our model is able to identify brand perceptions depicted in the brand imagery.

In summary, these results provide further support for the notion that IBBI measure reflects consumer brand perceptions.

Table 6. Correlations between Model-predicted IBBI from Consumer- and Firm-Created Images on Instagram and Survey-based Measures of Brand Perceptions from Instagram Users

		BAV vs. Survey	Consumer IBBI vs. Survey	Firm IBBI vs. Survey
Apparel	glamorous	0.9503***	0.5824***	0.6325***
	rugged	0.9338***	0.6831***	0.6630***
	healthy	0.8600***	0.0842	0.1941
	fun	0.6486***	0.5672***	0.4914**
Beverage	glamorous	0.9238***	0.5001**	0.5743**
	rugged	0.5485***	0.7899***	0.6645***
	healthy	0.9482***	0.7127***	0.5350**
	fun	0.8714***	0.2130	0.3648*

Note: The shaded cells represent instances of improvement over the correlations between consumer- and firm-image-based IBBI and the BAV data reported in Table 5. Average number of respondents per brand is 148 (min=87, max =175). Cost of the survey data collection is \$547.20. *p < 0.05, **p < 0.01, ***p < 0.001.

4.4 A Case Study of the Underwear Brands: Firm and Consumers' Brand


Portrayal on Social Media

Firms engage in branding through advertising, social media, and product packaging to create a particular brand image in consumers' minds. The imagery from a brand's official account reflects the firm's positioning efforts and the desired perceptions and associations the firm wants consumers to have. BrandImageNet model allows examination of consumers' and firms' portrayal of brands on social media. By comparing perceptual attributes expressed in the official brand images with those in consumer-created images, one may be able to examine whether firm portrayal matches consumers' portrayal of the brands and make inferences on whether the firms' brand-positioning objectives have been achieved.

Although we find high and significant correlations between IBBI scores from consumer- and firm-created images across all attributes we study, we do observe few notable differences for some brands and attributes. Figure 4 presents a chart of IBBI scores for glamorous in the apparel category. Consider, for example, Victoria's Secret, Joe Boxer, and Hanes brands. All three brands come from the same apparel subcategory: underwear. Victoria's Secret has the highest IBBI for glamorous for firm images and a slightly lower IBBI for consumer images. Joe Boxer has a similar

medium level of IBBI for both firm and consumer images. Hanes has a very large differential between the firm-based and consumer-based IBBI (firm images are much more glamorous). These numbers suggest consumer perceptions of Joe Boxer are consistent with brand positioning. Consumer perceptions of Victoria's Secret fall just short of the company's positioning on glamorous. Consumer perceptions reflected in their portrayal of Hanes, however, are highly inconsistent with how the company portrays the brand.

Figure 4. IBBI Scores for Glamorous in the Apparel Category



We have validated these patterns using human judges to rule out the possibility that they are simply an artifact of model-prediction error. We selected a random set of 500 consumer- and 500 firm-created images for each brand. For Joe Boxer and Hanes, we use all available firm-created images because less than 500 firm images are available on Instagram for these brands. A total of 2,297 images are included in this test. Ten Level 3 judges on Figure Eight platform evaluated each of the 2,297 images. Table 7 presents a summary of our findings showing consistency between our model and human judgments.

Table 7. Victoria’s Secret, Joe Boxer, and Hanes Case Study

(a) IBBI Scores: BrandImageNet Model and Human Judges			
		BrandImageNet ^a	Human Judges ^b
Victoria’s Secret	Consumer images (N=500)	0.39	0.47
	Firm images (N=500)	0.43	0.63
Joe Boxer	Consumer images (N=500)	0.20	0.11
	Firm images (N=134)	0.21	0.21
Hanes	Consumer images (N=500)	0.16	0.12
	Firm images (N=163)	0.30	0.36

(b) IBBI Score Differentials between Firm and Consumer Images: Difference (T-stat, p-value)			
		BrandImageNet	Human Judges
Victoria’s Secret: Firm vs. Consumer images		0.04 (2.25, 0.02)	0.16 (7.92, <.001)
Joe Boxer: Firm vs. Consumer images		0.01 (0.24, 0.81)	0.10 (5.86, <.001)
Hanes: Firm vs. Consumer images		0.14 (7.61, <.001)	0.24 (13.98, <.001)

