

# How Do People Differ? A Social Media Approach

Vincent Wong and Yaneer Bar-Yam

*New England Complex Systems Institute*

*210 Broadway, Suite 101, Cambridge, MA 02139, USA*

(Dated: August 9, 2017)

## Abstract

Research from a variety of fields including psychology and linguistics have found correlations and patterns in personal attributes and behavior, but efforts to understand the broader heterogeneity in human behavior have not yet integrated these approaches and perspectives with a cohesive methodology. Here we extract patterns in behavior and relate those patterns together in a high-dimensional picture. We use dimension reduction to analyze word usage in text data from the online discussion platform Reddit. We find that pronouns can be used to characterize the space of the two most prominent dimensions that capture the greatest differences in word usage, even though pronouns were not included in the determination of those dimensions. These patterns overlap with patterns of topics of discussion to reveal relationships between pronouns and topics that can describe the user population. This analysis corroborates findings from past research that have identified word use differences across populations and synthesizes them relative to one another. We believe this is a step toward understanding how differences between people are related to each other.

## I. INTRODUCTION

Since the advent of social media, abundant public data has been produced, and there has been a concerted research effort to understand the people behind the data. Previous studies have adopted different perspectives from disparate fields, including Natural Language Processing (NLP) and personality psychology. Work in sentiment analysis focuses on accurately extrapolating emotional content from text [1–6]. Other NLP methods involve algorithms that determine topics of conversation [7, 8] or demographics from language [9–11]. On the other hand, approaches from personality psychology incorporate predetermined personality traits to obtain correlations between word usage, demographics such as gender, and personality types [6, 12–18]. A number of studies have used social media data including Facebook and Twitter [16, 18–20]. Furthermore, research outside of these fields uses social media data to understand the dynamic behaviors of users [21, 22] or to detect medical conditions in individuals, including mental illnesses [23, 24] and heart disease [25].

A challenge that remains unaddressed is the connection between linguistics and personality beyond specific correlations of predefined variables. As a first step in this direction, we use multidimensional scaling (MDS) to investigate differences in individuals by characterizing the high dimensional space of their word usage. We calculate vectors of word counts (feature vectors) and reduce the dimensionality of the space of these vectors to observe patterns. The most important dimensions best capture the differences between individuals as combinations of word usage choices. We investigate content and characteristics of user posts to gain an understanding of what the reduced dimensions reveal about the behavior of users. This data-driven approach is unsupervised and requires no prior assumptions about the axes of personality traits. We analyze data from Reddit, the content rating and discussion website.

Our analysis finds similar correlations as previous research, provides new insights, and relates them to each other. Fig. 1 shows the set of users plotted in the two most important dimensions. Colors show 1st person pronoun use as an example of patterns of differences among them. By looking at multiple ways individuals differ in word usage across these two dimensions, we can identify the codependencies among different behavioral traits, including those that may be associated with personality [6, 15–18]. This is a preliminary study in using this type of analysis on social media data.

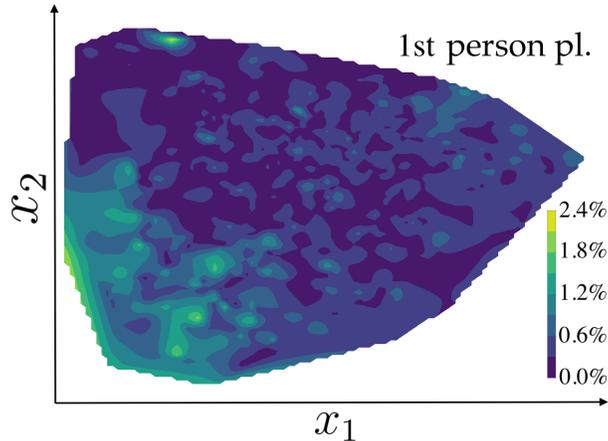


FIG. 1: **Differences in user pronoun use** plotted in the space of the two most prominent dimensions in which individuals differ from each other in word usage ( $x_1, x_2$ ). The region that is colored is the area of the space that is populated by individuals. Color represents the percent of each user’s dictionary that is first person plural pronouns (‘we’, ‘us’, ‘our’, ‘ours’, ‘ourselves’, etc). Areas in the bottom left have a much higher use of these pronouns. Pronouns were not used in the determination of the prominent dimensions of variation so the variation in pronoun use is linked to other word use variations.

