

Studying Personality through the Content of Posted and Liked Images on Twitter

Sharath Chandra Guntuku
NTU Singapore
sharathc001@e.ntu.edu.sg

Weisi Lin
NTU Singapore
wslin@ntu.edu.sg

Jordan Carpenter
Duke University
jmc51@duke.edu

Wee Keong Ng
NTU Singapore
wkn@pmail.ntu.edu.sg

Lyle H. Ungar
University of Pennsylvania
ungar@cis.upenn.edu

Daniel Preoțiu-Pietro
University of Pennsylvania
danielpr@sas.upenn.edu

ABSTRACT

Interacting with images through social media has become widespread due to ubiquitous Internet access and multimedia enabled devices. Through images, users generally present their daily activities, preferences or interests. This study aims to identify the way and extent to which personality differences, measured using the Big Five model, are related to online image posting and liking. In two experiments, the larger consisting of ~1.5 million Twitter images both posted and liked by ~4,000 users, we extract interpretable semantic concepts using large-scale image content analysis and analyze differences specific of each personality trait. Predictive results show that image content can predict personality traits, and that there can be significant performance gain by fusing the signal from both posted and liked images.

1 INTRODUCTION

People are increasingly using social media platforms in order to keep records of their daily activities, preferences and interests. This interaction allows researchers to study how people differ in the type of content they post or like. With the ubiquitous nature of camera-enabled mobile devices and Internet in many regions of the world, the complexity of content has increased in the recent years from simple text updates to images focused around platforms such as Twitter or Instagram [2].

In psychology, a key category of individual differences in behavior is represented by personality, with the Five Factor Model being the most widely used model for representing personality [3, 16]. Under this model, personality consists of five dimensions – extraversion, agreeableness, conscientiousness, openness to experience and neuroticism. Recent studies uncovered the relationship between these traits and users’ online behavior, such as language use [26] or ratings on videos [9, 27].

The aim of this paper is to study how personality is related to the content of posted and liked images on social media. For

instance, extraversion is in part an orientation toward engaging with the external, social world, which may reflect in posting images of groups by users high in this trait. Users high in openness to experience have an appreciation for art, which could lead them to post or like sketches or images containing musical instruments.

Prior research suggests that personality is strongly expressed on a platform which offers users sufficient self-expression and freedom of control [8]. Previous work in personality analysis through images either have focused on profile images which is a matter of self-presentation [10, 15, 21, 37] or use both small data sets and shallow image features [25, 29, 38], thus limiting the scope and interpretation of results [6, 28, 33, 36].

The main contributions of this work are to:

- Analyze the content of images using interpretable aesthetic and semantic features using a large-scale data set, an order of magnitude larger than previous work;
- Study posted and liked images both separately and jointly;
- Build predictive models of personality traits from image content.

While there has been previous work on predicting personality profiles using images liked on Flickr [11, 28], this is the first attempt, to the best of our knowledge, to analyze and predict personality using content from both posted and liked images on a large-scale social media data set.

Computational models that predict users’ personality from their online footprints have several applications. For social science research, these methods offer data-driven insights into human and group behaviors which can be used to generate new hypotheses for testing and can be used to unobtrusively measure large populations over time. Commercial applications include improving targeted online marketing, increasing acceptance of HCI systems and personalized recommendations.

2 DATA

We use two Twitter data sets in our experiments which differ in size and the method of acquiring labels. Both studies received approval from the IRB of our institution.

First, we use a data set (\mathcal{D}_1) containing 436 Twitter users whose Big Five personality scores were computed based on the NEO-PI-R inventory [3]. For each user we collected up to 3200 of their most recent tweets and downloaded all images that were embedded in these posts. In total, we downloaded 579,929 tweets which contained 34,875 embedded photos across 232 users. A total of 161 users posted at least 10 photos and we only include these in our experiments.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

WebSci'17, June 25-28, 2017, Troy, NY, USA.

