Modeling and Visualizing Locus of Control with Facebook Language

Kokil Jaidka,[1] **Anneke Buffone,**[2] **Johannes Eichstaedt,**[2]
Masoud Rouhizadeh,[3] **Lyle H. Ungar**[1]

[1] Computer & Information Science, University of Pennsylvania
[2] Department of Psychology, University of Pennsylvania
[3] School of Medicine, Johns Hopkins University
jaidka@sas.upenn.edu

Abstract

A body of literature has demonstrated that users' psychological traits such as personality can be predicted from their posts on social media. However, there is still a gap between the computational and descriptive analyses of the language features associated with different psychological traits, and their use by social scientists and psychologists to make deeper behavioral inferences. In this study, we aim to bridge this gap with a visualization that situates the language associated with one psychological trait in the context of other psychological dimensions. We predict Locus of Control (LoC), an individual's perception of personal control over events in their lives, from their Facebook language (F1=0.82). We then look at how language explains the relationship of LoC with consciousness and emotional stability.

Introduction

The abundance of social media data presents researchers with a unique opportunity to profile users and communities from the language they write. Many researchers have explored users' social media language to infer user attributes including age and gender (Schwartz et al. 2013; Jaidka, Guntuku, and Ungar 2018), personality (Plonsky, Erev, and others 2017; Rieman et al. 2017), and mental and physical health (Jaidka, Guntuku, and Ungar 2018). However, it is not known whether the relationships between different user traits connect with each other, and whether these relationships can be inferred on the basis of language alone. Our study aims to fill this gap by visualizing language in terms of interdependent psychological traits, to identify new relationships and facilitate new inferences. We focus on **Locus of Control (LoC)**, a construct that reflects the extent to which people ascribe the cause or control of events in their lives to themselves or the external factors (Rotter 1966):

- **Externals** tend to view the control of events as beyond their grasp, or attribute control to outside forces. They feel controlled by others or their circumstances.
- **Internals** tend to ascribe the control of events to themselves. They feel in control of their decisions and their circumstances.

Locus of Control is closely linked to stress, a primary factor linked to poor job performance, which has allegedly cost

corporations millions of dollars. Its consequences such as absenteeism have been found to impair productivity, communication, and trust within the workplace (Birnbaum et al. 2010). Organizations worldwide are attempting to better understand issues of employee well-being, including their locus of control and the closely related trait of self-efficacy. The British government conducted the Whitehall study to measure self-efficacy and job satisfaction (Marmot et al. 1991), and Gallup Inc. has conducted thousands of studies of worker well-being and satisfaction (Harter, Schmidt, and Keyes 2003). We anticipate that an employee's language could provide insight into their LoC, and thus allow unobstrusive, cost-effective estimates of employee health and job performance[1].

Setup and Motivation: In order to understand the subtler facets of internals and externals, which may not be captured in survey-items, it is important to analyze LoC's relationship with other personality traits. Prior work has investigated its association with the Big Five Personality Taxonomy (John and Srivastava 1999), finding a strong positive correlation between LoC and emotional stability (Judge et al. 2002) and a strong positive correlation of LoC with conscientiousness (Zuckerman et al. 1993). In this paper we investigate how these relationships between LoC and personality traits manifest themselves in language use. We use an *open-vocabulary approach* that mines linguistic concepts from corpora by first, grouping words occurring in the same context as 'topics' using Latent Dirichlet Allocation, and then relating them to other well-known personality traits, as was demonstrated by Park et. al (2016).

This paper makes two main contributions: (1) We build a model that predicts locus of control from language [2] and (2) We explore the relationship of LoC with two key personality traits (consciousness and emotional stability), looking at how language explains these connections.