## II. METHODS

From a dataset of Reddit posts, we obtained a sample of 1500 users who had made posts or comments at least 50 times in the month of January 2015 [26]. For each user, we created a user dictionary consisting of the counts of each word they used across all of their posts in that month. We determined words by space delineation, and we removed all punctuation as well as Reddit-specific text formatting (e.g. the use of asterisks to make text bold). Common words which primarily serve grammatical function and do not contribute to topical content—known as stop words in computational linguistics—were removed from the dataset (see Appendix A). We construct “feature vectors” for each user where each dimension corresponds to a word in the total dictionary of words used by all users, and nonzero entries are counts of words from each user dictionary.

We apply Multidimensional Scaling (MDS) to the set of feature vectors using Torgerson-Gower scaling. We generate a matrix of Euclidean distances between user vectors pairwise for all users, resulting in an  $N \times N$  matrix where  $N$  is the number of users and the diagonal of the matrix is zero. We plot the users in the space of the eigenvectors with the largest eigenvalues. These composite dimensions can be considered to best characterize the users

across the space of word usage.

When plotting users this way, bots were revealed because of their strong differentiation from human users. We removed bots from manually selected regions that show extreme separation from other users. We confirmed that they were bots by inspection. We used the resulting cropped list of 1462 users for our analysis.

This method is similar to the factor analysis used in identifying the Five Factor model in personality psychology [12–14, 27]. Unlike work done in the framework of that model, the data used here does not restrict itself to a predetermined set of adjectives. It differs also in the use of social media data rather than survey data.

To investigate the space of users in the reduced dimensions, we compiled statistical measures on each user’s dictionaries. We use the following measures: dictionary entropy, unique word count, and percent contribution of pronouns by their person, gender, and number (first-, second-, and third-person singular and first- and third-person plural).

We calculated the dictionary entropy as

$$S_u = - \sum_w \frac{n_w}{\sum_w n_w} \log\left(\frac{n_w}{\sum_w n_w}\right) \quad (1)$$

where  $w$  is an index labeling each word and  $n_w$  is the count of that word in the user dictionary. The dictionary entropy can be interpreted as the word diversity. An investigation of raw word counts yielded similar results.

To determine the topics discussed in the data, we use Latent Dirichlet Allocation (LDA) because it maps directly to the bag-of-words approach used for feature vectors. LDA assumes that each user feature vector (also called a “document”) is generated probabilistically from a distribution of topics, and attempts to derive these topics. The same word can appear in multiple topics. We used collapsed Gibbs sampling to redistribute words to topics given their presence in documents. We generated 100 topics of 40 words each using user feature vectors as documents over 1000 iterations of LDA, with stop words removed. We calculate the percent contribution of a topic to each user’s dictionary based on what words appear in topics. We labeled topics based on their content for ease of reference where possible.