© 2017 Copyright held by the owner/author(s). Publication rights licensed to ACM. 978-1-4503-4896-6/17/06...\$15.00

DOI: <http://dx.doi.org/10.1145/3091478.3091522>

A second data set (\mathcal{D}_2) is an order of magnitude larger and consists of 4132 Twitter users. Since we do not have personality computed via surveys for this data set, based previous work on personality analysis from profile images [15], we use an automatic text-regression method to assign each user scores for the Big Five personality traits [22]. The model was trained on a sample of over 70,000 Facebook users, using tokens and topics extracted from status updates as features, achieving a validation predictive performance of $r \sim .35$ on average for all five traits [22], which is considered a high correlation in measuring internal states [18]. We downloaded 3200 most recent user tweets, leading to a data set of 5,547,510 tweets, out of which 700,630 contained images across 3498 users. We also downloaded 3,135,764 tweets liked by these users on Twitter (an action previously known as favorite), out of which 909,861 contained an image. We exclude the texts associated with the tweets which contain images when predicting personality to limit any potential confound. Then, for our analysis, we excluded users who posted less than 20 photos, and who had less than 10 images their liked tweets.

3 FEATURES

We extract the following features from each image. To compute user level features, we perform mean feature pooling across all images posted/liked by a user.

Colors. Research has shown that colors can invoke emotions or psychological states [12]. Here, we use the HSV (Hue–Saturation–Value) space to filter grayscale images (4.99%) from subsequent color analysis. We then compute **saturation** and **brightness** followed by **Pleasure**, **Arousal** and **Dominance** [17], along with **hue count**. Hence, we applied the following formulas to extract the three dimensions of color affect: **Pleasure** = $.69 \cdot V + .22 \cdot S$, **Arousal** = $-.31 \cdot V + .60 \cdot S$, **Dominance** = $-.76 \cdot V + .32 \cdot S$ [34]. We also compute the standard deviation of the HSV values, which we will use for prediction only. In addition, we compute a 6-bin histogram which splits pixels into the primary and secondary colors as well as a 12-bin histogram containing also the tertiary colors. Finally, we compute the ratio of pixels with warm colors. The histogram features are used in prediction only. A subset of these features were successfully used in a smaller experiment on social media image content and personality [6, 33].

Image Content. Images can have very diverse content beyond faces, especially in the case of images embedded within tweets on a user’s Twitter feed. With the goal of aiding our analysis section, we first use the **Imagga** tagging API¹ as our content analysis engine, which was successfully used by past research [7]. We labeled all images with the Imagga Tagging API and generated for each image a *bag-of-tags* out of the top-10 predicted tags, based on the developers’ recommendations. We removed all tags that occurred less than 200 times in our large data set, leaving us with 1,299 distinct tags. With the aim of to increasing interpretability and decreasing sparsity, we learn tag clusters that contain frequently co-occurring tags. We use a procedure that was originally applied to words in a tweet, which produced accurate results for prediction and analysis [14, 24]. The procedure involves first computing the

¹<http://docs.imagga.com/#auto-tagging>

Feature	Ope.	Con.	Ext.	Agr.	Neu.
Colors					
Grayscale	.039	-.130	-.128	-.152	.262
Brightness	-.108	.040	.124	.027	-.020
Saturation	-.017	.023	.102	.076	-.077
Pleasure	-.017	.032	-.079	.037	-.024
Arousal	-.007	.005	.119	.048	-.054
Dominance	.005	-.013	.113	.010	-.021
Hue Count	-.094	.040	.118	.085	-.103
Content					
% Posts with People	-.106	.109	.116	.082	-.059
# Posted Images	.068	-.025	.094	-.141	.093

Table 1.A Posted Images

Feature	Ope.	Con.	Ext.	Agr.	Neu.
Colors					
Brightness	-.081	.068	.070	.087	-.093
Saturation	.049	.159	.062	.122	-.142
Pleasure	-.022	-.034	-.029	.018	.042
Arousal	.052	.150	.065	.096	-.139
Dominance	.043	.110	.055	.049	-.107
Hue Count	-.071	.030	.040	.066	-.092
Content					
% Posts with People	-.116	-.054	.117	-.041	.023
# Liked Images	-.049	-.131	.009	-.068	.159

Table 1.B Liked Images

Table 1: \mathcal{D}_2 : Pearson correlations between color features extracted from posted and liked images on Twitter and user personality, controlled for age and gender. Significant positive correlations are highlighted with green and negative correlations with red ($p < .01$, two-tailed t-test).

