pera

Why language?
Language is one of the greatest abilities of humans. Language contains dozens of conscious and subconscious markers that act as a mirror of your life. The words, style or tone that people use, or not use, to express themselves contain a wealth of information.

Natural language processing is a branch of artificial intelligence the enables computers to understand and interpret human language and has been used in countless applications to get insights from texts or their authors.

Pera is at the forefront of these developments, and we are the first to use natural language processing to identify potential in people based on open-ended questions in a digital interview. This provides great insights into people's potential, while also providing a great candidate experience at the same time.

Why does it work?
We have conducted millions of digital interviews to collect language from employees and candidates. Pera Labs has spent years of research and development, working together with the lab of Prof. Daelemans from the University of Antwerp, to develop the methodology that drives the digital interview.

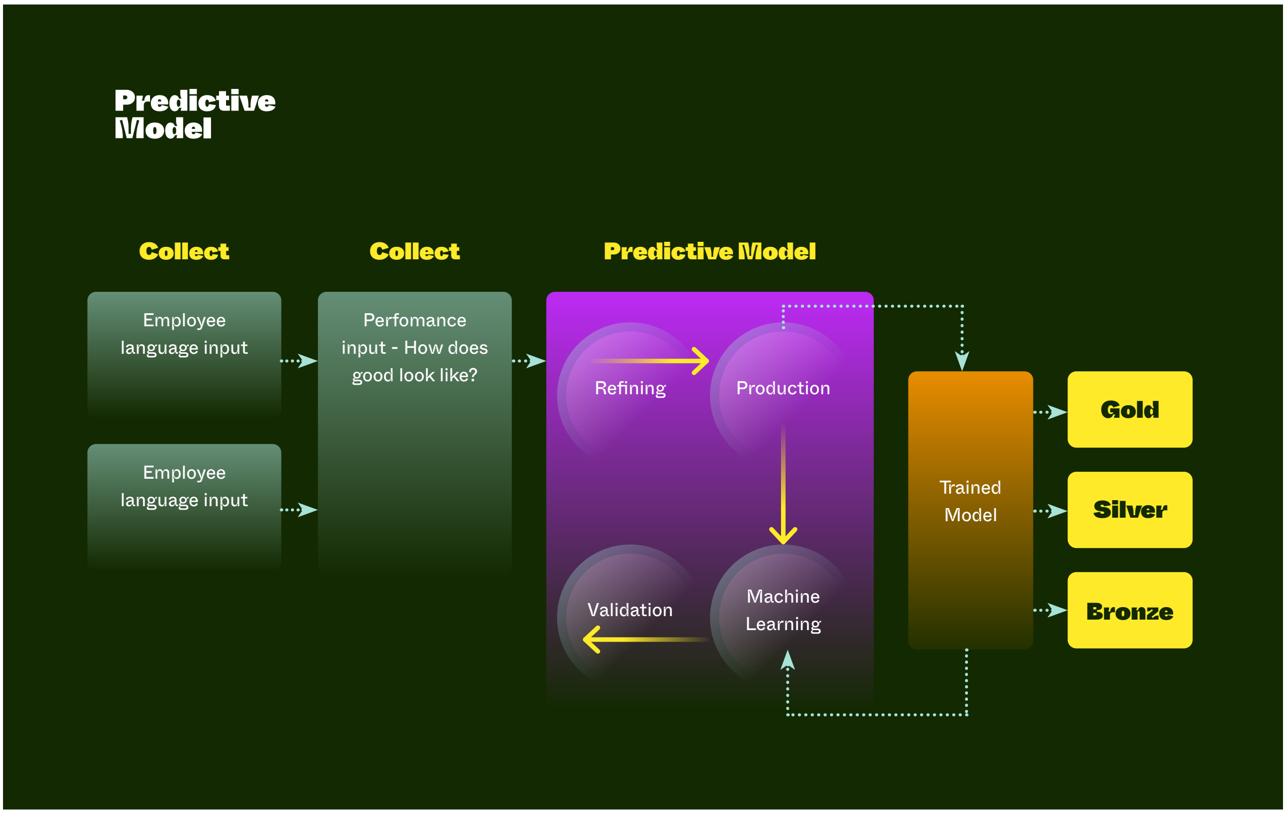

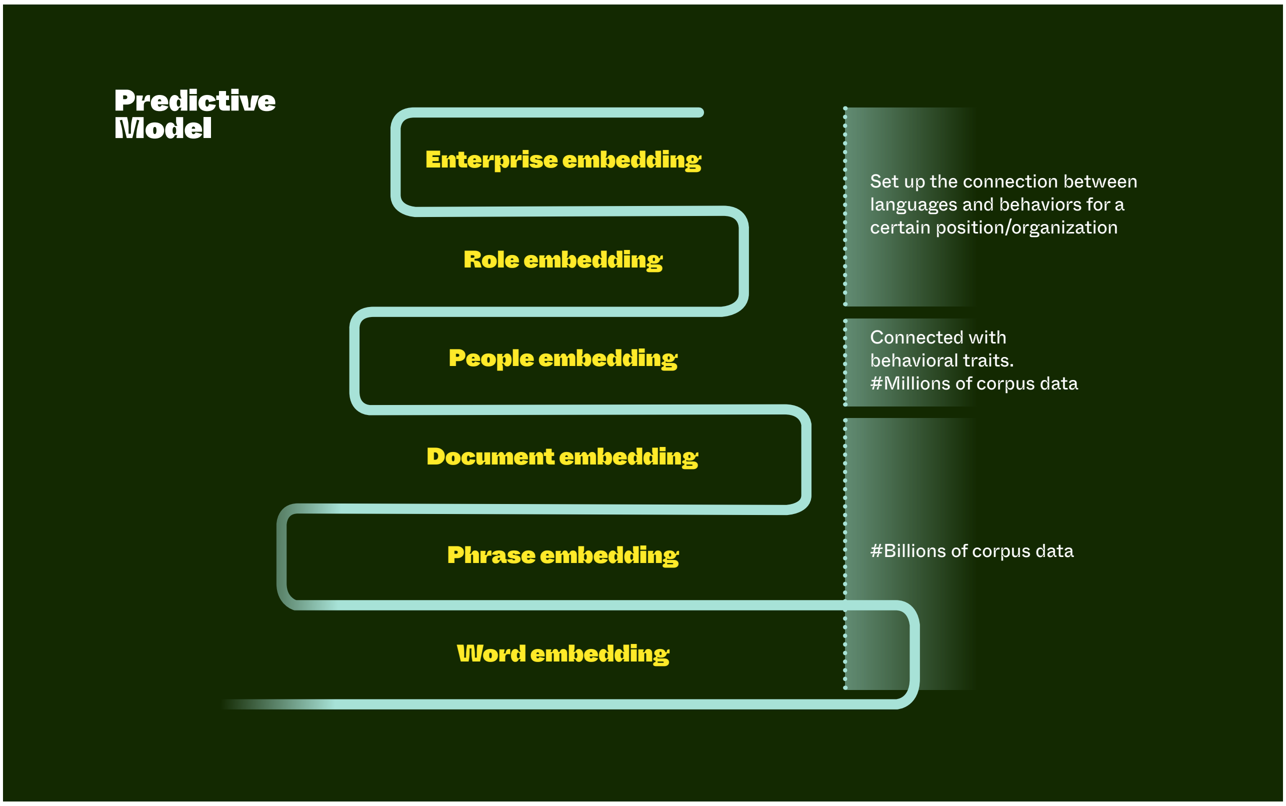

Unbiased by design
The approach we take to train our models is robust against bias.
We have taken several precautions to fight bias in our models:

- Our algorithms don't look at gender or cultural background and they also don't look at characteristics in your language that show strong correlations to your gender or cultural background.
- Training labels are collected with great care. The competency labels that serve as input to our algorithms are collected from employees with at least 6 months of working experience and (ideally) rated by both their managers and peers on competencies that lead to success in their job or organization.
- We evaluate if our models are fair. Gender bias is never found. We have investigated the impact of different cultures and countries in cooperation with Gert Jan Hofstede.