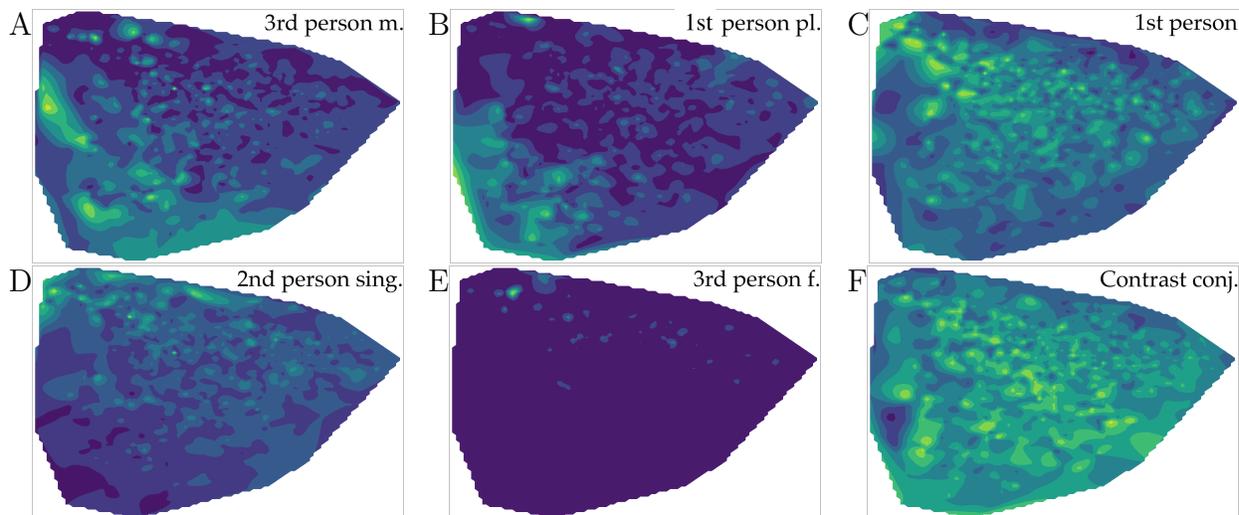


FIG. 2: **Word use metrics plotted in the space of the two most prominent dimensions of variation among individuals.** Each plot shows the same region of space colored by a different metric. Areas of the space with higher levels of usage of pronouns or other words are shown by lighter colors (color bar as in Fig. 1). A. Male third-person pronouns (‘he’, ‘his’, ‘him’, ‘himself’). B. First person plural pronouns (‘we’, ‘us’, ‘our’, ‘ourselves’). C. First person pronouns (‘I’, ‘me’, ‘myself’, ‘my’). D. Second person singular pronouns (‘you’, ‘your’, ‘yours’, ‘yourself’). E. Female third-person pronouns (‘she’, ‘her’, ‘hers’, ‘herself’). Third person female pronouns are underrepresented in Reddit because the site is male-dominated. F. An example of the analysis of another category of words, here contrasting conjunctions (‘but’, ‘though’, ‘however’).

### III. RESULTS

We constructed a reduced dimensional representation of word usage that captured the primary variations across user dictionaries excluding common words known as “stop words” (see Methods). The locations of users in the two most prominent dimensions of variation are shown in Fig. 2. We investigated the individual differences in use of pronoun types which may have implications for variations in personality [6, 15–18]. The colors in the first five panels show the percent of types of pronouns in user dictionaries. Since we are constructing our understanding directly from data, even elementary observations are important. We can see that people have significant differences in their use of pronouns. Moreover, individuals in different areas of the reduced space tend to have similar pronoun use. The bottom and left regions have a greater incidence of third-person male pronouns (‘he’, ‘his’, ‘himself’). The lower left itself has greater incidence of first-person plural pronouns (‘we’, ‘ours’, ‘ourselves’). The middle and upper left areas have a greater use of singular first-person pronouns (‘me’, ‘myself’, ‘I’). The top and top left has greater use of second-person

pronouns (‘you’, ‘yourself’). The few users with high use of third-person female pronouns (‘she’, ‘her’, ‘herself’) are primarily near the top.

The observation that pronoun use varies systematically with the dimensions that capture the greatest variance has important implications for our understanding of individual differences. In particular, the primary ways in which individuals differ include significant differences in pronoun use. It is worth noting that pronouns were included among the stop words and were removed from the dimensional reduction process, so the differences in pronoun use were not included in determining these dimensions. The observation of systematic variation implies that the other words people use are linked to pronoun use across the dimensions.

The spatial structures of the variations reflect not just individual or correlational changes in pronoun use, but also the codependency among these usages. For example, first-person plural pronouns (Fig. 2B) are high in a subset of areas of third-person male pronouns (Fig. 2A). Third-person female pronouns (Fig. 2E) appear to be localized in a subset of areas of high second-person singular pronoun use (Fig. 2D). Third-person male and first-person plural pronouns (Fig. 2A and B) occur in areas distinct from second-person singular and third-person female pronouns (Fig. 2D and E). However, the area of high first person pronoun use (Fig. 2C) overlaps with both groups but is distinct from them.

