Vinciarelli, A., and Mohammadi, G. (2014) More personality in personality computing. IEEE Transactions on Affective Computing, 5 (3). pp. 297-300. ISSN 1949-3045

Copyright © 2014 IEEE

A copy can be downloaded for personal non-commercial research or study, without prior permission or charge

Content must not be changed in any way or reproduced in any format or medium without the formal permission of the copyright holder(s)

When referring to this work, full bibliographic details must be given

<http://eprints.gla.ac.uk/100508>

Deposited on: 07 January 2015

Enlighten – Research publications by members of the University of Glasgow_
<http://eprints.gla.ac.uk>

More Personality in Personality Computing

Alessandro Vinciarelli *Member, IEEE*, and Gelareh Mohammadi

Abstract—By explicitly describing what has been done in the past, surveys implicitly outline what can (and sometimes should) be done in the future. The insightful commentary by Wright contributes significantly to this latter aspect, especially when it comes to aligning Personality Computing with the latest developments in Personality Science. This response article tries to progress in such a direction by discussing on Wright's suggestions from a computing science point of view.

Index Terms—Personality, Automatic Personality Perception, Automatic Personality Recognition, Automatic Personality Synthesis

1 INTRODUCTION

IT is a pleasure for us to reflect on Wright's insightful commentary. His introduction to the latest developments in Personality Psychology helps us to see more clearly where Personality Computing is and where it can go. In particular, we find it interesting that Wright often adopts the expression *Personality Science* rather than the more traditional *Personality Psychology*. It encourages us to think of personality as a common ground where multiple disciplines, including computing and psychology, can contribute and mutually benefit from each other: progress in personality theory should help to build more effective personality machines and vice versa.

Our Survey of Personality Computing [1] provides an extensive snapshot of the state-of-the-art¹. Furthermore, it shows that three main problems - automatic recognition, perception and synthesis of personality traits - encompass the wide spectrum of approaches and scenarios considered in the domain. Wright's commentary [2] focuses on Personality Theory issues that so far have been neglected in Personality Computing, but need to be considered for the field to progress. In particular, he attracts attention to the hierarchical structure of personality traits, the problem of "true" personality, the person-situation debate, and the stability of traits over time.

It is interesting to observe that many issues in Wright's commentary seem to have a parallel in computing science. For example, *deep learning* approaches - statistical models aimed at representing data at different levels of abstraction [3], [4] - might be suitable for capturing the hierarchical structure underlying personality traits. Furthermore, the person-situation debate echoes, at least in its general aspects, the problem of modeling *context* [5],

probably one of the oldest problems of computing science that still waits a satisfactory conceptual framework.

The above confirms the potential of personality as an optimal ground for interdisciplinary work, but an important caveat should be pointed out: There is often a wide gap between low-level information accessible to computers - physical measurements extracted from sensor data or observable cues generated artificially - and high-level information like personality traits. The consequence is that computing approaches act sometimes as a *bottleneck* with respect to richness and complexity of psychological theories, i.e. they require simplifications of reality in order to work.

Our survey provides an interesting example of this aspect: while personality scores account for a *continuum* between opposite extremes (e.g., introversion and extraversion), the actual task performed by many approaches is to predict whether someone is above or below an arbitrary threshold score². In other words, there are many cases where the gap between data and traits is so wide, that machine intelligence methodologies (signal processing, machine learning, natural language processing, etc.) cannot achieve satisfactory results without converting personality scores into binary variables, easier to deal with³.

This *bottleneck effect* probably contributes to why Personality Computing appears to lag behind the latest developments in Personality Theory. It is true that it takes time for advances in Psychology to reach the computing community, as Wright correctly points out, but it is true as well that the integration of new, more advanced personality models does not necessarily result into tangible technological improvements. Furthermore, not every theoretic finding is compatible with the constraints imposed by the use of technology.

