



Predicting Personality from Social Media Text

Jennifer Golbeck

Human–Computer Interaction Lab, University of Maryland, College Park, MD, USA
golbeck@cs.umd.edu

Abstract:

This paper replicates text-based Big Five personality score predictions generated by the Receptiviti API—a tool built on and tied to the popular psycholinguistic analysis tool *Linguistic Inquiry and Word Count* (LIWC). We use four social media datasets with posts and personality scores for nearly 9,000 users to determine the accuracy of the Receptiviti predictions. We found Mean Absolute Error rates in the 15–30% range, which is a higher error rate than other personality prediction algorithms in the literature. Preliminary analysis suggests relative scores between groups of subjects may be maintained, which may be sufficient for many applications.

Keywords: social media, personality, replication, psycholinguistics

The manuscript was received 03/10/2016 and was with the authors 2 months for 2 revisions.

Introduction

Psycholinguistic text analysis is the study of the relationship between people's language use and their cognitive, emotional, and psychological traits. The *Linguistic Inquiry and Word Count* tool (abbr. LIWC) (Tausczik & Pennebaker, 2010) is a widely tested, validated, and applied system for performing psycholinguistic text analysis. It performs a statistical analysis of a person's word use by categories, including: positive and negative emotions; cursing; function words (e.g., articles, adverbs, nouns, etc.); focus on past, present, and future; and personal concerns (e.g., work, money, and religion).

In 2015, a new version of LIWC was released along with *Receptiviti*, a commercial brand of the tool¹. In addition to the LIWC categories, Receptiviti includes insights about personality, thinking style, authenticity, and relationships.

In this paper, we focus on the personality attributes in Receptiviti. The Receptiviti API provides estimates of Big Five personality traits for a user based on a text sample they have authored. The Big Five personality profile comprises five attributes: *openness* (to new experiences), *conscientiousness*, *agreeableness*, *extroversion* (or *introversion*), and *neuroticism* (measuring emotional stability) (John, Donahue, & Kentle, 1991). Note that these predictions do not necessarily model causal relationships between linguistics and personality, but rather use linguistic traits to create statistical models that estimate scores that subjects would receive on the Big Five Personality Inventory.

This study is a conceptual replication of a number of studies that predict Big Five Personality scores from social media text, specifically Golbeck, Robles, Edmondson, & Turner (2011) and Quercia, Kosinski, Stillwell, & Crowcroft (2011). While there are no publications about how the Big Five estimates are computed in Receptiviti, we have access to the tool so we can see the results of their method applied to

¹ From the receptiviti.com website: "Receptiviti was launched in conjunction with LIWC2015 and established as its commercial side to make LIWC more accessible to the software development and data science communities, and the growing number of businesses that want to incorporate its capabilities into their technologies."

our data. In this replication, we use their tool to make predictions, but we use four different datasets to evaluate their accuracy.

1 Related Work

There is a large and growing body of work that examines how Big Five personality traits manifest on social media and how to predict user's Big Five test scores from their online profiles.

Ryan & Xenos (2011) studied over 1,300 Australian Facebook users and had them take the Big Five Personality Inventory (BFPI) along with other standard personality tests. They found that Facebook users in their sample tended to be more extroverted, but less conscientious, than non-users. In the converse of that analysis, researchers have studied how personality traits correlated with certain behaviors online. Gosling, Augustine, Vazire, Holtzman, & Gaddis (2011) found that extroverts tended to use and engage with Facebook more often, and that users extended their personalities into their online interactions. This was an extension and validation of their prior work (Vazire & Gosling, 2004), which found that websites reflected valid personality information about the author.

There are also many papers that describe techniques for building predictive models that can infer the scores users would receive on the BFPI based on their social media profiles. These techniques leverage many types of profile data, including social network structure (Staiano et al., 2011) as well as profile activity and interaction patterns (Adali & Golbeck, 2014; Ferwerda, Schedl, & Tkalcic, 2015; Golbeck et al., 2011; Markovikj, Gievska, Kosinski, & Stillwell, 2013). It is especially relevant to this replication that many studies build their models based on analysis of the text users post in their updates (Farnadi, Zoghbi, Moens, & De Cock, 2013; Park et al., 2015; Quercia et al., 2011; Kosinski, Stillwell, & Graepel, 2013).