In the last panel of Fig. 2, we show the use of contrasting conjunctions (‘but’, ‘though’, ‘however’). Contrasting conjunctions are widely distributed across two dimensions. However, there is a small patch on the left just below center where usage is substantially lower. Correlations between the two primary dimensions of variation and the measures of contrasting conjunctions would not reveal this behavior. They would fail to capture the spatial variation of the vertical dimension in the middle left of Fig. 2F because of aggregation across the horizontal dimension as well as the up-down-up variation of the left region. They would also fail to capture the variations in the horizontal dimension due to aggregation in the vertical. The highly nonlinear variations across the space would not be captured in correlations in either dimension.

We investigated the words people use (not including stop words) and identified topics of their conversation and how they vary across users. We used Latent Dirichlet Allocation to extract 100 topics of conversation with 40 words associated with each topic (see Methods). We found the percent contribution of each topic to each user dictionary. The percent contri-

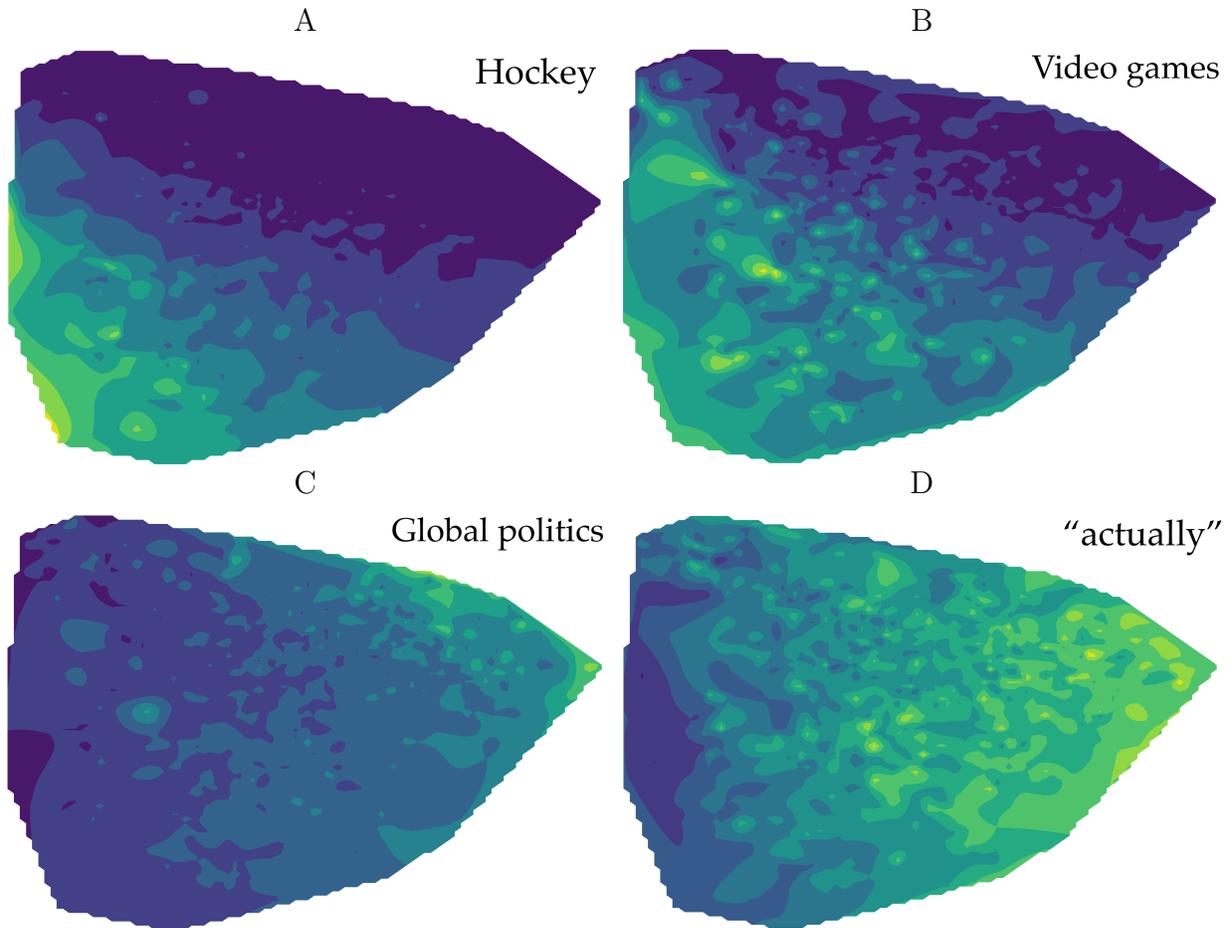


FIG. 3: **Topics of discussion in the space of the two most prominent dimensions of word usage differences** A. Hockey-related words (‘NHL’, ‘hockey’, ). B. Video game related words (‘game’, ‘enemy’, ‘kill’, ‘points’). Both A and B exemplify the way that many of the sport/game related topics extracted with LDA coincide in the population. C. Global politics (‘world’, ‘politics’, ‘government’, ‘money’) D. An assortment of common words determined by LDA to be a distinct topic labeled by the word “actually.” This analysis shows that this word group is meaningful even if it is non-trivial to determine the concept underpinning it.

bution of topics across users in the two most prominent dimensions of word usage variation is shown in Fig. 3 for each of four example topics named by inspection of representative words for that topic. The “hockey” topic contains words related to the sport of hockey, such as ‘hockey’, ‘nhl’, ‘puck’, ‘ice’, ‘bruins’, and ‘canucks’. The “video games” topic has words such as ‘damage’, ‘enemy’, ‘items’, ‘attack’, and ‘armor’. The topic “global politics” is characterized by words related to interactions between states at the global scale, such as ‘war’, ‘religion’, ‘people’, ‘money’, and ‘government’. The fourth topic, “actually”, is labeled by a prominent word from the topic because a clear identifier for the topic was not readily

