



Personality Recognition from Facebook Text

Barbara Barbosa Claudino da Silva and Ivandré Paraboni^(✉)

School of Arts, Sciences and Humanities,
University of São Paulo, São Paulo, Brazil
{barbara.barbosa.silva, ivandre}@usp.br

Abstract. This work concerns a study in the Natural Language Processing field aiming to recognise personality traits in Portuguese written text. To this end, we first built a corpus of Facebook status updates labelled with the personality traits of their authors, from which we trained a number of computational models of personality recognition. The models include a range of alternatives ranging from a standard approach relying on lexical knowledge from the LIWC dictionary and others, to purely text-based methods such as bag of words, word embeddings and others. Results suggest that word embedding models slightly outperform the alternatives under consideration, with the advantage of not requiring any language-specific lexical resources.

Keywords: Big Five · Personality recognition

1 Introduction

The increasing complexity of computer systems has been accompanied by the development of ever more sophisticated human-machine communication methods. Current systems are capable of interpreting and reproducing a wide range of human behaviour, including emotions and feelings. These temporary manifestations of character are however heavily influenced by a stable set of patterns of human behaviour that are largely foreseeable. These patterns - or traits - constitutes what is generally understood as human personality [1].

The computational recognition of human personality is at the heart of the design of the so-called intelligent systems, and will be the focus of the present work as well. Fundamental personality traits may be recognised through a range of methods proposed in the Psychology field. Among these, the most popular are those based on the lexical hypothesis, which establishes that personality traits are observable in the words that we use to communicate. This approach has been refined from an initial survey of 4,500 traits identified in the 1930s to produce, independently and simultaneously in several studies, a stable framework known as the Big Five model of human personality [2].

The Big Five model comprises five key dimensions (or traits) - Openness to experience, Conscientiousness, Extraversion, Agreeableness and Neuroticism -

which are widely accepted as an adequate basis for the representation of human personality [3]. From a computational perspective, the Big Five model is central to human-computer interaction studies in general and, given its linguistic motivation, it makes also a suitable theoretical basis for Natural Language Processing (NLP) research. Knowing an individual's personality traits (e.g., from their social network updates) enables the production of personalised content in many ways, including the presentation of more appealing website or the generation of targeted advertisement, among many other possible applications.

Based on these observations, this paper presents a study of automatic Big Five personality recognition for the Brazilian Portuguese language. More specifically, we built a corpus of Facebook text labelled with personality information, and designed a number of experiments involving alternative text representations and supervised machine learning methods to recognise each of the Big Five traits. In doing so, our goal was to determine which representation and method would provide best results for our target language and domain.

The remainder of this article is organised as follows. Section 2 introduces a number of basic concepts related to the Big Five model and personality inventories, and briefly discusses the related work on Big Five recognition from text. Section 3 presents our current work, comprising the corpus construction and the experiments that were conducted. Section 4 presents our results, and Sect. 5 draws a number of conclusions and discusses further studies.

2 Related Work

The Big Five personality model [2] comprises five fundamental dimensions of human personality: Openness to experience, Conscientiousness, Extraversion, Agreeableness and Neuroticism. Each of the five dimensions is modelled as a scalar value representing the degree to which an individual expresses a given personality trait or not. Thus, for example, a high value for Extraversion indicates an extrovert individual, whereas a low value for this dimension indicates an introvert.

Big Five personality dimensions may be estimated by many well-known methods proposed in the Psychology field, being the most popular the use of inventories of personality such as the 44-item Big Five inventory - or BFI - that has become popular in Computer Science studies as well. The BFI was originally developed for the English language, but it was subsequently replicated in dozens of other languages, including Brazilian Portuguese. In particular, the study in [3] validated a Brazilian Portuguese version of the BFI called IGFP-5 by presenting a factorial analysis involving a sample of 5,089 respondents from the five regions of Brazil. This inventory will also be adopted in the present study, as discussed in Sect. 3.

The computational recognition of personality from text tends to follow a traditional methodology of supervised [4] or semi-supervised [5] machine learning. The task may be modelled as a classification problem (e.g., deciding whether an individual is an introvert or not), as a regression problem (e.g., determining

the scalar value of a dimension of personality), or as a ranking problem (e.g., ordering a set of individuals according to a dimension of interest.)

One of the first large-scale initiatives to recognise personality traits for the English language was the work in [6]. This consisted of an experiment involving 2,263 essays written by 1,200 students who completed a personality inventory, but it was limited to the lower and upper ends of the scale for Extraversion and Neuroticism. Essay words were grouped into four categories of well-defined psychological meaning: function (articles, prepositions, etc.), cohesion (demonstratives etc.), evaluation (terms that evaluate the validity, likelihood, acceptance, etc.) and judgement (terms that express the author’s attitude in relation to content.) The texts were represented by the frequencies of each category, and the binary classes Extraversion and Neuroticism were classified using SVMs, with a maximum accuracy of 58%.

