

Predicting Human Trustfulness from Facebook Language

Mohammadzaman Zamani
Computer Science Department
Stony Brook University
mzamani@cs.stonybrook.edu

Anneke Buffone
Department of Psychology
University of Pennsylvania
buffonea@sas.upenn.edu

H. Andrew Schwartz
Computer Science Department
Stony Brook University
has@cs.stonybrook.edu

Abstract

Trustfulness — one’s general tendency to have confidence in unknown people or situations — predicts many important real-world outcomes such as mental health and likelihood to cooperate with others such as clinicians. While data-driven measures of *interpersonal trust* have previously been introduced, here, we develop the first language-based assessment of the personality trait of *trustfulness* by fitting one’s language to an accepted questionnaire-based trust score. Further, using trustfulness as a type of case study, we explore the role of questionnaire size as well as word count in developing language-based predictive models of users’ psychological traits. We find that leveraging a longer questionnaire can yield greater test set accuracy, while, for training, we find it beneficial to include users who took smaller questionnaires which offers more observations for training. Similarly, after noting a decrease in individual prediction error as word count increased, we found a word count-weighted training scheme was helpful when there were very few users in the first place.

1 Introduction

Trust, in general, indicates confidence that an entity or entities will behave in an expected manner (Singh and Bawa, 2007). While trust has been computationally explored as a property of relationships between people, i.e. *interpersonal trust* (Golbeck et al., 2003; Colquitt et al., 2007; Murray et al., 2012), few have considered *trustfulness* – a personality trait of an individual indicating their tendency, outside of any other context, to trust in people, institutions, and situations (Nannestad, 2008).

Trustfulness is tied to many real world and social outcomes. For example, it predicts individual health (Helliwell and Wang, 2010), and how

likely one is to join or to cooperate in diverse social groups (Uslaner, 2002; Stolle, 2002), and individual mental health and well-being (Helliwell and Wang, 2010). The importance of trustfulness is thought to be increasing as modern societies are increasingly interacting online with unknown people (Dinesen and Bekkers, 2016). This suggests it could be increasingly important in a clinical domain where has been shown to be essential in securing a strong and effective patient-client bond (Brennan et al., 2013; Lambert and Barley, 2001). Trait trustfulness also relates to self-disclosure which in turn greatly aids the clinician in her provision of care (Steel, 1991). Provider trust also likely is important to effectively treat a patient, especially in online therapeutic sessions, as it signals trustworthiness and care, but research on this topic remains sparse.

Unfortunately, traditional trustfulness measurement options (e.g. surveys) are expensive to scale to large populations and repeated assessment (i.e. in clinical practice) and they carry biases (Baumeister et al., 2007; Youyou et al., 2017). Researchers are actively searching for alternative behavior-based methods of measurement (Nannestad, 2008).

Language use in social media offers a behavior from which one can measure psychological traits like trust. Over the last five years, more and more researchers are turning to Facebook or Twitter language to develop psychological trait predictors, fitting user language to psychological scores from questionnaires (Schwartz and Ungar, 2015). According to standard psychometric validity tests, such language-based approaches have been found to rival other accepted measures, such as questionnaires and assessments from friends (Park et al., 2015). However, while language-based predictive models for many traits now exist, none have considered a model for trustfulness— a trait which

some have argued is now of marked importance as modern societies are increasingly interacting on-line with unknown people (Dinesen and Bekkers, 2016). Further, across such trait prediction work, little attention has been paid to the role of (1) *questionnaire-size* – how many questions are used to assess an individual’s trait, and (2) *word count* – how many words the user has written from which the language-based predictions are made.¹

Here, we answer the call for more behavior-based trait measurement (Baumeister et al., 2007; Youyou et al., 2017), by developing language-based (a behavior) predictive model of trustfulness fit to questionnaire scores, and we seek to draw insights into the role of word count and questionnaire size in predictive modeling.