apparent, and it contains the words ‘actually’, ‘very’, ‘more’, ‘anything’, and ‘lot’. (Full topics in Appendix B)

The topic “hockey” occurs more frequently in the bottom left area of the reduced dimensional space (Fig. 3A). Words related to video games occur in the bottom and left half (Fig. 3B). Words related to global politics occur in the dictionaries of users in the rightmost area (Fig. 3C). The topic labeled by the word “actually” is predominantly found in the users toward the right region.

The bottom and left areas that have a greater use of words from sport- and game-related topics also corresponds with a greater use of male third-person pronouns (Fig. 2A). This may indicate that people who discuss sports and games are also the people who refer to men most often, which may be due to the fact that both sports and video games are male-dominated subjects. In addition, these users have a greater use of first-person plural pronouns (Fig. 2B), which may be linked to a greater incidence of discussions about teamwork.

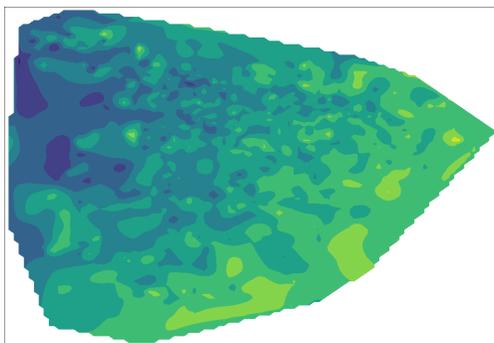


FIG. 4: **Entropy of user dictionaries** plotted as a function of the two most prominent dimensions of individual differences.

Finally, in Fig. 4 we show the entropy of user dictionaries as a function of the two most prominent dimensions of word usage variation. The entropy is higher in the lower right region of the reduced space, indicating that those users have more varied dictionaries.

All results shown in Fig. 2 - 4 are statistically significant with the exception of the use of third-person female pronouns in Fig. 2E. The statistical significance was determined by randomizing the metric values across users and comparing this to the original results with the null hypothesis that the results are unrelated to the position in the reduced space. We calculate the distribution of variances from a sliding window across the interpolated space and applied an independent t-test to compare the variances between the random and

original measures. We find a  $p < 0.001$  for all measures investigated, except third-person female pronouns which failed to reject the null hypothesis ( $p \approx 0.5$ ).