Relevant section from white paper

Introduction
Pera has developed technology that helps organizations improve their selection process by overcoming the challenges of human bias and siloed data. The technology relies heavily on the consensus that personality traits and competencies are important predictors for job performance (e.g. Barrick and Mount, 1991; Hurtz and Donovan, 2000; Jackson and Rothstein 1991), and person-organization fit (e.g. O'Reilly III et al., 1991; Kristof-Brown et al., 2005; and Gardner, et al. 2012).

Traditionally, psychometric personality questionnaires would be administered to candidates to evaluate their personality traits. Obviously, when a job is at stake, the responses to such questionnaires may not be entirely truthful, because responses may reflect a candidate's perception of the ideal candidate rather than them-selves (Furnham, 1990).

To make personality assessments more robust to dishonest answers, Pera's technology aims to predict these personality traits and competencies from answers to open-ended questions. Recent advances in natural language processing and machine learning have enabled efficient and reliable estimation of relevant traits and competencies from linguistic markers subconsciously produced in a person's language.

Personality trait prediction from text

Pennebaker and King (1999) were the first to investigate correlations between frequencies of word categories (e.g. positive emotion words, negative emotion words, pronouns) and personality traits. Using multiple writing samples of several hundred college students they found modest correlations to self-reports of Big Five personality dimensions. Their approach of using word categories to analyse texts became known as Linguistic Inquiry and Word Count (LIWC) and has since become a popular approach to study associations between personality and language use in different contexts, including directed writing assignments (Hirsh and Peterson, 2009), recording of day-to-day speech (Mehl et al., 2006), structured interviews (Fast and Funder, 2008), and online blogs (Yarkoni, 2010).

The technique of word categories provides insight into associations between personality and language use, but the reported correlations are typically too low to reliably infer author personality from text. Nowson and Oberlander (2006) found that using n-grams (sequences of n items typically used to capture word collocations) resulted in more accurate predictions of personality and gender from online blogs than LIWC. Schwartz et al. (2013) showed that an open-vocabulary approach on a large corpus containing 700 million words, phrases, and topic instances collected from Facebook messages of 75,000 volunteers, provided insights and accuracies that could not be obtained with closed-vocabulary word-category analyses such as LIWC.

Apart from predictive ability, another consideration when training language-based predictive models is their susceptibility to deception. It may be undesirable if an introvert could be classified as an extravert by deliberately using words that are mainly used by extraverts, e.g. party and beach (Schwartz et al., 2013). An approach to mitigate such straightforward deception attempts is to focus on how someone writes, rather than what they write. This method is known as computational stylometry and involves feature types such as simple character n-grams, punctuation, token n-grams, semantic and syntactic class distributions and patterns, parse trees, complexity, and vocabulary richness measures, and even discourse features (Daelemans, 2013). Stylometric features have been used to predict personality traits from student essays (Luyckx and Daelemans, 2008), transcribed video blogs (Verhoeven and Daelemans, 2014), and twitter messages (Verhoeven et al., 2016).

More recently, deep learning techniques have enabled computers to efficiently learn semantic vector representation of words, sentences, and paragraphs from large corpora (Mikolov et al., 2013; Pennington et al., 2014, Le and Mikolov, 2014, and Devlin et al., 2018). By representing words, sentences, or paragraphs as (sequences of) dense N-dimensional vectors, significant performance gains have been reported in various natural language processing problems including sentiment classification, machine translation, and question-answer systems (Young et al., 2018).

Not surprisingly, Majumder et al. (2017) report that a neural network using these word-level vector embeddings outperforms traditional approaches (e.g. n-grams, closed-vocabulary, and open-vocabulary approaches) in terms of accuracy for Big Five personality traits. IBM personality insights, a commercial service to extract personality characteristics from text, is no longer using a LIWC-based model for predictions but is currently using a machine learning algorithm operating on word-level vector embeddings (IBM Personality Insights, 2019).