Data We recruited our subjects – adults in the United States – via Qualtrics, a platform (similar to Amazon's Mechanical Turk) for deploying surveys and recruiting partici-

---

[1] http://www.pewinternet.org/2016/06/22/social-media-and-the-workplace/

[2] Available for non-commercial use at http://www.wwbp.org/data.html.

pants. All procedures were approved by the Institutional Review Board of the University of Pennsylvania. Our survey comprised demographic questions (age, gender, race, education and income brackets as per the items in the National Census), standardized psychological scales to measure Locus of Control (Brim, Ryff, and Kessler 2004), stress (Cohen, Kamarck, and Mermelstein 1983) and items from the United States Centers for Disease Control about their general health: number of days of work missed in the past year (rating from 0 (0-2 days) to 4 (16+ days), general health (rating from 0 (very healthy) to 5 (frequently ill)). Question order was randomized and scores were reverse-coded so that a higher score reflected better health.

Our analysis is based on the 2348 individuals who consented to share their Facebook data and posted at least 500 words[3]. 705 participants self-identified as female. The mean age of the sample is 36 and the median age is 38.

Locus of Control: We estimate user-level LoC by using the Sense of Control facet items from the MIDUS survey (Brim, Ryff, and Kessler 2004) which is a national survey of continental U.S. residents on psychological and social factors. Participants indicated how accurately the short phrases described themselves on a scale of 0 = very inaccurate to 4 = very accurate. The final LoC measure was based on a factor analysis and comprised six survey items with an internal reliability of 0.74 and an R^2 of 0.48 (mean=14.9; median=16).

We consider that an individual can either be an "internal" or an "external", so we modeled LoC as a binary variable, categorizing 1658 users with an LoC > 12 as *internals*, and the 690 users with an LoC $<= 12$ as *externals*. 12 was chosen as the midpoint of the maximum obtainable score, i.e. 24; however; no significant differences in predictive performance were observed if a mean, median or tercile split were performed instead. LoC is weakly positively correlated with age, being male ($r = 0.05$, $p < 0.05$) and income ($r = 0.14$, $p < 0.001$), which corroborates previous findings about its relationship with socioeconomic status. LoC is not significantly correlated with race or the total number of words posted on social media.

Table 1 shows why LoC is important for studies in health and management: it is a better predictor of absenteeism, health and especially stress than user demographics. Internals are likely to enjoy better health and suffer less stress at work, while externals are more likely to fall sick, miss work and experience high levels of stress.

Table 1: Pearson correlation of LoC with absenteeism, general health and stress ** $= p < 0.05$, two-tailed.

Traits	No. Workdays Missed	General Health	Self-Reported Stress
Locus of Control	**-.28****	**.28****	**-.53****
Age	.003	-.006	-.10**
Gender	.02	.02	.002
Education	-.01	.14**	-.15**
Income	-.06**	.20**	-.17**

---

[3]Studies have recommended analyzing at least 500 words per individual for stable results (Sap et al. 2014)

Predictive models of LoC

We collected 1.2 million posts from 2348 users, averaging 839 posts and 2845 unique words per individual. Tokenizing these posts using the HappierFunTokenizer[4], which is customized for use on social media posts, produced a total of 6.6 million tokens.

Lexica (64 features): We represent each user as a frequency distribution of categories from the Linguistic Inquiry and Word Count (LIWC) (Pennebaker, Booth, and Francis 2007), which comprises lists of words denoting psychological concepts, parts-of-speech and emotions.

N-grams (3000 features): We use a bag-of-words representation to reduce each users posting history to a normalized frequency distribution over a vocabulary, retaining only the most frequent 1000 1-, 2- and 3-grams each.

LDA Topics (2000 features): We use 2000 social-media specific topics as a data-driven lexicon. These topics were modeled using Latent Dirichlet Allocation from approximately 18 million Facebook updates [5]. A user is represented as $usage(topic|user)$, in terms of their probability of using each of the 2000 topics, which is further modeled in terms of the $p(topic|word)$, the probability of the topic given the word from LDA. We use the joint probability, $p(word, topic)$, to determine a words prevalence in a topic.

Results: In Table 2, predictive performance for LoC is reported as an F1-score on held out data in 10-fold cross validation. Classifiers were trained on different sets of language features, after feature selection based on a univariate regression and randomized principal component analysis (PCA). Language features out-predict user traits at predicting LoC. The best performing classifier is the gradient-boosted classifier trained on n-grams and LDA topics[6].