#### IV. CONCLUSIONS

We obtained relationships between the behavior of people from the words that they use by applying dimensional reduction to posts from the discussion site Reddit. We found the two most prominent dimensions that captured the greatest variation in word use across user dictionaries. We isolated nonlinear patterns in the reduced space by comparing the pronouns people use, the topics people discussed, and the entropy of their word use. By combining the set of differences among individuals into a few primary dimensions of variation we can identify how different behavioral traits are codependent. This strategy for integrating behavioral differences may have implications for the association of behavior with personality as well as other ways of characterizing individual differences. We found that users who discuss sports and games also used third-person male pronouns ('he') and first-person plural pronouns ('we') more often than other users in the population. Those users were distinct from groups of individuals who use second-person singular pronouns ('you') and those who discuss global politics. This is the beginning of a more complete characterization of the dimensions of differences in humans.

We greatly appreciate the advice of Alfredo Morales and Irving Epstein.

## APPENDIX

### A. Stop words

“i”, “i’m”, “i’ve”, “i’ll”, “i’d”, “you”, “your”, “you’re”, “he”, “his”, “him”, “she”, “her”, “hers”, “we”, “us”, “our”, “ours”, “they”, “their”, “they’re”, “them”, “my”, “me”, “it”, “its”, “it’s”, “the”, “a”, “an”, “no”, “yes”, “not”, “good”, “bad”, “well”, “be”, “am”, “is”, “are”, “was”, “were”, “been”, “wasn’t”, “aren’t”, “weren’t”, “being”, “isn’t”, “have”, “has”, “had”, “hadn’t”, “hasn’t”, “haven’t”, “having”, “do”, “does”, “did”, “doing”, “done”, “don’t”, “doesn’t”, “didn’t”, “can”, “can’t”, “couldn’t”, “could”, “shouldn’t”, “should”, “cannot”, “get”, “gets”, “got”, “gotten”, “getting”, “say”, “said”, “saying”, “will”, “would”, “wouldn’t”, “go”, “goes”, “gone”, “went”, “going”, “come”, “coming”, “came”, “make”, “makes”, “made”, “know”, “knows”, “knowing”, “knew”, “known”, “take”, “takes”, “taking”, “took”, “taken”, “see”, “sees”, “saw”, “seeing”, “seen”, “look”, “looks”, “looked”, “looking”, “think”, “thinks”, “thought”, “thinking”, “use”, “uses”, “using”, “used”, “want”, “wants”, “wanted”, “wanting”, “give”, “gives”, “giving”, “gave”, “given”, “to”, “of”, “in”, “that”, “for”, “on”, “with”, “as”, “at”, “from”, “by”, “than”, “about”, “so”, “like”, “into”, “only”, “and”, “but”, “or”, “if”, “just”, “also”, “even”, “because”, “up”, “down”, “left”, “right”, “there”, “out”, “over”, “back”, “way”, “one”, “all”, “other”, “any”, “these”, “those”, “most”, “time”, “now”, “then”, “after”, “before”, “new”, “day”, “who”, “what”, “where”, “why”, “how”, “which”, “that”, “that’s”, “this”, “too”

### B. Topics

- “Hockey”: ‘game’, ‘team’, ‘season’, ‘hockey’, ‘goal’, ‘leafs’, ‘adnan’, ‘nhl’, ‘tonight’, ‘puck’, ‘goalie’, ‘win’, ‘pp’, ‘fans’, ‘canucks’, ‘hae’, ‘net’, ‘players’, ‘ice’, ‘player’, ‘games’, ‘jay’, ‘kessel’, ‘penalty’, ‘teams’, ‘goals’, ‘fan’, ‘playing’, ‘gonna’, ‘bruins’, ‘period’
- “Video games”: ‘game’, ‘damage’, ‘play’, ‘team’, ‘player’, ‘games’, ‘level’, ‘playing’, ‘enemy’, ‘items’, ‘attack’, ‘kill’, ‘against’, ‘weapons’, ‘played’, ‘armor’, ‘skill’, ‘weapon’, ‘server’, ‘win’, ‘mana’, ‘deck’, ‘cards’, ‘lane’, ‘item’, ‘hit’, ‘card’, ‘hp’, ‘base’, ‘gold’, ‘power’, ‘support’, ‘map’, ‘hero’