This text-based analysis is at the foundation of the Receptiviti API's personality prediction. While we do not have data on the kind of model they use, their work fits neatly into the space of work on predicting BFPI scores from social media data. While the Big Five predictions have not been published or validated publicly, many other attributes within the LIWC framework have been tested and validated (Tausczik & Pennebaker, 2010).

2 Method

The Receptiviti API allows users to submit a text sample. That text is analyzed and a graphical representation of some traits and a JSON-based result list is returned. Figure 1 shows the web-based interface to the API. To test the accuracy of the Big Five personality estimates produced by the Receptiviti API, we submitted social media-derived text from people for whom we had known Big Five scores.

We used four social media data sets collected in previous personality research studies. Each had Big Five personality scores obtained by administering a standard Big Five Personality Inventory to each subject (John et al., 1991). Two data sets (one from Facebook and another from Twitter) come from our previous research on predicting Big Five personality traits (Golbeck et al., 2011; Golbeck, Robles, & Turner, 2011). The other Facebook data sets are from the *myPersonality* project (Celli, Pianesi, Stillwell, & Kosinski, 2013). They provide one set of 10,000 status updates for 250 users with personality scores (which we will call the *myPersonality Small* data set), as well as a large database of 22 million status updates for 154,000 users, and personality scores for each user. Their data has been used to successfully model Big Five personality traits on a large scale (Kosinski et al., 2013) as well as numerous other papers on personality and social media, including Alam, Stepanov, & Riccardi (2013), Bachrach, Graepel, Kohli, Kosinski, & Stillwell (2014), Farnadi et al. (2013), and Iacobelli & Culotta (2013), to list just a few.

For our analysis, we merged the status update and personality databases to create a *myPersonality Large* data set.