In [7], an extended version of the same set of texts and inventories from [6] was considered. In this approach, learning features consisted of 88 word categories provided by the psycholinguistic dictionary LIWC (Linguistic Inquiry and Word Count) [8] and 26 attributes provided by the MRC (Medical Research Council) database [9] composed of 150,837 lexical items. An experiment was carried out to discriminate between the upper and lower ends of the Big Five dimensions, with maximum accuracy ranging from 50%–62% when using SVMs.

Studies as in [6,7] are based on word counts provided by psycholinguistic lexical resources. By contrast, studies as in [4,10] rely solely on the text itself by making use of n-grams models. In these studies, the objective was also to discriminate between individuals scores for four of the five dimensions of personality (except Openness to experience.) In both cases, Naive-Bayes and SVM classification were attempted. In [4] a set of 71 blogs was considered, with accuracy ranging from 45% (random) to 100% depending on the order of the n-gram model and class. In [10], the same experiment was repeated using a set of 1,672 blogs, with a maximum accuracy of 65%.

Finally, a note on NLP resources. Clearly, the computational problem of personality recognition from text is well developed for the English language. By making use of large scale resources such as the *myPersonality* corpus [11], the field has even experienced a number of dedicated scientific events in recent years, including the PAN-CLEF shared tasks series. In the case of the Portuguese language, by contrast, there is no obvious equivalent for the purpose of Big Five personality recognition, and we are not aware of any existing systems that may be regarded as a baseline.

3 Current Work

Given the lack of data and baseline systems for personality recognition in Brazilian Portuguese, we devised an exploratory study in which we first build a suitable corpus, and then we investigate a range of computational methods for the task. This study is described in the next sections.

3.1 Objectives

The objective of the present study is to develop supervised models of human personality recognition from Brazilian Facebook status updates, and to determine which of these models are more suitable for the task. In particular, we would like to investigate recent methods for text representation - namely, those based on word embeddings - as a possible alternative to standard personality recognition based on language-dependent psycholinguistic knowledge.

3.2 Data Acquisition

The computational models under discussion were built from a Facebook corpus labelled with Big Five information obtained from self-report IGFP-5 personality inventories [3]. To this end, a Facebook application was developed. The application requests users to fill in the 44 items of the IGFP-5 inventory as proposed in [12], and from which the five dimensions of personality are computed.

In addition to providing the personality inventories, the application simultaneously collects the user's status updates upon consent. Once the personality inventory is completed, a result page displayed details about the user's personality, and a brief explanation of each trait. The purpose of this page was however merely illustrative, that is, aiming to offer some kind of reward to the participant as a mean to possibly motivate them to further disseminate the Facebook application to their social circle.

As discussed in [7], the accuracy of this form of self-assessment is admittedly lower than third-party evaluation (i.e., performed by Psychology experts.) However, due to the costs of a large-scale professional evaluation of this kind, self-assessment remains the most common method in the field [6,7], and may be considered sufficient for the purposes of the present (exploratory) study as well.

We obtained data from 1,039 participants to create a corpus that is, to the best of our knowledge, the largest data set of this kind for the Brazilian Portuguese language. The corpus - hereby called *b5-post* - contains 2.2 million words in total, and it was subject to a number of pre-processing, spell-checking and normalisation procedures to be described elsewhere.

3.3 Computational Models

We follow a great deal of previous studies such as [4-7] in that we model personality recognition as five independent binary classification tasks, that is, one for each personality dimension of the Big Five model. This decision is motivated both by the type of application intended (i.e., we would like to recognise personality traits exclusively from text, and not from other already known traits), and also by the fact that the personality factors of the Big Five model are, by definition, highly independent [2].

In our models, individuals are classified as either positive or negative for each Big Five dimension based on the mean personality score for that class.

Table 1 presents the number and proportion of positive and negative individuals (or classification instances) in the corpus.

Table 1. Positive and negative learning instances

Class	Positive		Negative		Mean
Extraversion	412	52.9%	366	47.1%	3.15
Agreeableness	406	52.1%	372	47.9%	3.54
Conscientiousness	381	48.9%	397	51.1%	3.25
Neuroticism	389	50.0%	389	50.0%	3.17
Openness to exp.	408	52.4%	370	47.6%	3.78

As a means to provide an overview of possible strategies for personality recognition from text, we envisaged a range of models based on Random Forest classification with different text representations. These models are complemented by an alternative that makes use of long short-term memory neural networks (LSTMs). In the case of Random Forest, we use 10-fold cross validation over the entire dataset. In the case of the LSTMs models, the corpus was split into training (70%) and test (30%) subsets.