Still, Wright's commentary convincingly shows that we need more personality in Personality Computing. So far, the efforts have been technology driven, i.e. based

-
- A.Vinciarelli is with University of Glasgow and Idiap Research Institute.
E-mail: vincia@dcs.gla.ac.uk
 - Gelareh Mohammadi is with University of Geneva.
E-mail: gelareh.mohammadi@unige.ch

G.Mohammadi was with Idiap Research Institute and Ecole Polytechnique Fédérale de Lausanne when the work was done.

1. Overall, the survey describes in detail 81 works published in the computing literature.

2. Typically average or median of the scores observed in a corpus.

3. The survey describes 55 works on Automatic Personality Recognition and Perception and the binarization of the personality scores is performed in 27 cases.

on the computing component of the problem, whether this means to find better features, to develop more effective machine learning methodologies, to create more realistic artificial agents, to build more pervasive sensors or simply to collect more data. Wright's commentary points out that misconceptions about personality might limit technology effectiveness as well. Furthermore, it highlights research directions along which Personality Computing improvements can result from a deeper understanding of Personality Science.

In light of the above, this response article tries to establish links, whenever possible, between Wright's considerations and relevant issues in computing science. Our hope is to further contribute to the debate while suggesting new research directions in an area that we consider not only crucial for bridging the gap between humans and machines, but also fascinating and intellectually stimulating.

2 FOCUS ON OUTCOMES

One of the main messages of Wright's commentary is that personality traits are assessed because they are predictive of *outcomes of interest*, i.e. observable, possibly measurable aspects in the life of an individual (e.g., *"happiness, physical and psychological health, spirituality, and identity at an individual level; [...] quality of relationships with peers, family, and romantic others at an interpersonal level; [...] occupational choice, satisfaction, and performance, as well as community involvement, criminal activity, and political ideology at a social institutional level"* [6]).

From this point of view, the situation in Personality Computing is different whether we consider Automatic Personality Synthesis (APS) or the other two tasks outlined in our survey, namely Automatic Personality Recognition (APR) and Automatic Personality Perception (APP). Most APS works synthesize personality traits because these are predictive of users' reactions and behavior (see, e.g., [7], [8]). In the case of APR and APP, the works that involve the prediction of an outcome of interest are still an exception (see, e.g., [9], [10]). In most cases, APR and APP approaches predict personality traits independently of any outcome of interest.

The main reason behind such a state-of-affairs is probably that the gap between low-level features and personality traits (see Section 1) is difficult to bridge for current APR and APP technologies. In particular, the international benchmarking campaigns organized so far [11], [12] show that a wide range of advanced methodologies (signal processing, machine learning, natural language processing, etc.) still have difficulties in reliably predicting traits. Therefore, predicting outcomes of interest is probably a step out of the reach for the moment.

However, the efforts so far were based solely on the computing side. To the best of our knowledge, no attempts were made to refine Personality Science aspects. Wright's suggestions help to move in this latter direction

and this, besides making the field more mature from a methodological point of view, might be a source of performance improvement for Personality Computing.

2.1 Meta-Traits, Traits or Sub-Traits?

Most of the works described in our survey (76 out of 81) represent personality in terms of the basic Big-Five traits (*Openness, Conscientiousness, Extraversion, Agreeableness, Neuroticism*). The assessments are performed using the questionnaires most commonly applied in the literature, sometimes in their short version (e.g., see [13]) to limit as much as possible the amount of work needed to assess a large number of individuals. In this way, it was possible to perform experiments over large numbers of subjects, a crucial need for approaches based on statistics and machine learning.

Such a basic setup for personality representation and measurement has allowed the computing community to investigate the very possibility of performing tasks like APS, APR and APP. Furthermore, it has allowed the organization of large benchmarking campaigns leading to the first rigorous assessment of the state-of-the-art in the field (at least for specific tasks and types of data) [11], [12]. However, Wright points out that at least an important methodological issue has been missed: *"depending on the use for which a computer task or assessment is being designed, it is important to be precise about the level of the hierarchy one is interested in measuring and the inventory chosen to validate it"*.