Contributions. This work makes several key contributions. First, we introduce the first language-based assessment of *trustfulness* (henceforth “trust”), evaluated over out-of-sample trust questionnaires, enabling large-scale or frequently repeated trust measurement. We also (2) study the number of questions in the psychological survey to which one fits our model (in other words, finding which one matters more: number of questions in questionnaires or number of users who took it?), (3) explore the relationship between users’ word count and model error, and (4) introduce a weighting scheme to train on low word count users. All together, we add trustfulness, an important trait for clinical care, to an increasing battery of language-based assessments.

2 Background

Previous computational work on trust has focused on *interpersonal trust* – an expectation of trust concerning future behaviour of a specific person toward another known person. (Bamberger, 2010). Interpersonal trust is primarily focused on situations in which there are two known individuals (the truster and trustee) who share a history of previous interactions. Such trust, requires study of a history of interactions indicating how well each member participant might understand the others’ personalities (Kelton et al., 2008; Golbeck et al., 2003). Interpersonal trust has been studied especially in the context of online social networks where it is sometimes possible to track users

¹One often imposes a word count limitation — e.g. users must write at least 1,000 words (Schwartz et al., 2013) — but few have studied the relationship between word count and accuracy as we do here.

from first interactions (Kuter and Golbeck, 2007; DuBois et al., 2011; Liu et al., 2014, 2008). While some of these works have considered the amount of communication (Adali et al., 2010), content is rarely considered and none of these past works have attempted to measure the trait, *trustfulness*, as we do here.

Trustfulness (also referred to as “generalized trust”), in contrast with interpersonal trust, measures trust between strangers. As Stolle (2002) put it:

[Trustfulness] indicates the potential readiness of citizens to cooperate with each other and to abstract preparedness to engage in civic endeavors with each other. Attitudes of trustfulness extend beyond the boundaries of face-to-face interactions and incorporate people who are not personally known.

This version of trust has been tied to the belief in the average goodness of human nature (Yamagishi and Yamagishi, 1994), and it involves a willingness to be vulnerable and engage with random others despite interpersonal risks (Mayer and Davis, 1999; Rousseau et al., 1998). It has been shown predictive of individual mental health and physical well-being (Abbott and Freeth, 2008; Helliwell and Wang, 2010). For communities, trust is a key indicator of social capital (Coleman, 1988; Putnam, 1993), and it is highly predictive of economic growth (Delhey and Newton, 2005; Knack and Zak, 2003)

2.1 Trustfulness from Questionnaires

Trustfulness, just like other personality traits is typically measured with either questionnaires or behavioral observations during experiments (Er-misch et al., 2009). Data linking experiments on trust with individual linguistic data is not available or easily acquired, so we fit our language-based model of trust to a gold-standard of questionnaire-based trust. A variety of such questionnaires exist with high inter-correlation, including the Faith in People scale (Rosenberg, 1957), Yamagishi & Yamagishis (1999) Trust Scale, and the Trust Facet of the Agreeableness trait in the Big Five personality questionnaire (Goldberg et al., 2006). Here, due to its availability, we chose to fit our language-based trust predictor to the later of these questionnaires – the trust personality facet.

3 Data Set.

We use trust facet scores from the trait questionnaire of consenting participants of the MyPersonality study (Kosinski et al., 2015).² From this dataset we derive two versions of trust measurement scores: (1) using 10 questions of trustfulness (referred to as *10-question trust*), or (2) using a subset of 3 questions (referred to as *3-question trust*). Participants can either answer all 10 questions (as part of larger set of over 300 questions) or just answer the 3-question version (as part of a 100 questions). Each question is on a scale of 1 to 5, from totally disagree to completely agree. For example, the following are the questions for the 3-item version:

- *I believe that others have good intentions.*
- *I suspect hidden motives in others.**
- *I trust what people say.*

Some questions (e.g. * above) are “reverse scored” so a 1 becomes a 5 and vice-versa. One’s final trust score is based on taking the mean of the responses to the individual trust questions. Although 3-question trust is less accurate,³ it may be useful to enable training data from more users.