Pera's technology also exploits recent deep learning techniques to infer personality traits and competencies from natural language. Using a proprietary unsupervised learning technique on a large answer corpus, we learn dense N-dimensional vector embeddings that capture the stylistic as well as semantic characteristics of answers to open-ended questions. In turn, these vector embeddings in combination with supervised machine learning techniques enable accurate personality trait and competency prediction models to be learnt from relatively small training datasets.

References

- Argamon, S., Dhawle, S., Koppel, M., & Pennebaker, J. W. (2005, June). Lexical predictors of personality type. In Proceedings of the 2005 Joint Annual Meeting of the Interface and the Classification Society of North America (pp. 1-16).
- Barrick, M. R., & Mount, M. K. (1991). The big five personality dimensions and job performance: a meta-analysis. Personnel psychology, 44(1), 1-26.
- Cable, D. M., & Judge, T. A. (1996). Person–organization fit, job choice decisions, and organizational entry. Organizational behavior and human decision processes, 67(3), 294-311.
- Daelemans, W. (2013, March). Explanation in computational stylometry. In International Conference on Intelligent Text Processing and Computational Linguistics (pp. 451-462). Springer, Berlin, Heidelberg.
- Devlin, J., Chang, M. W., Lee, K., & Toutanova, K. (2018). Bert: Pre-training of deep bidirectional transformers for language understanding. arXiv preprint arXiv:1810.04805.
- Fast, L. A., & Funder, D. C. (2008). Personality as manifest in word use: correlations with self-report, acquaintance report, and behavior. Journal of personality and social psychology, 94(2), 334.
- Furnham, A. (1990). Faking personality questionnaires: Fabricating different profiles for different purposes. Current psychology, 9(1), 46-55.
- Gardner, W. L., Reithel, B. J., Cogliser, C. C., Walumbwa, F. O., & Foley, R. T. (2012). Matching personality and organizational culture: Effects of recruitment strategy and the Five-Factor Model on subjective person–organization fit. Management Communication Quarterly, 26(4), 585-622.
- Gill, A. J., Nowson, S., & Oberlander, J. (2009, March). What are they blogging about? Personality, topic and motivation in blogs. In Third International AAAI Conference on Weblogs and Social Media.
- Hirsh, J. B., & Peterson, J. B. (2009). Personality and language use in self-narratives. Journal of research in personality, 43(3), 524-527.
- Hurtz, G. M., & Donovan, J. J. (2000). Personality and job performance: The Big Five revisited. Journal of applied psychology, 85(6), 869.
- IBM Personality Insights (2019), The science behind the service. Retrieved from https://cloud.ibm.com/docs/services/personality-insights?topic=personality-insights-science.
- Kandola, B. (2009). The value of difference: Eliminating bias in organisations. BookBaby.
- Kristof-Brown, A., Barrick, M. R., & Franke, M. (2002). Applicant impression management: Dispositional influences and consequences for recruiter perceptions of fit and similarity. Journal of Management, 28(1), 27-46.
- Kristof-Brown, A. L., Zimmerman, R. D., & Johnson, E. C. (2005). Consequences of individual's fit at work: a meta-analysis of person-job, person-organization, person-group, and person-supervisor fit. Personnel psychology, 58(2), 281-342.
- Le, Q., & Mikolov, T. (2014, January). Distributed representations of sentences and documents. In International conference on machine learning (pp. 1188-1196).
- Lee, C. H., Kim, K., Seo, Y. S., & Chung, C. K. (2007). The relations between personality and language use. The Journal of general psychology, 134(4), 405-413.
- Luyckx, K., & Daelemans, W. (2008). Using syntactic features to predict author personality from text. Proceedings of Digital Humanities, 2008, 146-9.