Including an outcome of interest in the picture (see above) determines the need for selecting the right level of hierarchy (meta-traits, traits or sub-traits) and, correspondingly, for identifying the most suitable assessment questionnaire and scenario: *"[...] specific tasks are likely to capture specific facets of a given domain, but not all tasks may be fungible for all features. Relatedly, using brief personality trait inventories that only capture narrow definitions of the traits could lead one to develop a computer task that does not generalize well"*.

Ensuring that questionnaire, hierarchical level and outcome of interest are correctly aligned with each other is an issue that has been neglected in Personality Computing. Addressing these problems will make the field more correct from a methodological point of view and will probably help to improve the performance of current approaches. However, the *bottleneck effect* mentioned in Section 1 might play an important role here. As the difference between sub-traits becomes subtler, computing approaches might not have the necessary discrimination power.

On the other hand, a contribution towards addressing the problem might come from one of the most important machine learning trends of the last decade, i.e. the development of *deep learning* architectures capable of representing data at different levels of abstraction [3]. The ability of these architectures in extracting features from the data [4], i.e. measurements accounting for the

information the data contains, might prove effective not only in addressing Wright's indications, but also in helping Personality Psychology to progress on a data driven basis. In much the same way they were able to detect speech phonemes in an unsupervised fashion [14], deep neural networks might be capable of detecting sub-traits and meta-traits in behavioral data. Furthermore, the analysis of large volumes of data might reveal trait facets that so far were not accessible to psychological inquiry.

2.2 Assessments or Self-Assessments?

Our survey distinguishes between the prediction of self-attributed traits (APR), and the prediction of traits attributed by others (APP). Besides reflecting the state-of-the-art, the distinction aims at showing that the two tasks are different, should not be confused with each other and are both worth being addressed because they account for different phenomena, i.e. externalization and attribution.

APR was described as the task aimed at predicting the traits that the literature traditionally considered true (according to [13]). Such a distinction is particularly appealing in the computing community that is used to problems (e.g., Object Recognition or Optical Character Recognition) where data can be rated in objective and virtually unambiguous terms. For this reason, Wright's observation that *"there is no universal 'true' personality rating, and instead the question becomes for what purpose?"* is important not only because it helps to eliminate a misconception, the existence of *true* ratings, but also because it proposes a criterion to decide whether it is more appropriate to perform APR or APP, namely the outcome of interest to be predicted. Assessed traits are more predictive of certain outcomes while self-assessed ones are more predictive of others. Therefore, the choice between APR and APP depends on the final purpose of personality assessment.

3 FOCUS ON CONTEXT

Another important issue pointed out in Wright's commentary is the interplay between personality and situation (or context): *"The view here is that personality may have certain main effects, but it is often contextualized behavior we are interested in"*. The problem of context has been addressed extensively in the computing community as well⁴. The main goal is to make technology *context aware*, i.e. capable of working and interacting differently depending on the context (e.g., a phone should not ring during a meeting). The key-issue is how to encode and represent context, whether this means to identify a *"set of features of the environment surrounding generic activities"* [5] or understand *"how and why, in the course of their interactions, [...] people achieve and maintain a mutual*

Fig. 1. The plot shows the performance of the Speaker Trait Challenge participants over the different traits.

understanding of the context for their actions" [Ibid.]. To the best of our knowledge, the issue is still open for computing technologies in general and, in particular, for technologies dealing with social and psychological phenomena like personality [15].

Besides the difficulties in representing context, current approaches for human behavior understanding (e.g., facial expression analysis, speaker diarization, action recognition, etc.) are still sensitive to factors like illumination changes, environmental noise or sensor placement. Hence, technology is not sufficiently robust, at least on average, to physical changes in context. Due to the problems above, most of the Personality Computing work done so far has addressed the problem implicitly, i.e. by working on data where the context is relatively stable (e.g., meetings, radio programs, laboratory sessions, etc.) and the variability related to such a factor is limited.