Normalised Pointwise Mutual Information (NPMI) between all pairs of tags [1]. This measures the degree to which two tags are likely co-occur in the same context – here, image – and takes a maximum value of 1 if two tags always co-occur and a value of 0 if they occur as according to chance. We can use this metric as a similarity measure and compute a tag \times tag similarity matrix. This matrix is fed to the spectral clustering algorithm, a hard-clustering method appropriate for generating non-convex clusters [20, 30, 35], which performs a graph partitioning on the Laplacian of the similarity matrix. The number of clusters needs to be specified in advance, and, after an exploration and analysis of various numbers, we decided to use 400 clusters for the rest of the study. Once these clusters of semantically similar tags are created, we represent each image as a vector containing the normalized number of times a member of each tag cluster is detected. For each user, we derive a feature vector of *image content topics* as the normalized number of times each topic cluster is present in the tweet-embedded images.

Secondly, we use a pre-trained version of the 19-layer **VGG-Net** image classifier based on convolutional neural networks [32], to predict a class probability for 1,000 objects in the ImageNet tagset. However, as this tagset does not contain categories of core interest for our analysis (e.g., faces, building interiors, sports), we use it for personality prediction experiments only.

4 ANALYSIS

To uncover the image features associated with each personality trait, we perform univariate correlation tests between each feature and personality trait. We control for age and gender effects using