(c) IBBI Score Differentials between Brands: Difference (T-stat, p-value)			
		BrandImageNet	Human Judges
Victoria’s Secret vs. Joe Boxer	Consumer images	0.19 (12.14, <.001)	0.39 (22.67, <.001)
	Firm images	0.22 (8.11, <.001)	0.42 (15.75, <.001)
Victoria’s Secret vs. Hanes	Consumer images	0.23 (15.01, <.001)	0.35 (21.58, <.001)
	Firm images	0.13 (4.98, 0.03)	0.27 (10.43, <.001)
Joe Boxer vs. Hanes	Consumer images	0.04 (3.71, <.001)	-0.01 (-1.68, 0.09)
	Firm images	-0.09 (-3.26, 0.001)	-0.15 (-5.53, <.001)

Note: We collected human judgments using the procedure outlined in section 4.3.1. To contain the costs, we used 10 US-based Level 3 judges to evaluate each image. We obtain an AUC of 0.81 using majority-based human labels for each image as the “ground truth.” Agreement between human and model labels is 77%. Average agreement between each individual human judge and the majority of human judges is 83%. Data collection cost \$573.60.

^a The IBBI scores here are based on a random sample of 500 images and, as such, are different from the IBBI scores reported in Figure 3 and Table OA1 in the Online Appendix.

^b We compute human-based IBBI as follows. For each image, we compute the probability that it expresses glamor as the number of judges who label it as glamorous divided by number of judges evaluating the image (ten judges). Then, we compute the average of these numbers for a brand and use it as the human-based IBBI.

Consistent with our model predictions, human judges perceive firm and consumer images for Victoria’s Secret as highly glamorous—more glamorous than the images of Joe Boxer and Hanes. Importantly, both the model and the judges identify firm images as more glamorous than

consumer images. Joe Boxer's firm images are judged as medium on glamor for both firm and consumer images, and although the slight positive differential is not significant in model predictions, the judges perceive firm images as significantly more glamorous. Both the model and the human judges perceive firm images of Hanes as much more glamorous than consumer images of Hanes. Our model prediction scores are consistent with the human perceptions of the images. The observed differentials identified by our model in the consumer-created and firm-created brand images for a given brand and across brands are also present in human evaluations.

This case study considers differences in the depiction of brand attributes in consumer and firm images at one time point. Managers, however, can use the BrandImageNet model and our proposed IBBI metric to monitor and track changes in consumers' brand portrayal in real time to detect changes in consumers' depiction of their brand. Firms can also use this approach to track how fast consumers pick up new brand positioning or how they respond to a brand crisis.

Overall, our empirical analyses suggest that the visual content consumers post on social media reveals their brand perceptions. The method we propose can extract this information and help marketing managers monitor and identify changes in brand perceptions over time, identify discrepancies between intended brand positioning and consumers' perceptions, and better understand their brand positioning in the context of a competitive landscape.

5. Sensitivity Analyses

We undertook several sensitivity analyses to select the best model specification for our task of extracting perceptual brand attributes. For example, we considered different options for initializing our BrandImageNet model parameters for fine-tuning. We tested a set of models initialized with learned parameters from Karayev's Fine-Tuned Flickr-Style ConvNet model for image style recognition.¹⁰ Karayev's ConvNet is a multi-class classification model, which itself was fine-tuned on 80,000 Flickr images from the BVLC Caffe Reference model. Karayev's ConvNet could potentially be a better model to fine-tune for our application, because an image style recognition task is presumably more relevant to our perceptual-attribute-identification task than

¹⁰ https://github.com/BVLC/caffe/tree/master/models/finetune_flickr_style

ConvNet models developed for object identification. We do not find that to be the case. We find that initializing with the original BVLC Caffe Reference model parameters generates a model with the best performance on the validation set. It also has the best out-of-sample fit on our Flickr data and generates the highest correlations with the BAV data and with our brand perceptions survey of active Instagram users.

Further, we took steps to identify a learning rate for the best model performance on our validation data. Although the differences in model performance across alternative learning-rate specifications are somewhat minor, we find that a lower learning rate of .0001, rather than the BVLC's recommended learning rate of .001, generates the best overall model performance in both validation and hold-out samples (results available upon request).

We also evaluated our deep-learning-based BrandImageNet model performance against alternative machine-learning approaches. For example, we tested BrandImageNet model against a popular Support Vector Machines (SVM) approach (see Online Appendix for details of SVM application). SVM is not a deep-learning method. It requires the researcher to first define and extract a set of image features before undertaking image analyses. We selected a set of typical features popular in computer vision research related to the color, shape, and texture and trained a multi-label SVM classifier on these features to predict the presence of perceptual brand attributes using our sample of labeled Flickr images. Similar to our approach to training the BrandImageNet model, we trained the SVM classifier on 80% of the Flickr images, used 10% as a validation set to choose model hyperparameters, and kept the remaining 10% as a hold-out sample to evaluate model performance. Our SVM classifier achieves an average accuracy of 87% and an AUC of 0.72 across the four perceptual attributes we train it to identify (Online Appendix, Table B1). The SVM performance is lower than that of our BrandImageNet model.