From MyPersonality, we used a dataset containing 19,455 Facebook users who wrote at least 1,000 words across all of their status updates. We additionally included 6,590 users who had less than 1,000 words in some experiments. Totally 26,045 users took the Big Five questionnaire, answering at least the 3 trust-focused questions in it (short version). Among all the users, only 621 had completely answered all of the 10 trust related question (long version). Table 3 represents number of users in detail. It is worth mentioning that not only the participants consent for their Facebook and questionnaire data to be used in research, but also the data has been anonymized.

4 Method

We build a language-based model for the trait of trustfulness. From Facebook status updates,

²Procedures were approved by an academic institutional review board, and volunteers agreed to share their data for research purposes via informed consent.

³Based on an experiment across 1,041 participants on Amazon Mechanical Turk, 3-question trust had a Pearson correlation of 0.916 with 10-question trust, indicating it is a reasonable approximation (Buffone et al., 2017)

	Long Version	Short Version
Threshold-1000	438	19445
Threshold-10	621	26045

Table 1: Number of users who filled the long or short version questionnaire based on their word counts. Threshold-X means setting word count threshold to X. Long version represents users who had 10-question trust score, and short version includes users who had 3-question trust score.

we extracted two types of user-level lexical features, which have previously been shown to be effective for trait prediction (Park et al., 2015): (a) ngrams of length 1 to 3 and (b) LDA topics. To extract the ngrams from the text we used the *HappierFunTokenizer*. We did not apply any text normalization, as past work has found that often the forms in which people choose to write a word ends up being predictive about their personality (Schwartz et al., 2013). Two types of ngrams were extracted: one containing relative frequencies of each ngram ($\frac{freq(ngram,user)}{freq(*,user)}$) and the other simply a binary indicator of whether the user mentioned each ngram at all. Considering ngrams mentioned by at least 1% of the users, we obtained 50,166 ngrams features for each of the two types of ngrams. Topic features were derived from posteriors of Latent Dirichlet Allocation. We use the 2,000 LDA topic posteriors publicly available from Schwartz et al. (2013).⁴

We use a series of steps to avoid high dimensional issues and prevent overfitting. First, an occurrence threshold is applied to remove words that were used by less than 1% of people. Second, we select features with at least a small relationship with our trust labels according to having a univariate family-wise error rate < 60 . Third, we ran a singular value decomposition (in randomized batches) to effectively decrease the size of feature space and reduce colinearity across dimensions (Boutsidis et al., 2015). We performed this process based on the training data, and then applied the resulting feature reduction on the test data.

Each type of feature (i.e. ngram relative frequencies, booleans, and topics) is qualitatively and distributionally different from each other (Almodaresi et al., 2017). Thus, we perform

⁴Topic posteriors and the tokenizer, Happier Fun Tokenizing, are available from <http://wwbp.org/data.html>

reduction technique on ngrams, boolean ngrams and topics separately. This is so the comparatively few topic features are not likely to get lost among the relatively plentiful ngrams. At the end, we merge both types of features to build a single feature matrix (or an embedding with approximately 5% of the number of training observations). Similar feature reduction pipelines have been shown to perform well in language based predictive analytics (Zamani and Schwartz, 2017). We then use ridge regression to fit our dimensionally reduced feature set to the trust labels from the Big Five questionnaires.

Questionnaire Size and Word Count. While the 10-question trust score is more accurate, we have less than 1,000 users with this label. Our default setup has the users with 10-question trust as the test set while we train over the much larger set of users with only 3-question trust. We then experiment to determine if this setup is ideal.