r -post	r -like	Image Tag Clusters	r -post	r -like	Image Tag Clusters
Openness (+)			Openness (-)		
0.155 (1)	-	senior, old, mature, elderly, husband, grandma, retired, grandfather, retirement	0.207 (1)	0.135 (2)	ball, sports equipment, game equipment, basketball, basketball equipment...
0.148 (2)	0.112 (1)	art, cartoon, clipart	0.195 (2)	0.138 (1)	player, athlete, baseball, contestant, base, baseball equipment, ballplayer...
0.137 (3)	0.097 (3)	drawing, representation, diagram	0.154 (3)	0.109 (3)	football, helmet, back, football_helmet, crash_helmet
0.120 (4)	-	ancient, palace, castle, historic	0.119 (4)	0.090 (4)	game, puzzle, jigsaw_puzzle
0.119 (5)	0.079 (7)	decoration, sketch, tattoo, graffiti	0.117 (5)	-	clothing, garment, shirt, consumer_goods, jersey, sleeve
0.108 (6)	0.090 (4)	office, businessman, professional, corporate, executive, suit, handsome, men	0.109 (6)	-	tennis, racket
0.102 (7)	0.101 (2)	artwork	0.105 (7)	-	runner, diversion, track
-	0.089 (5)	pattern, texture, wallpaper, backdrop, shape, curve	0.099 (8)	0.078 (5)	child, boy, kid, little, childhood, baby, children
Conscientiousness (+)			Conscientiousness (-)		
0.153 (1)	0.076 (1)	classroom, whole	0.125 (1)	0.049 (5)	face, pretty, hair, model, fashion, sexy, lady, brunette, glamour
0.132 (2)	0.069 (2)	office, businessman, professional, corporate, executive, suit, handsome, men	0.111 (2)	0.086 (2)	art, cartoon, clipart
0.114 (3)	-	building, architecture, city	0.105 (3)	0.065 (4)	drawing, representation, diagram
0.104 (4)	-	structure, fountain, gas	0.101 (4)	0.084 (3)	expression, looking, beard, head
0.101 (5)	-	business	0.101 (5)	-	cute, eyes
0.099 (6)	0.059 (5)	paper, document, writing, menu, pen, fare, book	0.097 (6)	0.1 (1)	artwork
0.094 (7)	-	furniture, table	0.072 (7)	-	computer, equipment, technology, display, screen, monitor
0.078 (8)	0.062 (4)	home, room, house, interior	0.068 (8)	-	people, person, adult, caucasian, man, male, happy, portrait, attractive, smile...
Extraversion (+)			Extraversion (-)		
0.321 (1)	0.128 (1)	people, person, adult, caucasian, man, male, happy, portrait, attractive, smile...	0.253 (1)	0.164 (1)	art, cartoon, clipart
0.243 (2)	0.089 (3)	face, pretty, hair, model, fashion, sexy, lady, brunette, glamour	0.213 (2)	0.122 (2)	drawing, representation, diagram
0.182 (3)	-	women, group, friends, friend, girls, friendship, buddy	0.165 (3)	-	design, sign, icon, graphic, symbol, internet, set, web, button, icons
0.175 (4)	0.080 (4)	lifestyle	0.149 (4)	-	casual, silhouette, sport, laptop, covering, businessperson, element, light...
0.171 (5)	0.075 (5)	happiness, couple, together, love, fun, two, family	0.128 (5)	0.066 (6)	black, african, picture, dark
0.136 (6)	0.093 (2)	disco, cabaret, spot, ballroom	0.127 (6)	0.103 (3)	artwork
0.127 (7)	-	body, swimsuit, slim, bikini, maillot, tights	0.125 (7)	-	text, 3d
0.120 (8)	-	expression, looking, beard, head	0.094 (8)	0.70 (4)	color, motion, futuristic
Agreeableness (+)			Agreeableness (-)		
0.144 (1)	0.090 (1)	happiness, couple, together, love, fun, two, family	0.211 (1)	0.096 (1)	office, businessman, professional, corporate, executive, suit, handsome, men
0.108 (2)	0.080 (7)	performance, concert, stage, platform, musician	0.177 (2)	0.076 (3)	press, print_media
0.098 (3)	0.075 (5)	flower, floral, garden, flowers, petal, blossom, flora, bloom	0.132 (3)	0.077 (2)	business
0.094 (4)	-	cat, feline, pet, domestic, fur, kitten, kitty, pets, domestic_cat, whiskers, furry...	0.129 (4)	-	newspaper, money, currency, finance, cash, bank, banking, financial, bill
0.094 (5)	-	trees, season	0.101 (5)	-	work, success, manager, confident, successful
0.093 (6)	-	women, group, friends, friend, girls, friendship, buddy	0.099 (6)	-	signboard, billboard, scoreboard, street_sign
0.089 (7)	-	decoration, sketch, tattoo, graffiti	0.092 (7)	-	people, person, adult, caucasian, man, male, happy, portrait, attractive, smile...
0.081 (8)	0.081 (2)	music, guitar, stringed_instrument	0.081 (8)	-	publication, magazine, book_jacket, comic_book, jacket, wrapping
Neuroticism (+)			Neuroticism (-)		
0.107 (1)	0.057 (7)	paper, document, writing, menu, pen, fare, book	0.108 (1)	-	casual, silhouette, sport, laptop, covering, businessperson, element, light...
0.104 (2)	0.071 (3)	happiness, couple, together, love, fun, two, family	0.096 (2)	-	computer, equipment, technology, display, screen, monitor
0.104 (3)	0.77 (2)	cute, eyes	0.094 (3)	-	device, machine, slot_machine, slot, vending_machine
0.099 (4)	0.082 (1)	face, pretty, hair, model, fashion, sexy, lady, brunette, glamour	0.064 (4)	-	design, sign, icon, graphic, symbol, internet, set, web...
0.092 (5)	-	animal, dog, domestic_animal, canine	0.063 (5)	-	art, cartoon, clipart
0.082 (6)	0.058 (6)	cat, feline, pet, domestic, fur, kitten, kitty, pets, domestic_cat, whiskers, furry...	0.060 (6)	-	communication, telephone, phone, mobile, call, cellular_telephone
0.079 (7)	0.061 (5)	people, person, adult, caucasian, man, male, happy, portrait, attractive, smile...	0.057 (7)	-	space, digital, fractal, laser, optical_device, glow, render
0.060 (8)	0.062 (4)	sporting_dog, retriever, golden_retriever, labrador_retriever...			

Table 2: \mathcal{D}_2 : Pearson correlations and feature rank between personality trait and image tag clusters extracted from images embedded in posted tweets (r -post) and liked tweets (r -like) (maximum 8 per personality trait). All correlations are significant at $p < .01$, Simes corrected, two-tailed t-test. Results for every personality trait are controlled for age and gender and other four corresponding personality traits. Tags are sorted by occurrence in our data set within a tag cluster.

partial correlation so that our analysis is not skewed by demographic biases. For robustness, we perform Simes p-correction to account for running multiple significance tests [31]. Additionally, in content analysis, as some personality traits are inter-correlated (e.g. extraversion and agreeableness), we control for all other four traits in order to isolate the peculiarities of each personality dimension.