6. Discussion, Limitations, and Directions for Future Research

The rapidly growing volume of visual brand-related content posted on social media is a new valuable information source for marketers and brand managers to monitor, track, and better understand brand performance. As text-mining approaches have gained popularity in leveraging UGC for brand monitoring, image-mining approaches are still relatively new and underutilized.

Our study takes a step toward understanding and utilizing consumer-created images on social media. It bridges the image-mining and machine-learning literature with the branding literature by proposing an approach to online brand monitoring and market intelligence gathering from consumer-generated images.

The proposed approach enables managers to monitor how their brands are portrayed on image-based social platforms. We find that the brand-portrayal metrics derived from consumer-created brand images are strongly correlated with survey-based metrics of consumer brand perceptions. That is, consumers' portrayal of brands on social media contains valuable information about brands: it reflects consumer brand perceptions. Firms can use our BrandImageNet model to understand consumers' brand perceptions, monitor brand performance and the success of repositioning efforts, and identify gaps in positioning strategies.

Another potential area of application for this type of model is to study an individual consumer. One can apply the BrandImageNet model to the visual content created and shared by a consumer (rather than a brand) to better understand this individual consumer's personality and identity (the focal components of the beliefs about self). Marketing theory postulates that better fit between consumer identity and brand identity is important for brand acceptance and success in the marketplace. That is, matching a consumer's identity attributes with fitting firm-generated imagery or ad copy can increase click-through rates, engagement, liking, and the likelihood of purchase. The methods described in the manuscript can be used to analyze an individual consumer's public visual data stream to extract specific attributes (e.g., healthy, glamorous, etc.) represented in the imagery the consumer creates and/or posts on social media platforms. The learned insights can be leveraged for better ad targeting by matching consumer identity characteristics with ad creatives. This application can be useful for both search engines serving ads and for the advertisers.

Our study suggests interesting new directions for research and poses new research questions. For example: What visual properties and image-content elements determine whether an image conveys a particular perceptual attribute? Which image features drive consumers' perceptions? Are perceptions shaped primarily by the focal elements of the image (persons,

products, objects in the image), the background of the image, or the design elements of the image (color scheme, texture, etc.)? These questions are interesting and important, and answering them can help improve ad design and social media campaigns. Because the BrandImageNet model utilizes a deep-learning approach, which simultaneously identifies the features and links them to perceptual attributes in the images, we are not able to answer these questions in the current study. Although ConvNet models have better predictive performance than traditional methods with feature engineering (e.g., SVM), they do not lend themselves to interpretability that may be feasible with methods where researchers explicitly define image features. As such, we leave these questions for future research.

Future work can also adapt unsupervised machine-learning methods to help uncover brand attributes that consumers care about but that brand managers might not have considered. Visual content is very amenable to unsupervised machine-learning methods. Similar to text-mining application for extracting topics and uncovering consumer needs (Timoshenko and Hauser 2019), visual content can be used to uncover new product-usage situations, identify individual personality traits of the content creator or a consumer group, and help design products and communication strategies to better fit and address the needs of the individual or consumer segment.

7. Conclusion

Visual content is a ubiquitous part of modern life. It affects consumers' beliefs, preferences, and decision-making at multiple stages. As such, incorporating visual content analyses into marketing models and management processes is imperative. The method presented in this paper provides evidence that capturing and analyzing rich image data generated by consumers is possible and can generate valuable insights for marketers. We hope our research provides a foundation and encourages more future research in this emerging area.