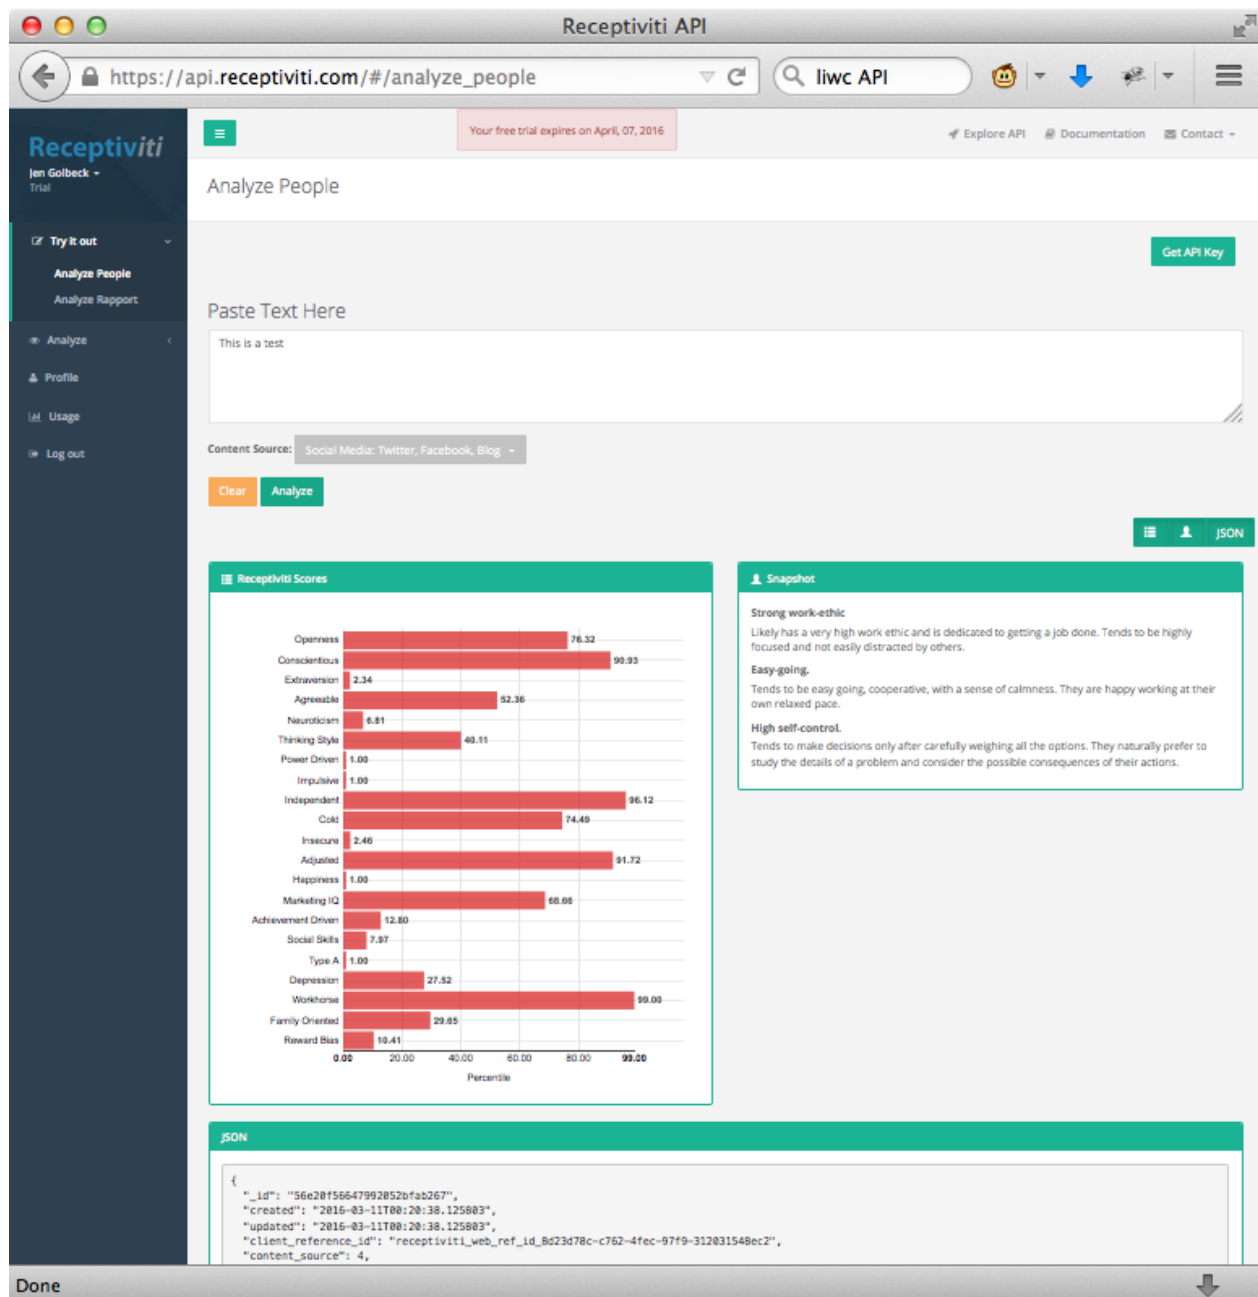


Figure 1. The Receptiviti API web interface, showing graph- and JSON-based results for a text sample.

The data sets were comprised of text posted by users on the social media platform. We aggregated each user's social media posts into a single "document". Because the Receptiviti API requires at least 300 words for statistical accuracy, we dropped all subjects who had fewer than 300 words. (We dropped 3 users from the *Twitter* data set, 193 from the 2011 *Facebook* data set, 122 from the *myPersonality Small* data set, and 1,314 from the *myPersonality Large* data set.) Details on each data set, including the number of users included in this study and the average number of words per user, are shown in Table 1. In total, we had data from nearly 9,000 social media users.

Because the Receptiviti API limits the maximum number of words per submission, we limited the word count to 10,000 for users who had provided a lot of text. For these cases, we kept the most recent 10,000 words. We then submitted the user-representing document to the Receptiviti API, selected the Receptiviti

option indicating the text is derived from social media, and extracted the raw Big Five values for each user.