Previous work has suggested user attribute prediction benefits from an approximate minimum threshold of 1,000 words per user in order to get accurate estimates of one’s personality (Schwartz et al., 2013). Since our dataset contains 6,590 users with less than 1,000 words, we explore if we can include these users in an effective way to improve the model. To this end, we weight each users’ contribution to the loss function proportionate to the number of words she or he has written. We used two different weighting schemes, linear and logistic, as shown below, where wc is the word count, and T_{max} and T_{min} are 1,000 and 200 respectively.

$$W_{linear} = \frac{\min(T_{max}, \max(0, wc - T_{min}))}{T_{max} - T_{min}}$$

$$W_{logistic} = \frac{1}{1 + \exp(-W_{linear})}$$

Thus, users with more than 1,000 words are weighted 1 while those with less than 200 words are weighted 0 (we settled on these min and max values based on our study of the mean error per word count – Figure 1).

5 Evaluation

We focus on evaluating our language model by comparing the performance of our model on prediction of 10-question trust vs. 3-question trust

labels. We did this comparison in 3 settings: (1) train and test on 10-question trust score, (2) train and test on 3-question trust score, and (3) train on 3-question and test on 10-question trust score.

For the first setting, where all users answered the same number of questions, we performed a 10-fold cross-validation. For the second and third settings, we consider all users with 10-question trust score as our test group and the remaining users which only had 3-question trust score but not the 10-question trust as the train group. This enables us to first determine how well a model trained on 3-question trust performs in not only predicting 3-question trust itself, but also the 10-question trust, and compare the later with the model which is trained on small group of users with 10-question trust. In all these three experiments, we considered 1,000 as the threshold for word count, and used the same group of users as the test group. We present result as both *mean squared error* and *dissattenuated correlation* which accounts for measurement error: $r_{dis(a,b)} = \frac{r_{a,b}}{\sqrt{r_{a,a}r_{b,b}}}$ where $r_{a,a} = .70$ the reliability of the trust questionnaire (Kosinski et al., 2015) and $r_{b,b} = .70$ the expected reliability of the trust language-based measurement based on evaluations of language-based personality assessment reliability (Park et al., 2015) (every r on the right-hand side of the equation is a Pearson product-moment correlation coefficient).

train label	test label	Pearson r_{dis}	MSE
10-question	10-question	0.259	0.719
3-question	3-question	0.426	0.776
3-question	10-question	0.494	0.662

Table 2: Comparing the language model performance on 3-question trust score vs. 10-question trust score. Pearson r_{dis} is dissattenuated Pearson r and MSE is the mean squared error.

As shown in table 2, our model’s r_{dis} with only limited 10-item data is 0.259, suggesting we cannot learn a very accurate model by training on such a small number of users. Comparing the second and third settings, we see the result of testing on 10-question trust score outperforms the 3-question trust score by 0.07 margin in dissattenuated Pearson r and MSE by a margin of 0.11. To further understand why 10-question trust seems to be easier to predict, we calculate the variance for both 3-question and 10-question trust, yielding $\sigma^2 = 0.85$ and $\sigma^2 = 0.72$ respectively. This suggests that 10-question trust has less noise than 3-question trust. Due to these results, in all of the following experi-

Features	Pearson r_{dis}	MSE
sentiment (baseline)	0.279	0.717
ngr_r	0.453	0.681
ngr_b	0.411	0.688
topics	0.458	0.677
word2vec	0.449	0.678
ngr_r + ngr_b + topics	0.494	0.662
ngr_r + ngr_b + topics + sent	0.483	0.666

Table 3: Comparing the performance of our language model with sentiment as baseline, using different feature sets: ngr_r: ngrams as relative frequency, ngr_b: ngrams as boolean variables. Bold indicates the best performance. Pearson r_{dis} is dissattenuated Pearson r and MSE is the mean squared error.

ments we only train on 3-question trust labels and test on 10-question trust labels.

We next evaluate the performance of our trust model by comparing to two baseline models. Because positiveness is associated with trust (Heliwell and Wang, 2010), we consider a baseline of sentiment scores using the NRC hashtag sentiment lexicon, an integral part of the best system participating in SemEval-2013 (Mohammad et al., 2013). We also compare it to clusters of words derived from word2vec embeddings (Mikolov et al., 2013) using spectral clustering (Preoțiu-Pietro et al., 2015).