References

- Archak, N., Ghose, A., Ipeirotis, P. G. (2011). Deriving the pricing power of product features by mining consumer reviews. *Management science*, 57(8), 1485-1509.
- Bengio Y. (2012) Deep learning of representations for unsupervised and transfer learning. *Proceedings of ICML Workshop on Unsupervised and Transfer Learning*, 17-36.
- Bengio Y., Bergeron A., Boulanger-Lewandowski N., Breuel T., Chherawala Y., Cisse M., Erhan D., Eustache J., Glorot X., Muller X., et al. (2011) Deep learners benefit more from out-of-distribution examples. *Proceedings of the Fourteenth International Conference on Artificial Intelligence and Statistics*, 164-172.
- Chevalier J.A., Mayzlin D. (2006) The effect of word of mouth on sales: Online book reviews. *Journal of Marketing Research*, 43(3):345-354.
- Culotta A., Cutler J. (2016) Mining brand perceptions from twitter social networks. *Marketing Science*, 35(3):343-362.
- Deng J., Dong W., Socher R., Li L., Li K., Fei-Fei L. (2009) Imagenet: A large-scale hierarchical image database. *2009 IEEE Conference on Computer Vision and Pattern Recognition*, 248-255, ISSN 1063-6919, URL <http://dx.doi.org/10.1109/CVPR.2009.5206848>.
- Dhar S., Ordonez V., Berg T.L. (2011) High level describable attributes for predicting aesthetics and interestingness. In *CVPR 2011*: 1657-1664. IEEE.
- Diehl K., Zauberger G., Barasch A. (2016) How taking photos increases enjoyment of experiences. *Journal of Personality and Social Psychology*, 111(2):119-140.
- Donahue J., Jia Y., Vinyals O., Hoffman J., Zhang N., Tzeng E., Darrell T. (2014) Decaf: A deep convolutional activation feature for generic visual recognition. *International conference on machine learning*, 647-655.
- Gabel S., Guhl D., Klapper D. (2019) P2v-map: Mapping market structures for large retail assortments. *Journal of Marketing Research*, 56(4): 4557-580.
- Gardner B.B., Levy S.J. (1955) The product and the brand. *Harvard business review*, 33(2):33-39.
- Giannakopoulos T., Papakostas M., Perantonis S., Karkaletsis V. (2015) Visual sentiment analysis for brand monitoring enhancement. In *2015 9th International Symposium on Image and Signal Processing and Analysis (ISPA)*: 1-6. IEEE.

- Girshick R., Donahue J., Darrell T., Malik J. (2014) Rich feature hierarchies for accurate object detection and semantic segmentation. Proceedings of the IEEE conference on computer vision and pattern recognition, 580-587.
- Goodfellow I., Bengio Y., Courville A. (2016) Deep Learning (MIT Press), <http://www.deeplearningbook.org>.
- Jia Y., Shelhamer E., Donahue J., Karayev S., Long J., Girshick R., Guadarrama S., Darrell T. (2014) Caffe: Convolutional architecture for fast feature embedding. Proceedings of the 22nd ACM international conference on Multimedia, 675-678 (ACM).
- Jiang Y.G., Wang Y., Feng R., Xue X., Zheng Y., Yang H. (2013) Understanding and predicting interestingness of videos. In Twenty-Seventh AAAI Conference on Artificial Intelligence. 2013.
- Karayev S., Trentacoste M., Han H., Agarwala A., Darrell T., Hertzmann A., Winnemoeller H. (2014) Recognizing image style. Proceedings of the British Machine Vision Conference (BMVA Press), URL: <http://dx.doi.org/http://dx.doi.org/10.5244/C.28.122>.
- Keller K.L., Lehmann D.R. (2006) Brands and branding: Research findings and future priorities. Marketing Science, 25(6):740-759.
- Krizhevsky A., Sutskever I., Hinton G.E. (2012) Imagenet classification with deep convolutional neural networks. Advances in neural information processing systems, 1097-1105.
- LeCun Y., Bengio Y., Hinton G. (2015) Deep learning. Nature, 521(7553):436-444.
- Liu X., Lee D., Srinivasan K. (2019) Large scale cross category analysis of consumer review content on sales conversion leveraging deep learning. Journal of Marketing Research, forthcoming.
- Liu Y. (2006) Word of mouth for movies: Its dynamics and impact on box office revenue. Journal of Marketing, 70(3):74-89.
- McAuley J., Leskovec J. (2012) Image labeling on a network: using social-network metadata for image classification. European Conference on Computer Vision, 828-841 (Springer).
- Mizik N., Jacobson R. (2008) The financial value impact of perceptual brand attributes. Journal of Marketing Research, 45(1):15-32.
- Netzer O., Feldman R., Goldenberg J., Fresko M. (2012) Mine your own business: Market-structure surveillance through text mining. Marketing Science, 31(3):521-543.