Table 2: Predictive performance (F-1 score) for Locus of Control, trained on different feature sets.

Feature Set	Gradient-Boosted Classification
Age, Gender and Income	.54
LIWC	.78
Parts of Speech	.79
N-grams	.81
LDA Topics	.81
N-grams + Topics	**.82**

Visualization

In this section, we describe the approach followed to relate LoC with other personality traits by using a larger corpus: the MyPersonality dataset (Kosinski, Stillwell, and Graepel 2013), which comprises approximately 15 million Facebook status updates shared after informed consent by over 70000

---

[4]https://github.com/dlatk/happierfuntokenizing

[5]available at http://www.wwbp.org/data.html

[6]The learning rate is set at 0.1 and the number of estimators is set to 500, with a maximum tree depth of 5.
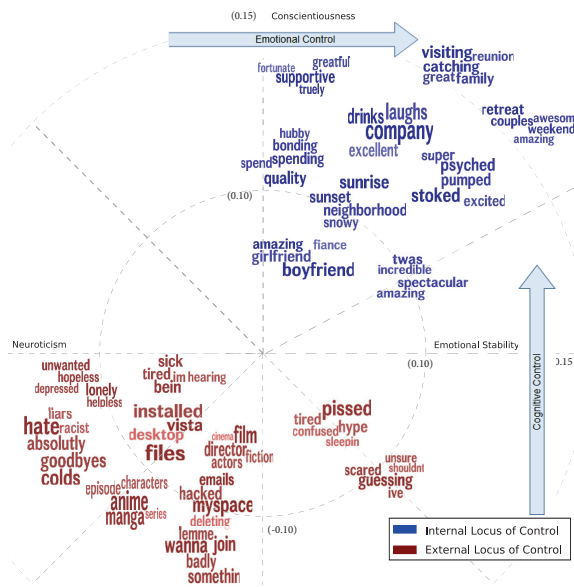
Figure 1: Locus of Control-linked language topics correlated with emotional stability and conscientiousness. Larger words are more frequent; darker words have a higher Pearson's correlation with internal (blue) or external (red) LoC.

participants who also took a variety of questionnaires, including the International Personality Item Pool proxy for the Big Five Personality taxonomy (John and Srivastava 1999).

Identifying language correlates of personality traits: We replicate our feature extraction approach to represent the language of MyPersonality users in terms of LDA topics, and then perform a correlation analysis between the 2000 topics and the five personality traits (i.e., openness, conscientiousness, extraversion, agreeableness and neuroticism or emotional stability).

Predicting users' personality traits: We use a pre-trained language-based predictive model built on the MyPersonality dataset to predict the personality scores for our set of users. This model was provided as a free resource by Schwartz et. al (Schwartz et al. 2013), who reported a predictive performance of $r > 0.3$ for all five traits, which is considered high accuracy in psychology for measuring internal states. LoC is moderate-to-weakly correlated with conscientiousness ($r = 0.15$) and emotional stability ($r = 0.10$) with $p < 0.01$ after controlling for age and gender, which corroborates previous results by Zuckerman et. al (1993). Conscientiousness and emotional stability are moderately correlated with each other ($r = 0.29, p < 0.01$). In the psychology literature, conscientiousness is considered an embodiment of 'cognitive control'. Emotional stability, the negative of neuroticism, is considered a proxy for 'emotional control'. Accordingly, we chose to study the language of LoC in the Emotional Stability - Conscientiousness space.

Language analysis of LoC and Personality

We follow the approach described by Park et. al (2016) to place LoC-linked LDA topics within the axes of the Big Five personality traits. We retain those topics which are significantly correlated with LoC after Bonferroni correction ($p < 0.05$). We identify the 55 topics which are also significantly correlated with emotional stability (after controlling for age, gender and conscientiousness) and conscientiousness (after controlling for age, gender and emotional stability) in the myPersonality dataset. We represent each topic as a pair of coordinates, reflecting its partial correlation with emotional stability and conscientiousness.