We carried out dozens of experiments involving alternative text representations and machine learning algorithms. For brevity, however, the present discussion will be limited to six of the best-performing alternatives and relevant baseline models, hereby called BoW, Psycholinguistics, word2vec-cbow-600, word2vec-skip-600, doc2vec and LSTM-600. Details are provided as follows.

- **BoW:** A bag-of-words model retaining 22,612 terms after lemmatisation.
- **Psycholinguistics:** LIWC-BR word counts [13] and psycholinguistic properties for Brazilian Portuguese [14]. LIWC categories include words that indicate emotions, social relations, cognitive processes, etc. Psycholinguistic properties include word age of acquisition, concreteness, etc.
- **word2vec-cbow-600:** A word2vec cbow model [15] of size 600, trained from a corpus of 50k tweets [16] using the vector component-wise average with the text corpus in lower case.
- **word2vec-skip-600:** A word2vec skip-gram model [15] of size 600, with the same features as the above cbow model.
- **doc2vec:** A doc2vec model [17] in which a document is defined as the set of all Facebook status updates written by each participant.
- **LSTM-600:** A Keras embedding model of size 600 built from the *b5-post* corpus, with combination provided by the LSTM hidden layers.

4 Results

The six models - and a Majority class baseline - were applied to the recognition of the five personality dimensions, resulting in 35 binary classifiers. Table 2 reports mean F1 scores for each model and class (i.e., each personality trait.)

Table 2. Mean F1 scores results

Model	EXT	AGR	CON	NEU	OPE
<i>Majority baseline</i>	0.33	0.34	0.33	0.32	0.34
BoW	0.60	0.53	0.57	0.53	0.55
Psycholinguistics	0.61	0.53	0.57	0.54	0.53
word2vec-cbow-600	0.61	0.56	0.58	0.55	0.59
word2vec-skip-600	0.61	0.55	0.59	0.55	0.57
doc2vec	0.60	0.58	0.57	0.55	0.55
LSTM-600	0.58	0.53	0.45	0.54	0.59

Based on these results, a number of observations are warranted. First, we notice that no single model is capable of providing the best results for all five classes. This may suggest that not all personality traits are equally accessible from text (or at least not in our domain.) Second, we notice that the Majority baseline and BoW never outperform the other models. This may suggest that the use of word embeddings is indeed a suitable approach to the task.

Looking at the classes individually, we notice that Extraversion results are slightly superior to those observed for the other classes, a result that has already been suggested in previous studies devoted to the English language [7]. As for the other classes, we notice that best results are divided between various methods, with a small advantage for the CBOW architecture.

5 Discussion

This article presented an exploratory study on the computational problem of human personality recognition from social network texts in the Brazilian Portuguese language. A corpus of texts labelled with personality information was collected, and subsequently used as training data for a range of supervised machine learning models of personality recognition.

Results suggest that different personality traits may be more or less evident from (Facebook) text, and that there is no single best-performing model for all traits. Despite our relatively small dataset, we notice that models based on word embeddings seem to outperform those based on lexical resources and, perhaps more importantly, we notice that these methods do not require language-specific resources such as psycholinguistic databases.

As future work, we consider improving the models based on word embeddings by making use of deep neural networks such as our current LSTM model. Despite the relatively weak results reported in our initial experiments, we believe that further fine-tuning of the network hyper parameters may provide more significant results in this regard.

The original *b5-post* corpus has been made publicly available for research, and has been reused on a number of related projects. Details regarding the corpus

are discussed in [18]. An experiment comparing personality recognition models based on Facebook and other textual sources is presented in [19]. The corpus has also been applied to the task of author profiling (i.e., for predicting author's gender, age group and others) in [20]. Finally, a pilot experiment investigating alternative models of personality appeared in [21].

Acknowledgements. The second author received supported from grant # 2016/14223-0, São Paulo Research Foundation (FAPESP).