- Park C.W., Jaworski B.J., MacInnis D.J. (1986) Strategic brand concept-image management. *Journal of Marketing*, 50:135-145.
- Pavlov E., Mizik N. (2019) Increasing consumer engagement with firm-generated social media content: The role of images and words. Working Paper, University of Washington.
- Pedregosa F., Varoquaux G., Gramfort A., Michel V., Thirion B., Grisel O., Blondel M., Prettenhofer P., Weiss R., Dubourg V., Vanderplas J., Passos A., Cournapeau D., Brucher M., Perrot M., Duchesnay E. (2011) Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research*, 12:2825-2830.
- Raghubir P., Greenleaf E.A. (2006) Ratios in proportion: what should the shape of the package be? *Journal of Marketing*, 70(2):95-107.
- Stadlen, A. (2015), "Find every photo with Flickr's new unified search experience," <https://blog.flickr.net/en/2015/05/07/flickr-unified-search/>, May 7, 2015.
- Timoshenko A., Hauser J.R. (2019) Identifying customer needs from user-generated content. *Marketing Science*, 38(1):1-20.
- 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*, 51(4):463-479.
- Wedel M., Pieters R., eds. (2007) *Visual marketing* (New York: Lawrence Erlbaum Associates).
- Wedel M., Pieters R. (2014) The buffer effect: The role of color when advertising exposures are brief and blurred. *Marketing Science*, 34(1):134-143.
- Xiao L., Ding M. (2014) Just the faces: Exploring the effects of facial features in print advertising. *Marketing Science*, 33(3):338-352.
- Yosinski J., Clune J., Bengio Y., Lipson H. (2014) How transferable are features in deep neural networks? *Advances in neural information processing systems*, 3320-3328.
- Zhang H., Korayem M., Crandall D.J., LeBuhn G. (2012) Mining photo-sharing websites to study ecological phenomena. *Proceedings of the 21st international conference on World Wide Web*, 749-758 (ACM).
- Zhang J., Wedel M., Pieters R. (2009) Sales effects of attention to feature advertisements: a Bayesian mediation analysis. *Journal of Marketing Research*, 46(5):669-681.
- Zhang M., Luo L. (2018) Can user generated content predict restaurant survival: deep learning of Yelp photos and reviews. Available at SSRN 3108288.

Zhang, S., Lee, D., Singh, P. V., Srinivasan, K. (2018). How much is an image worth? Airbnb property demand estimation leveraging large scale image analytics. Available at SSRN <https://ssrn.com/abstract=2976021>.

Appendix

Table A. BrandImageNet Model Architecture and Hyperparameters

Layer	Size	Kernel size	Stride	Pad
Input Image	$227 \times 227 \times 3$			
Convolution 1	$55 \times 55 \times 96$	$11 \times 11 \times 3$	4	0
ReLU1	$55 \times 55 \times 96$			
MaxPooling1	$27 \times 27 \times 96$	3×3	2	0
Norm1	$27 \times 27 \times 96$			
Convolution2	$27 \times 27 \times 256$	$5 \times 5 \times 96$	1	2
ReLU2	$27 \times 27 \times 256$			
MaxPooling2	$13 \times 13 \times 256$	3×3	2	0
Norm2	$13 \times 13 \times 256$			
Convolution3	$13 \times 13 \times 384$	$3 \times 3 \times 256$	1	1
ReLU3	$13 \times 13 \times 384$			
Convolution4	$13 \times 13 \times 384$	$3 \times 3 \times 384$	1	1
ReLU4	$13 \times 13 \times 384$			
Convolution5	$13 \times 13 \times 256$	$3 \times 3 \times 384$	1	1
ReLU5	$13 \times 13 \times 256$			
MaxPooling5	$6 \times 6 \times 256$	3×3	2	0
FC6	4096			
ReLU6	4096			
Dropout6	4096	dropout rate = 0.5		
FC7	4096			
ReLU7	4096			
Dropout7	4096	dropout rate = 0.5		
FC8	4			
Sigmoid8	4			

**Visual Listening In: Extracting Brand Image
Portrayed on Social Media**

Online Appendix

Table OA1. Image-based Brand Image (IBBI) Metrics from Consumer Photos – Apparel

Apparel	glamorous	rugged	healthy	fun
Abercrombie & Fitch	0.213	0.096	0.193	0.197
Adidas	0.179	0.104	0.195	0.215
Ann Taylor	0.278	0.073	0.174	0.163
Banana Republic	0.230	0.084	0.185	0.169
Carhartt	0.162	0.127	0.158	0.197
Converse	0.155	0.117	0.183	0.230
Dockers	0.155	0.103	0.199	0.208
Dolce & Gabbana	0.249	0.075	0.225	0.169
Eddie Bauer	0.130	0.179	0.148	0.180
Gap	0.169	0.084	0.220	0.219
Gucci	0.226	0.079	0.221	0.167
Guess	0.221	0.082	0.216	0.170
Hanes	0.161	0.084	0.234	0.173
Hollister	0.218	0.092	0.194	0.191
J. Crew	0.216	0.085	0.197	0.186
Joe Boxer	0.211	0.083	0.209	0.216
Juicy Couture	0.266	0.063	0.221	0.166
Kenneth Cole	0.194	0.098	0.200	0.175
Levi's	0.155	0.120	0.182	0.201
L.L. Bean	0.149	0.115	0.180	0.184
Lucky Brand Jeans	0.222	0.102	0.182	0.170
Old Navy	0.191	0.077	0.252	0.230
OshKosh	0.113	0.073	0.254	0.224
Prada	0.220	0.079	0.236	0.148
Ralph Lauren	0.187	0.094	0.198	0.201
Tommy Hilfiger	0.204	0.096	0.196	0.196
Under Armour	0.163	0.104	0.247	0.224
Urban Outfitters	0.202	0.092	0.166	0.176
Victoria's Secret	0.395	0.047	0.194	0.163