Results of image color analysis on the \mathcal{D}_2 data set are presented for posted and liked images separately in Tables 1.A and 1.B respectively. The content analysis results on posted and liked images are presented in Table 2. In Table 2, the correlations for the top-k topics are shown along with the rank associated with each topic for the trait in brackets. The same set of experiments on personality correlations was conducted on the \mathcal{D}_1 data set, but the correlations are not robust to randomness due to the small sample size. This shows the need for this behavior to be studied using larger samples. In Tables 1.A and 1.B we included an additional feature coding the percentage of posts that contained people, identified through the presence of a curated list of Imagga Clusters in the top 10 predicted ones (46.4% of the total number of images). Another feature coded the raw number of images posted or liked by a user.

Openness to Experience. Color analysis shows openness is associated with posting and liking images in grayscale and having low hue count, which indicate artistic and well-composed images specific of this trait. Users high in this trait also prefer and post less bright images and post a greater number of images, while liking fewer. Image content analysis (Table 2) reveals two major themes for users high in openness to experience. First, posting images with drawings (*art, drawing, decoration, artwork* – clusters denoted by their first tag) indicates engagement with aesthetic domains, a known characteristic of openness to experience [4] and especially abstract art [5]. Secondly, these users post fewer images containing people and focus more on objects, but when people are present they represent seniors or businessmen (*senior, office*). Users low in openness prefer posting images of sports or games (*player, football, game, tennis*), which are conventional interests in the US.

Conscientiousness. Conscientiousness is a trait associated with orderliness and preference for planned behaviour. Users high in this trait do not prefer grayscale images and prefer photos of people rather than objects, suggesting a preference for generally normative and conventional posting behavior. Conscientious users are

also the most cautious in liking images from other people, perhaps because this action is to a degree impulsive, which is atypical of conscientiousness. Content analysis reveals images containing either landscapes (*building, structure*) or formal, office environments (*classroom, office, paper, home*), the latter related to overall better job performance of users high in conscientiousness [13]. On the other hand, low conscientiousness is portrayed by images of drawings (*art, drawing, artwork*), which again shows a dislike for unconventional posts, as well as a specific type of close-up face images (*face, expression, cute*).

Extraversion. Extraversion is associated with both posting and liking images that are bright, saturated and more chromatically complex as revealed through a high hue-count. This trait is characterized by engagement with other people and, in accordance to theory, users who score high are specifically interested in posting and liking images containing people in general as well as a broad range of types of people (*people, face, women, happiness*). In addition, extraverts are perceived as being energetic and this is reflected by images portraying active interests (*disco*). Introversion is portrayed through drawings (*art, drawing*) or abstract symbols (*design, casual*) and by posting fewer images in general.

Agreeableness. Agreeableness shows similar patterns to extraversion in color associations, which is expected based on their common positive emotion core. Users high in agreeableness are characterized by kindness, generosity and optimism, and their images show a breadth of aspects of the good life: friends (*happiness, women*), pets (*cat*), flowers (*flower*) or music (*performance*). Liked images also reveal similar preferences for music (*music*). On the other hand, low agreeableness is associated with images of office environments (*office, work, business*), text-related content (*press, newspaper, signboard*) and fewer photos of people overall.

Neuroticism. Users high in neuroticism post grayscale images and their images are less saturated and less chromatic diversity by containing fewer hues. Perhaps surprisingly, neuroticism is associated with posting more images, although images of objects rather than people are preferred. Most striking is the preference for liking images from others, which is a more passive way of engagement. Through content analysis, we observe that high neuroticism is disclosed through posting images containing documents (*paper*) or animals (*cat, animal*). Although they are less likely to post pictures of people overall, they favor close-ups of people (*face, cute, people*). Users low in neuroticism are characterized by posting images of art and abstract symbols (*art, design*) and tech items (*casual, device*). Users low in neuroticism had no significant association with liked image content, perhaps caused by the tendency to like more images and, hence, more heterogeneous content. Overall, we highlight that in almost all traits, similar image content tendencies are specific of both posted and liked images.