Data Set	Number of Users*	Average # Words Per User
Facebook (myPersonality Small)	127	1,046
Facebook (myPersonality Large)	8,569	6,171
Facebook (2011)	69	463
Twitter	47	1,867

*Users with at least 300 words.

The personality values from Receptiviti were not on the same 1–5 scale as our input data; instead, each personality trait had its own scale (although still linear). We normalized the Receptiviti scores and converted the scores to the 1–5 scale². This allowed direct comparison between the known values and the predicted values for each subject.

For reference, Table 2 shows the mean scores—both actual and predicted by Receptiviti—for each personality trait in each data set.

Data Set	Scores	Open.	Consc.	Extra.	Agree.	Neuro.
Twitter	<i>Actual</i>	4.043	3.459	3.343	3.783	2.681
	<i>Receptiviti</i>	2.711	3.139	3.100	3.530	2.522
Facebook (myPersonality Small)	<i>Actual</i>	4.126	3.520	3.298	3.672	2.606
	<i>Receptiviti</i>	2.841	3.078	2.953	3.064	2.705
Facebook (2011)	<i>Actual</i>	3.793	3.399	3.315	3.752	2.750
	<i>Receptiviti</i>	2.820	3.209	3.303	3.548	2.563
Facebook (myPersonality Large)	<i>Actual</i>	2.807	3.669	3.575	3.461	3.917
	<i>Receptiviti</i>	2.494	2.832	3.049	2.796	2.891

Actual = Mean values for actual observed BFPI score
Receptiviti = Mean values for scores predicted by Receptiviti

3 Results

We computed the *Mean Absolute Error* (MAE), *Root Mean Squared Error* (RMSE), and *Pearson Correlation Coefficient* between the known and predicted values for each of the five attributes in each data set. Results are shown in Table 3. Note that the Big Five scores are on a 1–5 scale; thus the maximum possible error is 4, and an MAE of 1.0 represents an error of 25%.

All four data sets produced similar results. We found the error had a fairly wide range across the five personality traits. Openness was the most error-prone prediction, with an MAE greater than 1.0 (>25%). Conscientiousness and Neuroticism had the best prediction performance, with an MAE less than 20%. Correlations between predicted and known values for a given personality trait ranged from weak to near-zero.

² This was a simple linear conversion, and we verified this method with researchers at Receptiviti. We have chosen not to republish the internal documentation they shared with us.

Because the average number of words per user is low in the three smaller data sets—sometimes bordering the minimum Receptiviti requires for statistical significance—we wondered if the performance would be significantly better for users who had more text available for analysis. To test this, we compared performance on users with more than 2,000 words to those with fewer than 2,000 words. We found no statistical differences on any trait in any of the three smaller data sets (for $p < 0.05$). This is further supported by the fact that the large myPersonality data set, with higher average words per person, did not have better performance than the smaller data sets.

Table 3. Correlation and Error Values for Known and Predicted BFPI Scores

Data Set	Statistic	Open.	Consc.	Extra.	Agree.	Neuro.
Facebook (myPersonality Small)	MAE	1.279	0.728	0.891	0.689	0.539
	RMSE	1.540	1.006	0.908	1.012	0.837
	Correl.	0.358	0.249	0.368	0.412	0.377
Facebook (myPersonality Large)	MAE	1.450	0.790	0.872	0.971	0.678
	RMSE	1.596	0.957	1.052	1.177	0.836
	Correl.	0.204	0.202	0.220	0.239	0.182
Facebook (2011)	MAE	1.042	0.661	0.658	0.905	0.719
	RMSE	1.296	0.769	0.980	0.983	0.869
	Correl.	-0.347	-0.065	0.236	-0.354	-0.179

MAE = Mean Average Error
 RMSE = Root Mean Squared Error
 Correl. = Pearson Correlation Coefficient (between known and predicted values)

4 Discussion

From these four different data sets we obtained remarkably similar results. The Receptiviti API had similar performance on each personality attribute across the data sets. This suggests that the prediction algorithm is working consistently.