Results: The emotional stability-conscientiousness plot in Figure illustrates some of the topics as word clouds with the prominent words which had the highest effect sizes. The sizes and shades of words reflect the frequency and correlation of words (against LoC) for the entire dataset. The **size** of a word represents its prevalence within the topic (larger is more frequent) according to the original LDA modeling (Schwartz et al. 2013), and the shade reflects its Pearson's correlation with LoC (darker is stronger). The partial correlations of topics linked to externals (red topics) ranged from -0.12 to 0.06, while the effect sizes for the topics linked to internals (blue topics) ranged from 0.14 to 0.06 (all significant at $p < 0.05$), which is a range typical of the study of the language correlates of psychological traits (Schwartz et al. 2013). In the plot, we have adjusted the word clouds in order to avoid overlap.

We explore this spatial distribution in two complementary ways. First, we visualize the pattern of differences by coloring topics according to whether they significantly correlate with internals (blue) or externals (red). Secondly, we compare the distributions of internal- and external-LoC topics within interpersonally distinct areas of the plot. Internals-linked topics dominated the more conscientious half and the more emotionally stable quadrant, which is consistent with previous findings that individuals with an internal sense of control are usually conscientious and emotionally stable (Zuckerman et al. 1993).

Internals reiterate their emotional connections with topics that mention the 'support' of family and relationships, with words like 'boyfriend', 'girlfriend' 'fiance' and 'supportive' and topics about 'celebration'. We posit that this sense of relatedness, belonging and well-being (Baumeister and Leary 1995) affects how internals perceive the world, leading them to behave more positively towards their environment, projecting positive emotions and excitement. Social buffering may reduce the extent to which internals experience psychological distress, which may be why internals are associated with better health in Table 1. Internals appear to be more self-aware and thankful ('supportive','grateful','blessed') and more plan-oriented ('saturday','plan','cookout') as they actively mention their plans for a day in the near future.

Externals-linked topics express greater self-focus and negative affect words compared to internals, including mentions of emotional grievances and wounds ('scars', 'wounds','heal') and helplessness ('confused'). Previous literature supports the association isolation and a lack of

a sense of connectedness with mental and physical illnesses and suicidal ideation (Pennebaker, Booth, and Francis 2007). Externals are more likely to discuss media, such as TV shows and movies, and the Internet, such as emails, 'logging in' and 'notifications'. Mood management theory (Zillmann 1988) suggests that these reflect an attempt to alter negative moods and meet emotional needs through media exposure.

Inferences: Two arrows at the top and the right of the plot indicate the direction of increasing cognitive control (with increasing correlation with conscientiousness) and increasing emotional control (with increasing correlation with Emotional Stability). Now, among the topics associated with internals, certain topics are associated with more and some with lesser cognitive control. Among the topics associated with externals, psychological distress (evidenced by topics mentioning 'lonely, 'depressed, 'helplessness) are possibly more strongly associated with their lack of emotional control, rather than denoting the lack of cognitive control.

Conclusion

Locus of control is better at predicting the number of work days missed, general health status and self-reported stress than models that combined demographic and socioeconomic variables. In recent years, governments and corporations have attempted to better understand issues of employee well-being including their locus of control and self-efficacy. In this paper, we demonstrate that a less expensive alternative to estimate employees' locus of control could be from their language on social media, which would offer high levels of spatial and temporal resolution with reasonable accuracy (F1 = 0.82). Our results also corroborate previous studies' findings of locus of control as a core self-evaluation concept (Judge et al. 2002) and an important determinant of work and professional outcomes.

Our findings highlight the value of data-driven, open-vocabulary-based language analysis in the social sciences to help identify novel behavior patterns associated with known psychological constructs. We are also able to differentiate aspects of emotional and cognitive control between internals and externals in terms of their relationships with other personality dimensions. Our approach borrows equally from computational linguistics and psychology and suggests that exploratory, data-driven language analyses may extend our understanding of emotions, thoughts, and behaviors associated with known constructs in psychological theory.