5 PREDICTION

We study the predictive power of image features in the task of predicting user personality traits. We treat this as a regression problem to which we apply a linear regression with Elastic Net regularization. As personality is text predicted on the \mathcal{D}_2 data set, we only perform prediction using text on the \mathcal{D}_1 data set where

Feature	# Feat.	Ope	Con	Ext	Agr	Neu
Colors [6, 33]	33	.093	.147	.022	.229	.085
Imagga Clusters	400	.044	.077	.183	.085	.404
VGG-net Classes [32]	1000	.056	.051	.061	.037	.222
Text	100	.168	.059	.223	.111	.261
All (Image)	3	.081	.177	.187	.229	.416
All (Image+Text)	2	.171	.178	.223	.230	.418

Table 3: \mathcal{D}_1 : Prediction results for personality traits with all features. Performance is measured using Pearson correlation in 10-fold cross-validation.

Feature	# Feat.	Ope	Con	Ext	Agr	Neu
Colors [6, 33]	33	.284	.352	.293	.317	.398
Imagga Clusters	400	.275	.364	.317	.221	.383
VGG-net Classes [32]	1000	.410	.383	.319	.198	.398
All (Image)	3	.448	.479	.369	.336	.503

Table 4.A Posted Images.

Feature	# Feat.	Ope	Con	Ext	Agr	Neu
Colors [6, 33]	33	.351	.383	.229	.286	.396
Imagga Clusters	400	.411	.492	.345	.335	.412
VGG-net Classes [32]	1000	.302	.388	.193	.008	.374
All (Image)	3	.468	.530	.366	.378	.467
Posts + Likes (Image)	2	.543	.566	.440	.433	.530

Table 4.B Liked Images.

Table 4: \mathcal{D}_2 : Prediction results for personality traits with all features on posted and liked images. Performance is measured using Pearson correlation in 10-fold cross-validation.

we have collected survey-based personality scores. We extracted word2vec [19] clusters from the text in \mathcal{D}_2 following [23]. Since VGG and Imagga Clusters are high-dimensional features, we employed PCA for dimensionality reduction by setting the explained variance ratio threshold as 95% for \mathcal{D}_2 , and choosing 51 principal components for \mathcal{D}_2 (due to the small size of the data set). We have also experimented with non-linear methods (Support Vector Machines with RBF kernel), but the results did not improve significantly.

To evaluate our results, we split our data into 10 folds and performed cross-validation on one held-out fold at a time. For all our methods, we tune the parameters of all our models on a separate validation fold. The overall performance is assessed using Pearson correlation of the predicted value to the self-reported score. Results are presented in Tables 3 and 4. Feature combination is performed using a linear ensemble over the individual prediction scores of each feature set. The same patterns hold when evaluating the results with Root Mean Squared Error and we omit them for clarity.

On the \mathcal{D}_1 data set, features extracted from images perform better than the text-derived features for conscientiousness, agreeableness and neuroticism. For the other traits, the text features perform better. It is to be noted that the small size of the data set might be acting as an artifact. To confirm the contribution of features from text and images, larger data sets with survey-based labels need to be acquired.

On the \mathcal{D}_2 data set, for image posts, image content features extracted from either VGG-Net or Imagga performed better than color features in all but one cases. Combining all features using an ensemble works better in predicting all traits, showing colors and content are to some extent complimentary. For image likes, Imagga features outperform VGG-Net classes and color features for all personality traits. Also, colors perform better than VGG-Net classes on all traits except for conscientiousness. By combining the features from the two different groups of images – posted and liked – we obtain significant performance gains (6% for conscientiousness and 10-15% for the remaining traits) over either image posts or likes considered independently. This suggests that both interactions that users have on social media – namely posting images and liking images reveal a more complete picture of users’ personality.

Prediction performance using images is overall high, especially for openness and neuroticism – even if excluding from text personality prediction the tweets containing images. Combined with results from the \mathcal{D}_1 data set, this indicates that the two modalities disclose both overlapping and complimentary information. However, a larger data set with survey-based personality would further reveal what can be captured through each modality.