References

1. Allport, F.H., Allport, G.W.: Personality traits: their classification and measurement. *J. Abnorm. Soc. Psychol.* **16**, 6–40 (1921)
2. Goldberg, L.R.: An alternative description of personality: the Big-Five factor structure. *J. Pers. Soc. Psychol.* **59**, 1216–1229 (1990)
3. de Andrade, J.M.: Evidências de validade do inventário dos cinco grandes fatores de personalidade para o Brasil. Ph.D. thesis, Universidade de Brasília (2008)
4. Oberlander, J., Nowson, S.: Whose thumb is it anyway? Classifying author personality from weblog text. In: COLING/ACL-2006 Poster Sessions, Sydney, Australia, Association for Computational Linguistics, pp. 627–634 (2006)
5. Celli, F.: Adaptive personality recognition from text. Ph.D. thesis, University of Trento (2012)
6. Argamon, S., Dhawle, S., Koppel, M., Pennebaker, J.W.: Lexical predictors of personality type. In: The Joint Annual Meeting of the Interface and the Classification Society of North America (2005)
7. Mairesse, F., Walker, M., Mehl, M., Moore, R.: Using linguistic cues for the automatic recognition of personality in conversation and text. *J. Artif. Intell. Res. (JAIR)* **30**, 457–500 (2007)
8. Pennebaker, J.W., Francis, M.E., Booth, R.J.: *Inquiry and Word Count: LIWC*. Lawrence Erlbaum, Mahwah (2001)
9. Coltheart, M.: The MRC psycholinguistic database. *Q. J. Exp. Psychol. Sect. A: Hum. Exp. Psychol.* **33**(4), 497–505 (1981)
10. Nowson, S., Oberlander, J.: Identifying more bloggers: towards large scale personality classification of personal weblogs. In: Proceedings of the International Conference on Weblogs and Social Media, Boulder, Colorado, USA (2007)
11. Kosinski, M., Matz, S., Gosling, S., Popov, V., Stillwell, D.: Facebook as a social science research tool: opportunities, challenges, ethical considerations and practical guidelines. *Am. Psychol.* **70**(6), 543–556 (2015)
12. John, O.P., Naumann, L.P., Soto, C.J.: *Paradigm Shift to the Integrative Big-Five Trait Taxonomy: History, Measurement, and Conceptual Issues*, pp. 114–158. Guilford Press, New York (2008)
13. Filho, P.P.B., Aluísio, S.M., Pardo, T.: An evaluation of the Brazilian Portuguese LIWC dictionary for sentiment analysis. In: 9th Brazilian Symposium in Information and Human Language Technology - STIL, Fortaleza, Brazil, pp. 215–219 (2013)
14. dos Santos, L.B., Duran, M.S., Hartmann, N.S., Candido, A., Paetzold, G.H., Aluísio, S.M.: A lightweight regression method to infer psycholinguistic properties for Brazilian Portuguese. In: Ekštejn, K., Matoušek, V. (eds.) TSD 2017. LNCS (LNAI), vol. 10415, pp. 281–289. Springer, Cham (2017). https://doi.org/10.1007/978-3-319-64206-2_32

15. Mikolov, T., Wen-tau, S., Zweig, G.: Linguistic regularities in continuous space word representations. In: Proceedings of NAACL-HLT-2013, Atlanta, USA, pp. 746–751. Association for Computational Linguistics (2013)
16. Ramos Casimiro, C., Paraboni, I.: Temporal aspects of content recommendation on a microblog corpus. In: Baptista, J., Mamede, N., Candeias, S., Paraboni, I., Pardo, T.A.S., Volpe Nunes, M.G. (eds.) PROPOR 2014. LNCS (LNAI), vol. 8775, pp. 189–194. Springer, Cham (2014). https://doi.org/10.1007/978-3-319-09761-9_20
17. Le, Q., Mikolov, T.: Distributed representations of sentences and documents. In: Proceedings of Machine Learning Research, PMLR, Beijing, China, vol. 32, no. 2, pp. 1188–1196 (2014)
18. Ramos, R.M.S., Neto, G.B.S., Silva, B.B.C., Monteiro, D.S., Paraboni, I., Dias, R.F.S.: Building a corpus for personality-dependent natural language understanding and generation. In: 11th International Conference on Language Resources and Evaluation (LREC-2018), ELRA, Miyazaki, Japan, pp. 1138–1145 (2018)
19. dos Santos, V.G., Paraboni, I., Silva, B.B.C.: Big Five personality recognition from multiple text genres. In: Ekštejn, K., Matoušek, V. (eds.) TSD 2017. LNCS (LNAI), vol. 10415, pp. 29–37. Springer, Cham (2017). https://doi.org/10.1007/978-3-319-64206-2_4
20. Hsieh, F.C., Dias, R.F.S., Paraboni, I.: Author profiling from facebook corpora. In: 11th International Conference on Language Resources and Evaluation (LREC-2018), ELRA, Miyazaki, Japan, pp. 2566–2570 (2018)
21. Silva, B.B.C., Paraboni, I.: Learning personality traits from Facebook text. IEEE Lat. Am. Trans. **16**(4), 1256–1262 (2018)