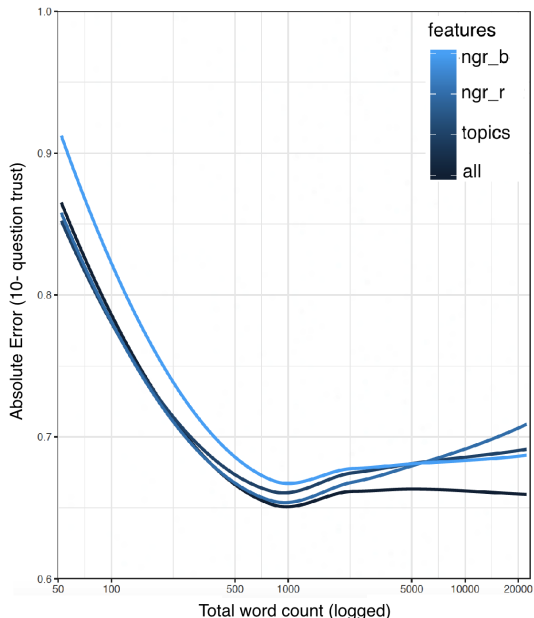


Figure 1: Effect of word count on error rate of the language model: ngr_b stands for binary-ngrams and ngr_r stands for relative-ngrams.

Table 3 demonstrates the predictive performance of our model in comparison to the sentiment and word2vec baselines. Our best model

($ngr_r + ngr_b + topics$) had an 8% reduction in mean squared error over sentiment, and achieved a Pearson correlation coefficient of $r_{dis} = .494$ which is considered a large relationship between a behavior (language use) and a psychological trait (Meyer et al., 2001) and just below state-of-the-art language-based assessments of other personality traits (Park et al., 2015).

In the next experiment we present how the error rate changes as a function of word count per user using various combinations of features. We trained 4 models using (1) relative-ngrams, (2) binary-ngrams, (3) topics, and (4) all features together. We predict the 10-question trust score of our test users and plot the test users error rate with respect to their word count, which is shown in figure 1. Overall, users’ trust score is more predictable as they use more words flattening out after 1000 words. Additionally, for users with few words, relative-ngrams and binary-ngrams are equally predictive and better than topics. For users with many words, the prediction power of binary-ngrams fades out, likely reflecting features being primarily ones. Similarly, topic-based models perform better for talkative users, likely because more words means better topic estimation.




Figure 2: Effect of increasing the number of training users, who have more than 1,000 word count, while there are 6,590 users with less than 1,000 word count in train set: “Threshold-1000” is training ridge-regression on users with at least 1,000 words, “threshold-200” is training ridge-regression on users with at least 200 words, “linear” is training weighted ridge-regression on users with at least 200 words, and finally “logistic” is training weighted ridge-regression on users with at least 200 words.




Figure 3: Unigrams most distinguish trust according to absolute value of (a) positive correlation and (b) negative correlated with 3-question trust score. Size of word indicates correlation strength, while color indicates frequency. All unigrams listed are significantly correlated at Benjamini-Hochberg corrected $p < .05$.

Now that we know word count is correlated with prediction error, we explore a word count weighting scheme that enables us to include 6,590 users with fewer than 1,000 words in training. Such users are included in three different ways, (1) without using any weight, (2) using linear weighting, and (3) using logistic weighting.

In figure 2 we compare the various model training setups at different training sizes. As shown, when we have just a few users with more than 1,000 words, including more users, but with low word count, improves the performance, no matter which models we exploit. However, as the number of users with more than 1,000 word count increases, injecting low word count users hurts the performance. In addition, the weighting scheme does not seem to help at all in this situation.