6 CONCLUSION

We analyzed image posting and liking preferences using interpretable aesthetic and semantic features using a data set an order of magnitude larger than previous work, with the aim of identifying the way and extent to which they can be related to personality differences measured using the Big Five model. We created tag clusters to analyze the content of images and identified idiosyncrasies of each personality trait, while controlling for all other four traits to isolate the peculiarities of each personality dimension. Finally, predictive models of personality traits show reliable accuracy on held-out data using image features even when text does not. Also, combining the signal from both posted and liked images leads to significant performance gain compared to individual interactions.

ACKNOWLEDGMENTS

The authors acknowledge the support of the Templeton Religion Trust, grant TRT-0048.

REFERENCES

- [1] Gerlof Bouma. 2009. Normalized (Pointwise) Mutual Information in Collocation Extraction. In *Biennial GSCL Conference*.
- [2] Anne Burdick, Johanna Drucker, Peter Lunenfeld, Todd Presner, and Jeffrey Schnapp. 2012. *Digital Humanities*. The MIT Press.
- [3] Paul T Costa and Robert R MacCrae. 1992. *Revised NEO Personality Inventory and NEO Five-factor Inventory: Professional Manual*. Psychological Assessment Resources.
- [4] Colin DeYoung, Lena Quilty, and Jordan Peterson. 2007. Between Facets and Domains: 10 aspects of the Big Five. *JPSP* (2007).
- [5] Gregory J Feist and Tara R Brady. 2004. Openness to Experience, Non-Conformity, and the Preference for Abstract Art. *Empirical Studies of the Arts* 22, 1 (2004).
- [6] Bruce Ferwerda, Markus Schedl, and Marko Tkalcic. 2016. Using Instagram Picture Features to Predict Users’ Personality (*MMM*).
- [7] Venkata Rama Kiran Garimella, Abdulrahman Alfayad, and Ingmar Weber. 2016. Social Media Image Analysis for Public Health (*CHI*).
- [8] Samuel D Gosling, Sei Jin Ko, Thomas Mannarelli, and Margaret E Morris. 2002. A Room with a Cue: Personality Judgments based on Offices and Bedrooms. *JPSP* (2002).
- [9] Sharath Chandra Guntuku, Weisi Lin, Michael James Scott, and Gheorghita Ghinea. 2015. Modelling the Influence of Personality and Culture on Affect and Enjoyment in Multimedia (*ACII*).