References

Baumeister, R. F., and Leary, M. R. 1995. The need to belong: desire for interpersonal attachments as a fundamental human motivation. *Psychological bulletin* 117(3):497.

Birnbaum, H. G.; Kessler, R. C.; Kelley, D.; Ben-Hamadi, R.; Joish, V. N.; and Greenberg, P. E. 2010. Employer burden of mild, moderate, and severe major depressive disorder: mental health services utilization and costs, and work performance. *Depression and anxiety* 27(1):78–89.

Brim, O. G.; Ryff, C. D.; and Kessler, R. C. 2004. *The MIDUS National Survey: An Overview.* University of Chicago Press.

Cohen, S.; Kamarck, T.; and Mermelstein, R. 1983. A global measure of perceived stress. *Journal of health and social behavior* 385–396.

Harter, J. K.; Schmidt, F. L.; and Keyes, C. L. 2003. Well-being in the workplace and its relationship to business outcomes: A review of the gallup studies. *Flourishing: Positive psychology and the life well-lived* 2:205–224.

Jaidka, K.; Guntuku, S. C.; and Ungar, L. 2018. Facebook vs. twitter: Differences in self-disclosure and trait prediction. In *Proceedings of the International AAAI Conference on Web and Social Media*.

John, O. P., and Srivastava, S. 1999. The big five trait taxonomy: History, measurement, and theoretical perspectives. *Handbook of personality: Theory and research* 2(1999):102–138.

Judge, T. A.; Erez, A.; Bono, J. E.; and Thoresen, C. J. 2002. Are measures of self-esteem, neuroticism, locus of control, and generalized self-efficacy indicators of a common core construct? *Journal of personality and social psychology* 83(3):693.

Kosinski, M.; Stillwell, D.; and Graepel, T. 2013. Private traits and attributes are predictable from digital records of human behavior. *Proceedings of the National Academy of Sciences* 110(15).

Marmot, M. G.; Stansfeld, S.; Patel, C.; North, F.; Head, J.; White, I.; Brunner, E.; Feeney, A.; and Smith, G. D. 1991. Health inequalities among british civil servants: the whitehall ii study. *The Lancet* 337(8754):1387–1393.

Park, G.; Yaden, D. B.; Schwartz, H. A.; et al. 2016. Women are warmer but no less assertive than men: Gender and language on facebook. *PloS one* 11(5):e0155885.

Pennebaker, J. W.; Booth, R. J.; and Francis, M. E. 2007. Linguistic inquiry and word count: Liwc [computer software]. *Austin, TX: liwc. net*.

Plonsky, O.; Erev, I.; et al. 2017. Psychological forest: Predicting human behavior. In *Proceedings of the 31st AAAI Conference on Artificial Intelligence*, AAAI, 656–662. AAAI Press.

Rieman, D.; Jaidka, K.; Schwartz, H. A.; and Ungar, L. 2017. Domain adaptation from user-level facebook models to county-level twitter predictions. In *Proceedings of the Eighth International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, volume 1, 764–773.

Rotter, J. B. 1966. Generalized expectancies for internal versus external control of reinforcement. *Psychological monographs: General and applied* 80(1):1.

Sap, M.; Park, G.; Eichstaedt, J.; Kern, M.; Stillwell, D.; Kosinski, M.; Ungar, L.; and Schwartz, H. A. 2014. Developing age and gender predictive lexica over social media. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, 1146–1151.

Schwartz, H. A.; Eichstaedt, J. C.; Kern, M. L.; et al. 2013. Personality, gender, and age in the language of social media: The open-vocabulary approach. *PloS one* 8(9):e73791.

Zillmann, D. 1988. Mood management: Using entertainment to full advantage. *Communication, social cognition, and affect* 31:147–171.

Zuckerman, M.; Kuhlman, D. M.; Joireman, J.; et al. 1993. A comparison of three structural models for personality: The big three, the big five, and the alternative five. *Journal of personality and social psychology* 65(4):757.