Table OA2. Image-based Brand Image (IBBI) Metrics from Consumer Photos – Beverages

Beverage	glamorous	rugged	healthy	fun
Bacardi	0.118	0.058	0.349	0.176
Bigelow's Tea	0.077	0.055	0.455	0.112
Bud Light	0.130	0.083	0.278	0.223
Budweiser	0.130	0.083	0.289	0.195
Coca-Cola	0.124	0.076	0.311	0.181
Coors Light	0.158	0.093	0.261	0.200
Corona	0.155	0.069	0.300	0.182
Dom Perignon	0.162	0.074	0.248	0.217
Dr Pepper	0.138	0.064	0.307	0.181
Fanta	0.106	0.047	0.338	0.225
Gatorade	0.098	0.102	0.285	0.254
Grey Goose	0.177	0.056	0.254	0.243
Honest Tea	0.069	0.049	0.441	0.109
Jack Daniel's	0.164	0.083	0.259	0.169
Korbel	0.129	0.068	0.345	0.144
Lipton	0.094	0.032	0.437	0.138
Minute Maid	0.084	0.048	0.409	0.179
Moet & Chandon	0.158	0.069	0.278	0.199
Monster	0.123	0.125	0.206	0.254
Mott's	0.080	0.051	0.442	0.157
NesQuik	0.074	0.032	0.463	0.156
San Pellegrino	0.075	0.054	0.426	0.137
Snapple	0.088	0.058	0.356	0.163
Swiss Miss	0.090	0.054	0.424	0.125
Tazo	0.085	0.035	0.482	0.129
Vitaminwater	0.100	0.056	0.416	0.138
Welch's	0.074	0.037	0.454	0.167

Table OA3. Image-based Brand Image (IBBI) Metrics from Firm Photos – Apparel

Apparel	glamorous	rugged	healthy	fun
Abercrombie & Fitch	0.245	0.122	0.145	0.162
Adidas	0.170	0.151	0.155	0.285
Ann Taylor	0.253	0.051	0.201	0.188
Banana Republic	0.187	0.079	0.193	0.174
Carhartt	0.123	0.179	0.137	0.198
Converse	0.130	0.122	0.153	0.304
Dockers	0.112	0.124	0.191	0.210
Dolce & Gabbana	0.344	0.070	0.136	0.212
Eddie Bauer	0.047	0.314	0.067	0.217
Gap	0.141	0.082	0.203	0.226
Gucci	0.249	0.076	0.149	0.205
Guess	0.261	0.070	0.172	0.249
Hanes	0.299	0.058	0.313	0.210
Hollister	0.214	0.101	0.191	0.212
J. Crew	0.164	0.080	0.208	0.221
Joe Boxer	0.208	0.064	0.297	0.228
Juicy Couture	0.280	0.063	0.180	0.210
Kenneth Cole	0.176	0.105	0.163	0.204
Levi's	0.144	0.139	0.158	0.212
L.L. Bean	0.058	0.238	0.103	0.173
Lucky Brand Jeans	0.232	0.115	0.138	0.192
Old Navy	0.245	0.075	0.198	0.244
OshKosh	0.166	0.079	0.254	0.258
Prada	0.270	0.085	0.111	0.204
Ralph Lauren	0.263	0.097	0.143	0.200
Tommy Hilfiger	0.264	0.096	0.149	0.225
Under Armour	0.133	0.132	0.199	0.312
Urban Outfitters	0.163	0.086	0.188	0.214
Victoria's Secret	0.426	0.048	0.158	0.214

Table OA4. Image-based Brand Image (IBBI) Metrics Based on Firm Photos – Beverages