To get an idea of the type of features signalling high and low trust predictions, we ran a differential language analysis (Schwartz et al., 2013) to identify the top 50, independently, most predictive features. Figure 3 show the word-clouds of both positively correlated and negatively correlated with 3-question trust score, limited to those passing a Benjamini-Hochberg False Discovery rate α of 0.01 (Benjamini and Hochberg, 1995). Many of the ngrams correspond with the definition of trustfulness, such as the pro-social words in the positive predictors (e.g. ‘friends’ ‘family’, ‘thanks’). On the other hand, many curse words can be seen among negative predictors.

6 Conclusion

We introduced the first language-based model for measuring *trustfulness* from language, and used it to study novel and useful aspects of the predictive modeling of user traits. First, we found that language use in social media can be used to predict trustfulness about as accurate as other personality traits. Then, we found that, in order to build a language model over questionnaires, including *more* users who took a *shorter* questionnaire can lead to improvement, in comparison to using *less* users who took a *longer* questionnaire. We also showed that the language model usually performs better in predicting users with more total word count, with error flattening out around 1,000 words, and that when there are few users (i.e. $< 1,000$) it is worth lowering the minimum word count threshold to include more users for training purpose. However, using a weighting scheme was not helpful.

Our scaleable measure of trust enables future work to investigate some interesting questions about trust, such as those involved in large-scale or frequent assessments. For example, this may allow for large-scale assessments of trait trustfulness of different patient populations or of samples of clinicians. Also, if clients were to opt into sharing of social media, therapists may be able to use this model to detect drops in patient trust which may help to understand when one is more receptive or not. Trends over time may help to signal interpersonal improvements or regressions, as well as negative interactions with others. It should be noted

that while trust is thought of as a relatively stable personality aspect or trait, some research suggests that it is malleable over time (Jones and George, 1998), so changes in trust over time could be another meaningful exploration for future study. Thus, the present model may be helpful for the generation of trustful chat bots, such as virtual assistants or therapeutic aids.