- “Global politics”: ‘people’, ‘money’, ‘against’, ‘government’, ‘world’, ‘state’, ‘police’, ‘years’, ‘country’, ‘law’, ‘power’, ‘support’, ‘system’, ‘many’, ‘war’, ‘religion’, ‘part’, ‘american’, ‘never’, ‘america’, ‘nothing’, ‘evidence’, ‘cops’, ‘believe’, ‘pay’, ‘public’, ‘rights’, ‘states’, ‘society’, ‘political’, ‘human’, ‘children’, ‘laws’, ‘tax’, ‘article’, ‘kill’, ‘religious’
- “actually”: ‘people’, ‘more’, ‘some’, ‘when’, ‘things’, ‘very’, ‘something’, ‘much’, ‘same’, ‘actually’, ‘point’, ‘lot’, ‘different’, ‘need’, ‘someone’, ‘without’, ‘thing’, ‘anything’, ‘may’, ‘many’, ‘work’, ‘own’, ‘might’, ‘read’, ‘reason’, ‘kind’, ‘while’, ‘problem’, ‘person’, ‘etc’, ‘find’, ‘less’, ‘still’, ‘through’, ‘understand’, ‘enough’, ‘case’

- 
- [1] J. Pennebaker, M. Mehl, K. Niederhoffer, Psychological aspects of natural language use: Our words, our selves, *Annual Review Psychology* **54**: 547-577 (2003), doi: 10.1146/annurev.psych.54.101601.145041.
- [2] Y. Yang, P. Cui, W. Zhu, S. Yang, User interest and social influence based emotion prediction for individuals, *MM '13 Proceedings of the 21st ACM International Conference on Multimedia*, pp. 785-788, doi:10.1145/2502081.2502204.
- [3] P. Dodds, K. Harris, I. Kloumann, C. Bliss, C. Danforth, Temporal patterns of happiness and information in a global social network: Hedonometrics and Twitter, *PLOS ONE* (2011), <http://dx.doi.org/10.1371/journal.pone.0026752>.
- [4] K.Z. Bertrand, M. Bialik, K. Virdee, A. Gros, Y. Bar-Yam, Sentiment in New York City: A high resolution spatial and temporal view, *arXiv:1308.5010* (2013).
- [5] Y. Lewenberg, Y. Bachrach, S. Volkova, Using emotions to predict user interest areas in online social networks, *Data Science and Advanced Analytics (DSAA)* **36678** (2015), doi: 10.1109/DSAA.2015.7344887.
- [6] H. Andrew Schwartz, L. Ungar, Data-driven content analysis of social media: A systematic overview of automated methods, *ANNALS, AAPSS* **659** (2015), doi: 10.1177/0002716215569197.
- [7] W.X. Zhao, J. Jiang, J. Weng, J. He, E.-P. Lim, H. Yan, X. Li, Comparing Twitter and traditional media using topic models, *ECIR'11 Proceedings of the 33rd European conference*

- on *Advances in Information Retrieval* (2011), pp. 338-349.
- [8] L. Xiang, Z. Xu, Q. Yang, L. Ru, Discovering user interest on Twitter with a modified author-topic model, *EEE/WIC/ACM International Conference on Web Intelligence and Intelligent Agent Technology* **01**: 422-429 (2011), doi:10.1109/WI-IAT.2011.47.
- [9] D. Preotiuc-Pietro, S. Volkova, V. Lampos, Y. Bachrach, N. Aletras, Studying user income through language, behavior, and affect in social media, *PLOS ONE* (2015), <http://dx.doi.org/10.1371/journal.pone.0138717>.
- [10] Y. Tausczik, J. Pennebaker, The psychological meaning of words: LIWC and computerized text analysis, *Journal of Language and Social Psychology* **29**(1): 24-54 (March 2010), doi:10.1177/0261927X09351676.
- [11] S. Volkova, B. Van Durme, Online Bayesian models for personal analytics in social media, *Twenty-Ninth AAAI Conference on Artificial Intelligence* (2015).
- [12] L. Goldberg, An alternative “description of personality:” The Big-Five factor structure, *Journal of Personality and Social Psychology* **59**(6): 1216-1229 (December 1990).
- [13] J. Digman, Personality structure: Emergence of the five-factor model, *Annu. Rev. Psychol.* **41**: 417-440 (1990).
- [14] L. Goldberg, The structure of phenotypic personality traits, *American Psychologist* **48**(1): 26-34 (January 1993).
- [15] W. Norman, Toward an adequate taxonomy of personality attributes: Replicated factor structure in peer nomination personality ratings, *The Journal of Abnormal and Social Psychology* **66**(6): 574-583 (June 1963).
- [16] A. Schwartz, J. Eichstaedt, M. Kern, L. Dziurzynski, S. Ramones, M. Agrawal, A. Shah, M. Kosinski, D. Stillwell, M.E.P. Seligman, L.H. Ungar, Personality, gender, and age in the language of social media: The open-vocabulary approach, *PLoS ONE* **8**(9): e73791 (2013).
- [17] A. Carey, M. Brucks, A. Kufner, N. Holtzman, F. große Deters, M. Back, B. Donnellan, J. Pennebaker, M. Mehl, Narcissism and the use of personal pronouns revisited, *Journal of Personality and Social Psychology* **109**(3): e1-e15 (September 2015).
- [18] G. Farnadi, G. Sitaraman, M. Rohani, M. Kosinski, D. Stillwell, M.-F. Moens, S. Davalos, M. De Cock, How are you doing? Emotions and personality in Facebook, *Proceedings of 22nd International Conference on User Modeling, Adaptation and Personalization (UMAP)* (2014).
- [19] G. Park, H.A. Schwartz, J. Eichstaedt, M. Kern, M. Kosinski, D. Stillwell, L. Ungar, M.