Overall, the MAE-based error rates generally fall in the range of 15–30%. Other personality prediction studies tended to have higher accuracy, with an MAE less than 0.5 (<12.5%) in most cases (Golbeck et al., 2011; Golbeck, Robles, & Turner, 2011; Quercia et al., 2011; Adali & Golbeck, 2012; Kosinski et al., 2013); however, they considered other information, not just text.

Given a “Best Guess” technique, which uses the mean observed personality score as the prediction for every user in a data set, the Best Guess technique significantly outperformed Receptiviti. For example, we computed this Best Guess for the *Facebook (myPersonality Large)* data set. The Best Guess MAE was significantly lower than the Receptiviti MAE ($p < 0.001$) in all cases, except for Neuroticism where there was no significant difference. Table 4 compares the MAE for both Receptiviti and Best Guess estimates.

Table 4. MAE Comparison of Receptiviti and “Best Guess” Predictions

Data Set	Open.	Consc.	Extra.	Agree.	Neuro.
Receptiviti	1.450	0.790	0.872	0.971	0.678
Best Guess	0.531	0.588	0.653	0.563	0.674

Predictions use the average observed personality score as the prediction for every user.
 Since this is a measure of error, lower scores are better.

This is not to suggest that the Receptiviti tool does not offer more insights than Best Guessing. Obviously, the correlations are stronger with the predicted values than with a Best Guess (which has a correlation of

0). Furthermore, predicting a precise value often may not be the main goal. One may want to compare the personalities of people in one group with another, in which case only the relative values are truly important. This replication study does not address that point in depth, however, we did undertake a preliminary analysis.

Related to this, one limitation of this study is that we aggregated social media posts to create “documents” of at least 300 words. Receptiviti may produce more accurate results when a single document or social media post of 300 words or more is analyzed. We did not have enough long social media posts to do a meaningful analysis of this type, but it is a question that could be analyzed in the future.

In our earlier work, we found a moderate correlation ($\rho \sim 0.3$) between extroversion scores and social network structure (Golbeck, Robles, & Turner, 2011). There was a positive correlation with a person’s Facebook friend count (extroverts tend to have more friends, as one might expect). We also found that social network density is negatively correlated with extroversion (extroverts’ friends know each other less often, perhaps because their friends are in many different social circles). When we checked for these relationships in the Receptiviti data, we found hints of the same patterns—those with more Facebook friends tended to have higher extroversion scores ($\rho \sim 0.2$), while those with higher social network density tended to have lower extroversion scores ($\rho \sim 0.13$). This is only a first glimpse at relative personality scores between groups. Systematically analyzing relative personality scores in social network data will require a deeper, careful analysis.

Acknowledgments

We would like to acknowledge and thank David Stillwell, Michal Kosinski, and the *myPersonality* project (mypersonality.org). They collected the myPersonality data sets referenced here.