References

- Stephen Abbott and Della Freeth. 2008. Social capital and health: starting to make sense of the role of generalized trust and reciprocity. *Journal of Health Psychology*, 13(7):874–883.
- Sibel Adali, Robert Escriva, Mark K Goldberg, Mykola Hayvanovych, Malik Magdon-Ismael, Boleslaw K Szymanski, William Wallace, Gregory Williams, et al. 2010. Measuring behavioral trust in social networks. In *Intelligence and Security Informatics (ISI), 2010 IEEE International Conference on*, pages 150–152. IEEE.
- Fatemeh Almodaresi, Lyle Ungar, Vivek Kulkarni, Mohsen Zakeri, Salvatore Giorgi, and H Andrew Schwartz. 2017. On the distribution of lexical features at multiple levels of analysis. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics*, volume 2, pages 79–84.
- Walter Bamberger. 2010. Interpersonal trust—attempt of a definition. *Scientific Report, Technical University Munich*.
- Roy F Baumeister, Kathleen D Vohs, and David C Funder. 2007. Psychology as the science of self-reports and finger movements: Whatever happened to actual behavior? *Perspectives on Psychological Science*, 2(4):396–403.
- Yoav Benjamini and Yoel Hochberg. 1995. Controlling the false discovery rate: a practical and powerful approach to multiple testing. *Journal of the royal statistical society. Series B (Methodological)*, pages 289–300.
- Christos Boutsidis, Anastasios Zouzias, Michael W Mahoney, and Petros Drineas. 2015. Randomized dimensionality reduction for k -means clustering. *IEEE Transactions on Information Theory*, 61(2):1045–1062.
- Nicola Brennan, Rebecca Barnes, Mike Calnan, Oonagh Corrigan, Paul Dieppe, and Vikki Entwistle. 2013. Trust in the health-care provider–patient relationship: a systematic mapping review of the evidence base. *International Journal for Quality in Health Care*, 25(6):682–688.
- Anneke Buffone, H Andrew Schwartz, Patrick Crutchley, Margaret L. Kern, Mohammadzaman Zamani, L. K. Smith, Johannes C. Eichstaedt, Lyle Ungar, and Martin E. P. M. Seligman. 2017. Measuring trust through large scale language analysis: Trust as an aspect of individuals and communities. *In Press*.
- James S Coleman. 1988. Social capital in the creation of human capital. *American journal of sociology*, 94:S95–S120.
- Jason A Colquitt, Brent A Scott, and Jeffery A LePine. 2007. Trust, trustworthiness, and trust propensity: a meta-analytic test of their unique relationships with risk taking and job performance. *Journal of applied psychology*, 92(4):909.
- Jan Delhey and Kenneth Newton. 2005. Predicting cross-national levels of social trust: global pattern or nordic exceptionalism? *European Sociological Review*, 21(4):311–327.
- Peter Thisted Dinesen and Rene Bekkers. 2016. The foundations of individuals generalized social trust: A review. In *Trust in Social Dilemmas*. Oxford University Press.
- Thomas DuBois, Jennifer Golbeck, and Aravind Srivasan. 2011. Predicting trust and distrust in social networks. In *Privacy, Security, Risk and Trust (PASSAT) and 2011 IEEE Third International Conference on Social Computing (SocialCom), 2011 IEEE Third International Conference on*, pages 418–424. IEEE.
- John Ermisch, Diego Gambetta, Heather Laurie, Thomas Siedler, and SC Noah Uhrig. 2009. Measuring people’s trust. *Journal of the Royal Statistical Society: Series A (Statistics in Society)*, 172(4):749–769.
- Jennifer Golbeck, Bijan Parsia, and James Hendler. 2003. *Trust networks on the semantic web*. Springer.
- Lewis R Goldberg, John A Johnson, Herbert W Eber, Robert Hogan, Michael C Ashton, C Robert Cloninger, and Harrison G Gough. 2006. The international personality item pool and the future of public-domain personality measures. *Journal of Research in personality*, 40(1):84–96.
- John F Helliwell and Shun Wang. 2010. Trust and well-being. Technical report, National Bureau of Economic Research.
- Gareth R Jones and Jennifer M George. 1998. The experience and evolution of trust: Implications for cooperation and teamwork. *Academy of management review*, 23(3):531–546.
- Kari Kelton, Kenneth R Fleischmann, and William A Wallace. 2008. Trust in digital information. *Journal of the American Society for Information Science and Technology*, 59(3):363–374.
- Stephen Knack and Paul J Zak. 2003. Building trust: public policy, interpersonal trust, and economic development. *Supreme court economic review*, 10:91–107.