- Seligman. Automatic personality assessment through social media language, *Journal of Personality and Social Psychology* **108**(6): 934-952 (June 2015), <http://dx.doi.org/10.1037/pspp0000020>.
- [20] R. Lambiotte, M. Kosinski, Tracking the digital footprints of personality, *Proceedings of the IEEE* **102**(12) (December 2014).
- [21] G. Miller, Social scientists wade into the Tweet stream, *Science* **333**: 1814-1815 (September 2011), doi:10.1126/science.333.6051.1814.
- [22] Y. Bachrach, R. Graepel, P. Kohli, M. Kosinski, D. Stillwell, Your digital image: Factors behind demographic and psychometric predictions from social network profiles, *Proceedings of the 2014 International Conference on Autonomous Agents and Multi-Agent Systems*, pp. 1649-1650, ISBN: 978-1-4503-2738-1, <http://dl.acm.org/citation.cfm?id=2616107>.
- [23] D. Preotiuc-Pietro, J. Eichstaedt, G. Park, M. Sap, L. Smith, V. Tobolsky, H.A. Schwartz, L. Ungar, The role of personality, age and gender in tweeting about mental illnesses, *Proceedings of the 2nd Workshop on Computational Linguistics and Clinical Psychology: From Linguistic Signal to Clinical Reality*, pp. 21-30 (2015).
- [24] G. Coppersmith, M. Dredze, C. Harman, K. Hollingshead, From ADHD to SAD: Analyzing the language of mental health on Twitter through self-reported diagnoses, *Proceedings of the 2nd Workshop on Computational Linguistics and Clinical Psychology: From Linguistic Signal to Clinical Reality*, pp. 1-10 (2015).
- [25] J. Eichstaedt, H.A. Schwartz, M. Kern, G. Park, D. Labarthe, R. Merchant, S. Jha, M. Agrawal, L. Dziurzynski, M. Sap, C. Weeg, E. Larson, L. Ungar, M. Seligman, Psychological language on Twitter predicts county-level heart disease mortality, *Psychol. Sci.* **26**(2):159-169 (February 2015), doi: 10.1177/0956797614557867.
- [26] Stuck\_In\_the\_Matrix, I have every publicly available Reddit comment for research. 1.7 billion comments @ 250 GB compressed. Any interest in this?, *Reddit* (2015), [https://www.reddit.com/r/datasets/comments/3bxlg7/i\\_have\\_every\\_publicly\\_available\\_reddit\\_comment/](https://www.reddit.com/r/datasets/comments/3bxlg7/i_have_every_publicly_available_reddit_comment/)
- [27] R. McCrae, O. John, An introduction to the five-factor model and its applications, *Journal of Personality* **60**(2): 175-215 (1992), 10.1111/j.1467-6494.1992.tb00970.x.