References

- Adali, S., & Golbeck, J. (2012). Predicting personality with social behavior. *Proceedings of the 2012 International Conference on Advances in Social Networks Analysis and Mining*, 302–309.
- Adali, S., & Golbeck, J. (2014). Predicting personality with social behavior: a comparative study. *Social Network Analysis and Mining*, 4(1), 1–20.
- Alam, F., Stepanov, E. A., & Riccardi, G. (2013). Personality traits recognition on social network-facebook. *Proceedings of the Workshop on Computational Personality Recognition*, 6–9.
- Bachrach, Y., Graepel, T., Kohli, P., Kosinski, M., & Stillwell, D. (2014). Your digital image: factors behind demographic and psychometric predictions from social network profiles. *The 2014 International Conference on Autonomous Agents and Multi-agent Systems*, 1649–1650.
- Celli, F., Pianesi, F., Stillwell, D., & Kosinski, M. (2013). Workshop on computational personality recognition (shared task). *Proceedings of the Workshop on Computational Personality Recognition*.
- Farnadi, G., Zoghbi, S., Moens, M. F., & De Cock, M. (2013). Recognising personality traits using facebook status updates. *Proceedings of the Workshop on Computational Personality Recognition at the 7th International AAAI Conference on Weblogs and Social Media*.
- Ferwerda, B., Schedl, M., & Tkalcic, M. (2015). Predicting personality traits with instagram pictures. *Proceedings of the 3rd Workshop on Emotions and Personality in Personalized Systems*, 7–10.
- Golbeck, J., Robles, C., Edmondson, M., & Turner, K. (2011). Predicting personality from twitter. *Privacy, Security, Risk, and Trust*, 149–156.
- Golbeck, J., Robles, C., & Turner, K. (2011). Predicting personality with social media. *Extended Abstracts on Human Factors in Computing Systems*, 253–262.
- Gosling, S. D., Augustine, A. A., Vazire, S., Holtzman, N., & Gaddis, S. (2011). Manifestations of personality in online social networks: Self-reported facebook-related behaviors and observable profile information. *Cyberpsychology, Behavior, and Social Networking*, 14(9), 483–488.
- Iacobelli, F., & Culotta, A. (2013). Too neurotic, not too friendly: Structured personality classification on textual data. *Proceedings of the Workshop on Computational Personality Recognition*, 19–22.
- John, O. P., Donahue, E. M., & Kentle, R. L. (1991). The Big Five Inventory--Versions 4a and 54. *Journal of Research in Personality*, 37(6), 504–528.
- Kosinski, M., Stillwell, D., & Graepel, T. (2013). Private traits and attributes are predictable from digital records of human behavior. *The National Academy of Sciences*, 110(15), 5802–5805.
- Kosinski, M., Bachrach, Y., Kohli, P., Stillwell, D., & Graepel, T. (2014). Manifestations of user personality in website choice and behaviour on online social networks. *Machine Learning*, 95(3), 357–380.
- Markovikj, D., Gievska, S., Kosinski, M., & Stillwell, D. (2013). Mining facebook data for predictive personality modeling. *Proceedings of the 7th International AAAI Conference on Weblogs and Social Media*, Boston, MA.
- Park, G., Schwartz, H. A., Eichstaedt, J. C., Kern, M. L., Kosinski, M., Stillwell, D. J., ... & Seligman, M. E. (2015). Automatic personality assessment through social media language. *Journal of Personality and Social Psychology*, 108(6), 934.
- Quercia, D., Kosinski, M., Stillwell, D., & Crowcroft, J. (2011). Our Twitter profiles, our selves: predicting personality with twitter. *Privacy, Security, Risk, and Trust*, 180–185.
- Ryan, T., & Xenos, S. (2011). Who uses Facebook? An investigation into the relationship between the Big Five, shyness, narcissism, loneliness, and Facebook usage. *Computers in Human Behavior*, 27(5), 1658–1664.
- Staiano, J., Lepri, B., Aharony, N., Pianesi, F., Sebe, N., & Pentland, A. (2012). Friends don't lie: inferring personality traits from social network structure. *The 2012 ACM Conference on Ubiquitous Computing*, 321–330.

- Tausczik, Y. R., & Pennebaker, J. W. (2010). The psychological meaning of words: LIWC and computerized text analysis methods. *Journal of Language and Social Psychology, 29*(1), 24–54.
- Vazire, S., Gosling, S. D. (2004) e-Perceptions: Personality Impressions Based on Personal Websites. *Journal of Personality and Social Psychology, 87*(1), 123–132.

About the Authors

Jennifer Golbeck. Jennifer Golbeck is an associate professor in the College of Information Studies at the University of Maryland where she directs the Social Intelligence Lab.

Copyright © 2016 by the Association for Information Systems. 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 full citation on the first page. Copyright for components of this work owned by others than the Association for Information Systems must be honored. Abstracting with credit is permitted. To copy otherwise, to republish, to post on servers, or to redistribute to lists requires prior specific permission and/or fee. Request permission to publish from: AIS Administrative Office, P.O. Box 2712 Atlanta, GA, 30301-2712 Attn: Reprints or via e-mail from ais@aisnet.org.