The limitations above are a considerable *bottleneck* towards the possibility of including context in the Personality Computing picture. However, two recent technological advances might help further progress with respect to such a situation: the diffusion of mobile devices equipped with multiple sensors [16] and the advent of Big-Data analytics [17]. Mobile devices can collect a large amount of contextual information (geographic position, proximity to other people, audio environment, etc.) for extended periods of time. Big-Data analytics can make sense of the data and might provide information about context and its effect on behavior. The potential of these technologies for Personality Computing has already been shown [18], [19]. However, no results were presented so far that take context into account and, furthermore, not every sensor necessary for human behavior understanding can actually be mounted on a mobile platform, another *bottleneck effect* that might limit the integration of Personality Science findings in computing.

3.1 All Traits in All Contexts?

One of Wright's indications is that *"[...] if a particular trait is of interest it must be given the opportunity to express itself"*. The context plays a major role from this point of view because it allows certain traits to emerge while keeping the others hidden or, at least, difficult to observe (provided that the subjects manifest a trait when they have it). While APS approaches seem to be aware of this point and tend to focus on Extraversion - probably because it is the trait most likely to make a difference for interactive technologies - APP and APR experiments often aim at predicting all Big-Five traits independently

4. The ACM Digital Library contains 13,369 articles, published between 1962 and 2014, where the word *"context"* appears in the title. In the case of IEEEExplore, the articles are 8,157, published between 1960 and 2014.

of the context where the data has been collected⁵. Furthermore, all traits tend to be addressed in the same way independently of how evident are the markers they leave in overt behavior.

One important reason behind the practice above is probably that approaches developed for one trait can easily be extended to the other traits (as long as the ratings are available). However, the notions of *relevance* and *availability* [20] pointed out by Wright can still help under at least two main respects. The first is that they can drive the design of experiments targeting specific traits (or trait-facets). The second is that they can help to interpret the results, especially when it comes to APP and APR.

Figure 1 provides an example of this latter possibility. The plot shows the results of the *Interspeech Speaker Trait Challenge* for the different participants (the larger the bubble the better the performance). The data consists of short speech samples extracted from radio news and the theory suggested by Wright [20] can help to understand why the results are more satisfactory for Extraversion and Conscientiousness. In turn, plots like the one in Figure 1 might help to quantify variables and factors that the theory mentioned above introduces in abstract terms.

3.2 Static or Dynamic?

The tight interplay between context and personality leads to another important issue highlighted in Wright's commentary, namely whether personality should be considered a stable construct or a process that involves changes and evolution over time: "*depending on how it is measured and aggregated, personality appears both highly stable and trait like, and also highly variable and adaptive to context*". The large majority of Personality Computing approaches considers personality stable at least at the time scale of the data used in the experiments (ranging between a few seconds and several weeks).

However, a few works take into account the possibility of observing personality changes during face-to-face interactions. In particular, the work in [21] adopts the concept of *personality state* to explain changes in attributed traits during meetings. Furthermore, at the moment this article is being written, an international benchmarking campaign is ongoing on prediction of various traits (including the Big-Five) assessed at different time steps⁶.

Both cases above consider controlled settings, but nowadays wearable sensors allow one to perform the same experiments in everyday life. As correctly pointed out by Wright, this is likely to become one of the most fertile grounds for collaboration between psychology and computing science. However, one issue seems to remain open, i.e. whether personality evolves according

to an underlying dynamics or not. In the first case, it will be possible to develop approaches capturing principles and laws behind the temporal changes (i.e., via temporal series analysis or chaos theory). In the second case, modeling changes might simply mean to apply prediction approaches like those adopted today at regular time steps.

4 CONCLUSIONS

Overall, the main contribution of Personality Computing so far has been to show that tasks like Automatic Personality Synthesis, Perception and Recognition are possible. Extensive experiments - including benchmarking campaigns where multiple approaches are compared according to rigorous protocols - provide solid evidence that personality traits can be inferred from measurable aspects of behavior. Furthermore, research on personality synthesis has shown that artificially generated cues can actually convey personality impressions.

With a few exceptions, all attempts to progress so far were based on the improvement of the sole computing methodologies (signal processing, machine learning, artificial agents, etc.). More work is still possible and needed in such a direction and, for example, APP approaches should predict how attributed traits are distributed rather than simply predicting the average of attributed traits.