Beverage	glamorous	rugged	healthy	fun
Bacardi	0.142	0.075	0.251	0.234
Bigelow's Tea	0.081	0.050	0.487	0.137
Bud Light	0.088	0.083	0.208	0.378
Budweiser	0.147	0.109	0.161	0.277
Coca-Cola	0.127	0.083	0.222	0.250
Coors Light	0.083	0.128	0.146	0.253
Corona	0.105	0.081	0.207	0.262
Dom Perignon	0.199	0.097	0.186	0.177
Dr Pepper	0.108	0.095	0.304	0.177
Fanta	0.124	0.060	0.214	0.326
Gatorade	0.101	0.111	0.211	0.282
Grey Goose	0.113	0.056	0.319	0.178
Honest Tea	0.061	0.061	0.405	0.155
Jack Daniel's	0.123	0.088	0.160	0.227
Korbel	0.080	0.071	0.421	0.109
Lipton	0.084	0.043	0.396	0.202
Moet & Chandon	0.191	0.071	0.228	0.202
Monster	0.089	0.157	0.132	0.322
NesQuik	0.113	0.092	0.240	0.354
San Pellegrino	0.072	0.037	0.475	0.134
Snapple	0.054	0.042	0.366	0.198
Tazo	0.101	0.035	0.465	0.112
Vitaminwater	0.115	0.036	0.168	0.461
Welch's	0.079	0.035	0.489	0.117

SVM with Human-Defined Image Features

We train a multi-label Support Vector Machine (SVM) classifier to identify perceptual brand attributes of glamorous, rugged, healthy, and fun, using the Flickr data we collected. SVM is one of the most widely used classification algorithms because it offers robust and accurate predictions, has a sound theoretical foundation, and can be trained on high-dimensional data even with relatively small sample sizes (Wu et al. 2008).

To analyze and classify image data with an SVM, we first define a set of image features to represent the images quantitatively. We define 13 features relating to color, shape, and texture (Table OA5). Many of these features are widely used in the computer vision literature, and are known to work well in object-recognition and -detection tasks.

Table OA5. List of Features Used in the SVM Classifier

Type	Feature
Color	RGB color histogram
	HSV color histogram
	L*a*b color histogram
Shape	Line: number of straight lines
	Line: percentage of parallel lines
	Line: histogram of line orientations & distances
	Line: histogram of line orientations
	Corner: percentage of global corners
	Corner: percentage of local corners
	Edge Orientation Histogram
	Histogram of Oriented Gradients (HOG)
Texture	Local Binary Pattern (LBP)
	Gabor

Color is one of the most important features of images. We use color histograms to represent the distribution of colors in each image. We compute color histograms in three standard color spaces: RGB, HSV, and L*a*b. In the RGB (red, green, blue) space, we compute a 192-dimensional color histogram, 64 bins for each of the red, green, and blue channels. That is, we first discretize colors in an image into bins, and then count the number of pixels whose color value falls into each bin (color range). In the HSV (hue, saturation, value) space, we compute a 256-

dimensional joint histogram, 16, 4, and 4 bins, respectively, in the hue, saturation, and value channels (e.g., Manjunath et al. 2001). In the L^*a^*b space (L^* stands for lightness, and a^* and b^* represent color), we compute a 784-dimensional joint histogram, 4, 14, and 14 bins, respectively, in the L , a , and b channels (e.g., Palermo et al. 2012).

Another important element of an image is the set of depicted shapes. Shape identification is a critical step in recognizing objects. It can also convey cognitive, symbolic, and perceptual meanings. We extract features related to lines and corners in the images. To compute features related to lines, we transform each image to grayscale, and then detect edges of the image using the Canny detector (Canny 1986). Then, we apply the Hough Line Transform (Duda and Hart 1972, Galamhos et al. 1999) to extract lines based on the detected edges. Each identified line is denoted by both its orientation and its distance to the top-left corner of the image, allowing us to construct the line features in Table OA1. To compute corner features, we extract corners using the Harris corner detector (Harris and Stephens 1988). We compute both the percentage of corners in the entire picture (global) and in each of the quadrants of the image (local). We also compute an edge-orientation histogram (64-dimensional feature) from Sobel operators and a histogram of oriented gradients (HOG, a 144-dimensional feature). HOG (Dalal and Triggs 2005) describes the shape within an image, using the local distributions of intensity gradients and edge direction.

The final set of features is related to texture. Texture captures the surface quality of an object, which one would normally sense through touch. Texture portrayed in an image can cause the viewer to imagine the sensation she might experience if she felt it. We extracted two features that are widely used as texture descriptors (e.g., Ahonen et al. 2006, Guo et al. 2010, Hamamoto et al. 1998, Weldon et al. 1996): the local binary pattern (LBP, a 26-dimensional feature) and the Gabor feature (a 32-dimensional feature). The LBP describes local texture patterns by encoding local pixel neighborhoods and comparing the lightness and darkness of a pixel with all its neighboring pixels (Ojala et al. 1996). By contrast, the Gabor feature captures texture at a coarser scale (Manjunath and Ma 1996). Each image is filtered using different Gabor filter kernels, and we use the mean and variance of the filtered image as features.