- Majumder, N., Poria, S., Gelbukh, A., & Cambria, E. (2017). Deep learning-based document modeling for personality detection from text. IEEE Intelligent Systems, 32(2), 74-79.
- Martin, J. (2014). For Senior Leaders, Fit Matters More than Skill. Harvard Business Review, Retrieved from https://hbr.org/2014/01/for-senior-leaders-fit-matters-more-than-skill.
- Mehl, M. R., Gosling, S. D., & Pennebaker, J. W. (2006). Personality in its natural habitat: Manifestations and implicit folk theories of personality in daily life. Journal of personality and social psychology, 90(5), 862.
- Mikolov, T., Chen, K., Corrado, G., & Dean, J. (2013). Efficient estimation of word representations in vector space. arXiv preprint arXiv:1301.3781.
- Nowson, S., & Oberlander, J. (2006, March). The Identity of Bloggers: Openness and Gender in Personal Weblogs. In AAAI spring symposium: Computational approaches to analyzing weblogs (pp. 163-167).
- Pennington, J., Socher, R., & Manning, C. (2014). Glove: Global vectors for word representation. In Proceedings of the 2014 conference on empirical methods in natural language processing (EMNLP) (pp. 1532-1543).
- Wikipedia (2019), List of official languages by country and territory. Retrieved from https://en.wikipedia.org/wiki/List_of_official_languages_by_country_and_territory.
- O'Reilly III, C. A., Chatman, J., & Caldwell, D. F. (1991). People and organizational culture: A profile comparison approach to assessing person-organization fit. Academy of management journal, 34(3), 487-516.
- Pennebaker, J. W., & King, L. A. (1999). Linguistic styles: Language use as an individual difference. Journal of personality and social psychology, 77(6), 1296.
- Reuters (2018). Amazon scraps secret AI recruiting tool that showed bias against women. Retrieved from https://www.reuters.com/article/us-amazon-com-jobs-automation-insight/amazon-scraps-secret-ai-recruiting-tool-that-showed-bias-against-women-idUSKCN1MK08G.
- Rynes, S. L., & Gerhart, B. (1993). Recruiter perceptions of applicant fit: Implications for individual career preparation and job search behavior. Journal of Vocational behavior, 43(3), 310-327.
- Schmidt, F. L., & Hunter, J. E. (1998). The validity and utility of selection methods in personnel psychology: Practical and theoretical implications of 85 years of research findings. Psychological Bulletin, 124(2), 262-274.
- Sullivan J. (2017). Ouch, 50% Of New Hires Fail! 6 Ugly Numbers Revealing Recruiting's Dirty Little Secret. Retrieved from https://www.ere.net/ouch-50-of-new-hires-fail-6-ugly-numbers-revealing-recruitings-dirty-little-secret/.
- Tett, R. P., Jackson, D. N., & Rothstein, M. (1991). Personality measures as predictors of job performance: a meta‐analytic review. Personnel psychology, 44(4), 703-742.
- Vancouver, J. B., & Schmitt, N. W. (1991). An exploratory examination of person‐organization fit: Organizational goal congruence. Personnel psychology, 44(2), 333-352.
- Verhoeven, B., & Daelemans, W. (2014, November). Evaluating content-independent features for personality recognition. In Proceedings of the 2014 ACM Multi Media on Workshop on Computational Personality Recognition (pp. 7-10). ACM.
- Verhoeven, B., Daelemans, W., & Plank, B. (2016). Twisty: a multilingual twitter stylometry corpus for gender and personality profiling. In Proceedings of the 10th Annual Conference on Language Resources and Evaluation (LREC 2016)/ Calzolari, Nicoletta [edit.]; et al. (pp. 1-6).
- Yarkoni, T. (2010). Personality in 100,000 words: A large-scale analysis of personality and word use among bloggers. Journal of research in personality, 44(3), 363-373.
- Young, T., Hazarika, D., Poria, S., & Cambria, E. (2018). Recent trends in deep learning based natural language processing. ieee Computational intelligenCe magazine, 13(3), 55-75.