However, Wright's commentary convincingly shows that further progress cannot be achieved without taking into account the findings of Personality Science. This response article is an attempt to move in such a direction while still identifying potential obstacles and problems. Needless to say, our considerations were made through the eyes of a computer scientist, curious about personality, but not fully competent about it. Everything we said should not be considered a conclusive answer, but an invitation to delve further in the problems pointed out by Wright. We welcome with excitement his encouragement in this sense, we definitely need more personality in Personality Computing!

ACKNOWLEDGMENTS

The research that has led to this work has been supported in part by the European Community's Seventh Framework Programme (FP7/2007-2013), under grant agreement no. 231287 (SSPNet), in part by the Swiss National Science Foundation via the National Centre of Competence in Research IM2 (Information Multimodal Information management) and Indo-Swiss Joint Research Project CCPP (Cross-Cultural Personality Perception).

REFERENCES

- [1] A. Vinciarelli and G. Mohammadi, "A survey of personality computing," *IEEE Transactions on Affective Computing*, 2014.

5. Our survey describes 55 works on APP and APR and 47 of them take all the traits into account.

6. emotion-research.net/sigs/speech-sig/maptraits

- [2] A. Wright, "Current directions in personality science and the potential for advances through computing," *IEEE Transactions on Affective Computing*, 2014.
- [3] Y. Bengio, "Learning deep architectures for AI," *Foundations and Trends in Machine Learning*, vol. 2, no. 1, pp. 1–127, 2009.
- [4] Y. Bengio, A. Courville, and P. Vincent, "Representation learning: A review and new perspectives," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 35, no. 8, pp. 1798–1828, 2013.
- [5] P. Dourish, "What we talk about when we talk about context," *Personal and Ubiquitous Computing*, vol. 8, no. 1, pp. 19–30, 2004.
- [6] D. Ozer and V. Benet-Martinez, "Personality and the prediction of consequential outcomes," *Annual Reviews of Psychology*, vol. 57, pp. 401–421, 2006.
- [7] C. Nass and K. Min Lee, "Does computer-synthesized speech manifest personality? experimental tests of recognition, similarity-attraction and consistency-attraction," in *Journal of Experimental Psychology: Applied*, vol. 7, no. 3, pp. 171–181, 2001.
- [8] S. Woods, K. Dautenhahn, C. Kaouri, R. Boekhorst, and K. Koay, "Is this robot like me? Links between human and robot personality traits," in *IEEE-RAS International Conference on Humanoid Robots*, 2005, pp. 375–380.
- [9] M. Tkalcic, T. Tasic, and J. Kosir, "Emotive and personality parameters in multimedia recommender systems," in *Proceedings of the IEEE International Conference on Affective Computing and Intelligent Interaction*, 2009.
- [10] G. Mohammadi, S. Park, K. Sagae, A. Vinciarelli, and L.-P. Morency, "Who is persuasive? The role of perceived personality and communication modality in social multimedia," in *Proceedings of the ACM on International Conference on Multimodal Interaction*, 2013, pp. 19–26.
- [11] B. Schuller, S. Steidl, A. Batliner, E. Noth, A. Vinciarelli, F. Burkhardt, R. V. Son, F. Weninger, F. Eyben, T. Bocklet, G. Mohammadi, and B. Weiss, "The INTERSPEECH 2012 speaker trait challenge," in *Proceedings of Interspeech*, 2012.
- [12] F. Celli, F. Pianesi, D. Stillwell, and M. Kosinski, "Workshop on Computational Personality Recognition (shared task)," in *Proceedings of the Workshop on Computational Personality Recognition*, 2013.
- [13] B. Rammstedt and O. John, "Measuring personality in one minute or less: A 10-item short version of the Big Five Inventory in English and German," *Journal of Research in Personality*, vol. 41, no. 1, pp. 203–212, 2007.
- [14] H. Lee, P. Pham, Y. Largman, and A. Ng, "Unsupervised feature learning for audio classification using convolutional deep belief networks," in *Proceedings of Conference on Neural Information Processing Systems*, 2009, pp. 1096–1104.
- [15] A. Vinciarelli, M. Pantic, and H. Bourlard, "Social Signal Processing: Survey of an emerging domain," *Image and Vision Computing Journal*, vol. 27, no. 12, pp. 1743–1759, 2009.
- [16] M. Raento, A. Oulasvirta, and N. Eagle, "Smartphones an emerging tool for social scientists," *Sociological Methods & Research*, vol. 37, no. 3, pp. 426–454, 2009.
- [17] M. Minelli, M. Chambers, and A. Dhiraj, *Big Data, Big Analytics: Emerging Business Intelligence and Analytic Trends for Today's Businesses*. John Wiley & Sons, 2012.
- [18] D. Olguin-Olguin, P. Gloor, and A. Pentland, "Capturing individual and group behavior with wearable sensors," in *Proceedings of AAAI Spring Symposium on Human Behavior Modeling*, 2009.
- [19] J. Staiano, F. Pianesi, B. Lepri, N. Sebe, N. Aharoni, and A. Pentland, "Friends don't lie - Inferring personality traits from social network structure," in *Proceedings of the ACM International Conference on Ubiquitous Computing*, 2012, pp. 321–330.
- [20] D. Funder, "On the accuracy of personality judgment: a realistic approach," *Psychological Review*, vol. 102, no. 4, p. 652, 1995.
- [21] J. Staiano, B. Lepri, R. Subramanian, N. Sebe, and F. Pianesi, "Automatic modeling of personality states in small group interactions," in *Proceedings of the ACM International Conference on Multimedia*, 2011, pp. 989–992.