- Michal Kosinski, Sandra C Matz, Samuel D Gosling, Vesselin Popov, and David Stillwell. 2015. Facebook as a research tool for the social sciences: Opportunities, challenges, ethical considerations, and practical guidelines. *American Psychologist*, 70(6):543.
- Ugur Kuter and Jennifer Golbeck. 2007. Sunny: A new algorithm for trust inference in social networks using probabilistic confidence models. In *AAAI*, volume 7, pages 1377–1382.
- Michael J Lambert and Dean E Barley. 2001. Research summary on the therapeutic relationship and psychotherapy outcome. *Psychotherapy: Theory, research, practice, training*, 38(4):357.
- Guangchi Liu, Qing Yang, Honggang Wang, Xiaodong Lin, and Mike P Wittie. 2014. Assessment of multi-hop interpersonal trust in social networks by three-valued subjective logic. In *INFOCOM, 2014 Proceedings IEEE*, pages 1698–1706. IEEE.
- Haifeng Liu, Ee-Peng Lim, Hady W Lauw, Minh-Tam Le, Aixun Sun, Jaideep Srivastava, and Young Kim. 2008. Predicting trusts among users of online communities: an epinions case study. In *Proceedings of the 9th ACM Conference on Electronic Commerce*, pages 310–319. ACM.
- Roger C Mayer and James H Davis. 1999. The effect of the performance appraisal system on trust for management: A field quasi-experiment. *Journal of applied psychology*, 84(1):123.
- Gregory J Meyer, Stephen E Finn, Lorraine D Eyde, Gary G Kay, Kevin L Moreland, Robert R Dies, Elena J Eisman, Tom W Kubiszyn, and Geoffrey M Reed. 2001. Psychological testing and psychological assessment: A review of evidence and issues. *American psychologist*, 56(2):128.
- Tomas Mikolov, Kai Chen, Greg Corrado, and Jeffrey Dean. 2013. Efficient estimation of word representations in vector space. *arXiv preprint arXiv:1301.3781*.
- Saif M Mohammad, Svetlana Kiritchenko, and Xiaodan Zhu. 2013. Nrc-canada: Building the state-of-the-art in sentiment analysis of tweets. *arXiv preprint arXiv:1308.6242*.
- Sandra L Murray, Shannon P Lupien, and Mark D Seery. 2012. Resilience in the face of romantic rejection: The automatic impulse to trust. *Journal of Experimental Social Psychology*, 48(4):845–854.
- Peter Nannestad. 2008. What have we learned about generalized trust, if anything? *Annu. Rev. Polit. Sci.*, 11:413–436.
- Gregory Park, H Andrew Schwartz, Johannes C Eichstaedt, Margaret L Kern, Michal Kosinski, David J Stillwell, Lyle H Ungar, and Martin EP Seligman. 2015. Automatic personality assessment through social media language. *Journal of personality and social psychology*, 108(6):934.
- Daniel Preotiu-Pietro, Vasileios Lampos, and Nikolaos Aletras. 2015. An analysis of the user occupational class through twitter content.
- Robert D Putnam. 1993. The prosperous community. *The american prospect*, 4(13):35–42.
- Morris Rosenberg. 1957. Occupation and values: Glencoe.
- Denise M Rousseau, Sim B Sitkin, Ronald S Burt, and Colin Camerer. 1998. Not so different after all: A cross-discipline view of trust. *Academy of management review*, 23(3):393–404.
- H Andrew Schwartz, Johannes C Eichstaedt, Margaret L Kern, Lukasz Dziurzynski, Stephanie M Ramones, Megha Agrawal, Achal Shah, Michal Kosinski, David Stillwell, Martin EP Seligman, et al. 2013. Personality, gender, and age in the language of social media: The open-vocabulary approach. *PLoS one*, 8(9):e73791.
- H Andrew Schwartz and Lyle H Ungar. 2015. Data-driven content analysis of social media: a systematic overview of automated methods. *The ANNALS of the American Academy of Political and Social Science*, 659(1):78–94.
- Sarbjee Singh and Seema Bawa. 2007. A privacy, trust and policy based authorization framework for services in distributed environments. *International Journal of Computer Science*, 2(2):85–92.
- Jennifer L Steel. 1991. Interpersonal correlates of trust and self-disclosure. *Psychological Reports*, 68(3-suppl):1319–1320.
- Dietlind Stolle. 2002. Trusting strangers—the concept of generalized trust in perspective. *Austrian Journal of Political Science*, 31(4):397–412.
- Eric M Uslaner. 2002. *The moral foundations of trust*. Cambridge University Press.
- Toshio Yamagishi, Masako Kikuchi, and Motoko Koguchi. 1999. Trust, gullibility, and social intelligence. *Asian Journal of Social Psychology*, 2(1):145–161.
- Toshio Yamagishi and Midori Yamagishi. 1994. Trust and commitment in the united states and japan. *Motivation and emotion*, 18(2):129–166.
- Wu Youyou, H Andrew Schwartz, David Stillwell, and Michal Kosinski. 2017. Birds of a feather do flock together: Behavior-based personality-assessment method reveals personality similarity among couples and friends. *Psychological Science*, page 0956797616678187.
- Mohammadzaman Zamani and H Andrew Schwartz. 2017. Using twitter language to predict the real estate market. *EACL 2017*, page 28.