- [10] Sharath Chandra Guntuku, Lin Qiu, Sujoy Roy, Weisi Lin, and Vinit Jakhethiya. 2015. Do others Perceive You as You want them To?: Modeling Personality based on Selfies. In *Workshop on Affect & Sentiment in Multimedia (MM)*.
- [11] Sharath Chandra Guntuku, Joey T Zhou, Sujoy Roy, Lin Weisi, and Ivor W Tsang. 2016. Who likes what, and why? Insights into Personality Modeling based on Image ‘Likes’. *IEEE Transactions on Affective Computing* (2016).
- [12] Kai-Qi Huang, Qiao Wang, and Zhen-Yang Wu. 2006. Natural Color Image Enhancement and Evaluation Algorithm Based on Human Visual System. *CVIU* (2006).
- [13] Gregory M Hurtz and John J Donovan. 2000. Personality and Job Performance: The Big Five Revisited. *Journal of Applied Psychology* 85, 6 (2000).
- [14] Vasileios Lamos, Nikolaos Aletras, Daniel Preotjuc-Pietro, and Trevor Cohn. 2014. Predicting and Characterising User Impact on Twitter (*EACL*).
- [15] Leqi Liu, Daniel Preotjuc-Pietro, Zahra Riahi Samani, Mohsen E. Moghaddam, and Lyle Ungar. 2016. Analyzing Personality through Social Media Profile Picture Choice (*ICWSM*).
- [16] Robert R McCrae and Oliver P John. 1992. An Introduction to the Five-factor Model and its Applications. *Journal of Personality* (1992).
- [17] Albert Mehrabian and James A Russell. 1974. *An Approach to Environmental Psychology*. The MIT Press.
- [18] GJ Meyer et al. 2001. Psychological Testing and Psychological Assessment: A Review of Evidence and Issues. *American Psychologist* (2001).
- [19] T Mikolov and J Dean. 2013. Distributed Representations of Words and Phrases and their Compositionality (*NIPS*).
- [20] Andrew Y. Ng, Michael I. Jordan, and Yair Weiss. 2002. On Spectral Clustering: Analysis and an Algorithm (*NIPS*).
- [21] Jie Nie, Lei Huang, Peng Cui, Zhen Li, Yan Yan, Zhiqiang Wei, and Wenwu Zhu. 2016. Social Media Profiler: Inferring Your Social Media Personality from Visual Attributes in Portrait. In *Pacific Rim Conference on MM*.
- [22] Gregory Park, H. Andrew Schwartz, Johannes Eichstaedt, Margaret Kern, Michal Kosinski, David Stillwell, Lyle Ungar, and Martin Seligman. 2015. Automatic Personality Assessment through Social Media Language. *JPSP* (2015).
- [23] Daniel Preotjuc-Pietro, Vasileios Lamos, and Nikolaos Aletras. 2015. An Analysis of the User Occupational Class through Twitter Content (*ACL*).
- [24] Daniel Preotjuc-Pietro, Svitlana Volkova, Vasileios Lamos, Yoram Bachrach, and Nikolaos Aletras. 2015. Studying User Income through Language, Behaviour and Affect in Social Media. *PLoS ONE* (2015).
- [25] Jitao Sang, Huaiwen Zhang, and Changsheng Xu. 2016. Visual BFI: An Exploratory Study for Image-Based Personality Test (*PCM*).
- [26] H Andrew Schwartz, Johannes C Eichstaedt, Margaret L Kern, Lukasz Dziurzynski, Stephanie M Ramones, Megha Agrawal, Achal Shah, Michal Kosinski, David Stillwell, and Martin Seligman. 2013. Personality, Gender, and Age in the Language of Social Media: The Open-vocabulary Approach. *PLoS ONE* (2013).
- [27] Michael James Scott, Sharath Chandra Guntuku, Weisi Lin, and Gheorghita Ghinea. 2016. Do personality and culture influence perceived video quality and enjoyment? *IEEE Transactions on Multimedia* 18, 9 (2016), 1796–1807.
- [28] Cristina Segalin, Dong Seon Cheng, and Marco Cristani. 2016. Social Profiling through Image Understanding: Personality Inference using Convolutional Neural Networks. *Computer Vision and Image Understanding* (2016).
- [29] Cristina Segalin, Alessandro Perina, Marco Cristani, and Alessandro Vinciarelli. 2016. The Pictures we like are our Image: Continuous Mapping of Favorite Pictures into Self-Assessed and Attributed Personality Traits. *IEEE Transactions on Affective Computing* (2016).
- [30] Jianbo Shi and Jitendra Malik. 2000. Normalized Cuts and Image Segmentation. *Transactions on Pattern Analysis and Machine Intelligence* (2000).
- [31] R John Simes. 1986. An Improved Bonferroni Procedure for Multiple Tests of Significance. *Biometrika* (1986).
- [32] Karen Simonyan and Andrew Zisserman. 2015. Very Deep Convolutional Networks for Large-Scale Image Recognition (*ICLR*).
- [33] Marcin Skowron, Marko Tkalcic, Bruce Ferwerda, and Markus Schedl. 2016. Fusing Social Media Cues: Personality Prediction from Twitter and Instagram (*WWW Companion*).
- [34] Patricia Valdez and Albert Mehrabian. 1994. Effects of Color on Emotions. *Journal of Experimental Psychology* (1994).
- [35] Ulrike von Luxburg. 2007. A Tutorial on Spectral Clustering. *Statistics and Computing* (2007).
- [36] Xiaoyu Xiong, Maurizio Filippone, and Alessandro Vinciarelli. 2016. Looking Good With Flickr Faves: Gaussian Processes for Finding Difference Makers in Personality Impressions (*MM*).
- [37] Yan Yan, Jie Nie, Lei Huang, Zhen Li, Qinglei Cao, and Zhiqiang Wei. 2015. Is your First Impression Reliable? Trustworthy Analysis using Facial Traits in portraits (*MMM*).
- [38] Huaiwen Zhang, Jiaming Zhang, Jitao Sang, and Changsheng Xu. 2017. A Demo for Image-Based Personality Test (*MMM*).