PLACE
PHOTO
HERE

Alessandro Vinciarelli is with the University of Glasgow where he is Senior Lecturer (Associate Professor) at the School of Computing Science and Associate Academic at the Institute of Neuroscience and Psychology (www.dcs.gla.ac.uk/vincia). His main research interest is in Social Signal Processing, the domain aimed at modelling analysis and synthesis of nonverbal behaviour in social interactions. In particular, Alessandro has investigated approaches for role recognition in multiparty conversations, automatic personality perception from speech, and conflict analysis and measurement in competitive discussions. Overall, Alessandro has published more than 100 works, including one authored book, five edited volumes, and 26 journal papers. Alessandro has participated in the organization of the IEEE International Conference on Social Computing as a Program Chair in 2011 and as a General Chair in 2012, he has initiated and chaired a large number of international workshops, including the Social Signal Processing Workshop, the International Workshop on Socially Intelligence Surveillance and Monitoring, the International Workshop on Human Behaviour Understanding, the Workshop on Political Speech and the Workshop on Foundations of Social Signals. Furthermore, Alessandro is or has been Principal Investigator of several national and international projects, including a European Network of Excellence (the SSPNet, www.sspnet.eu), an Indo-Swiss Joint Research Project (<http://www.idiap.ch/project/ccpp/>) and an individual project in the framework of the Swiss National Centre of Competence in Research IM2 (www.im2.ch). Last, but not least, Alessandro is co-founder of Klewel (www.klewel.com), a knowledge management company recognized with several awards.

PLACE
PHOTO
HERE

Gelareh Mohammadi graduated with a B.Sc degree in Biomedical Engineering from Amir-kabir University of Technology, Iran and a M.Sc. degree in Electrical Engineering from Sharif University of Technology, Iran in 2003 and 2006 respectively. She received a PhD in Electrical Engineering from Swiss Federal Institute of Technology in Lausanne (EPFL), Switzerland in 2013 where the main focus of her research was on personality perception from nonverbal vocal behavior. During 2012 she was a research assistant at Institute for Creative Technology (ICT), University of Southern California. From 2013 to 2014 she was a post-doctoral researcher at Idiap Research Institute in Martigny, Switzerland and she is currently a post-doctoral researcher at Behavioral Neurology and Imaging of Cognition (LabNIC) group, Department of Neuroscience, University of Geneva, Switzerland where she is working on analyzing the neural basis of emotions. Her research interests include social signal processing, affective computing, computational neuropsychology and neural analysis of behavior.