After extracting the features, we train a multi-label SVM to identify the presence of four perceptual brand attributes. For each perceptual attribute $a \in \{\text{glamorous, rugged, healthy, fun}\}$, our training set of images is of the form $\{(\mathbf{x}_i, y_i), i = 1, \dots, N\}$, where $\mathbf{x}_i \in R^D$ represents a D-dimensional visual-feature vector for image i , and $y_i \in \{-1, +1\}$ are class labels. The SVM finds the optimal hyperplane to separate positive and negative instances while leaving the largest possible margin on both sides of the decision boundary. In our paper, we used a L1-regularized L2-loss SVM:

$$\min_{\mathbf{w}_a} \|\mathbf{w}_a\|_1 + C \sum_{i=1}^N (\max(1 - y_i \mathbf{w}_a^T \mathbf{x}_i, 0))^2.$$

Note that our optimization function in the equation above differs from the classic SVM formulation (Cortes and Vapnik 1995) in two ways. First, we use squared hinge loss, or L2 loss, rather than the usual loss function. Second, we use L1 rather than L2 regularization, in order to impose sparsity, which is typical for problems with a very large number of features. We use scikit-learn (Pedregosa et al. 2011) for estimation.

We train the model using 80% of the Flickr images we collected, keeping 10% for validation and the remaining 10% as a hold-out. The parameter C is set using model performance on the validation set. Out-of-sample prediction results are presented in Table OA6.

Table OA6. Out-of-Sample Prediction Accuracy for SVM Model

	Accuracy	AUC
glamorous	85.8%	0.749
rugged	87.5%	0.742
healthy	87.2%	0.701
fun	88.4%	0.674
Mean	87.2%	0.717

Online Appendix References

- Ahonen T., Hadid A., Pietikainen M. (2006) Face description with local binary patterns: Application to face recognition. *IEEE transactions on pattern analysis and machine intelligence*, 28(12): 2037-2041.
- Canny J. (1986) A computational approach to edge detection. *IEEE Transactions on pattern analysis and machine intelligence*, (6): 679-698.
- Cortes C., Vapnik V. (1995) Support-vector networks. *Machine Learning*, 20(3): 273-297.
- Dalal N., Triggs B. (2005) Histograms of oriented gradients for human detection. 2005 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR'05), 1: 886-893 (IEEE).
- Duda R.O., Hart P.E. (1972) Use of the hough transformation to detect lines and curves in pictures. *Communications of the ACM*, 15(1): 11-15.
- Galamhos C., Matas J., Kittler J. (1999) Progressive probabilistic hough transform for line detection. In *Proceedings. 1999 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (Cat. No PR00149)*, 1: 554-560. IEEE.
- Guo Z., Zhang L., Zhang D. (2010) A completed modeling of local binary pattern operator for texture classification. *IEEE Transactions on Image Processing*, 19(6): 1657-1663.
- Hamamoto Y., Uchimura S., Watanabe M., Yasuda T., Mitani Y., Tomita S. (1998) A gabor filter-based method for recognizing handwritten numerals. *Pattern Recognition*, 31(4): 395-400.
- Harris C., Stephens M. (1988) A combined corner and edge detector. In *Alvey vision conference*, 15(50): 10-5244.
- Hashimoto A., Clayton M. (2009) *Visual Design Fundamentals: A Digital Approach* (Nelson Education).
- Manjunath B.S., Ma W.Y. (1996) Texture features for browsing and retrieval of image data. *IEEE Transactions on pattern analysis and machine intelligence*, 18(8): 837-842.
- Manjunath B.S., Ohm J.R., Vasudevan V.V., Yamada A. (2001) Color and texture descriptors. *IEEE Transactions on Circuits and Systems for Video Technology*, 11(6): 703-715.
- Ojala T., Pietikäinen M., Harwood D. (1996) A comparative study of texture measures with classification based on featured distributions. *Pattern recognition*, 29(1): 51-59.

Palermo F., Hays J., Efros A.A. (2012) Dating historical color images. European Conference on Computer Vision, 499-512 (Springer).

Weldon T.P., Higgins W.E., Dunn D.F. (1996) Efficient gabor filter design for texture segmentation. Pattern Recognition, 29(12): 2005-2015.

Wu X., Kumar V., Quinlan J.R., Ghosh J., Yang Q., Motoda H., McLachlan G.J., Ng A., Liu B., Philip S.Y., et al. (2008) Top 10 algorithms in data mining. Knowledge and information systems, 14(1): 1-37.