

ICCV 2021 Understanding Social Behavior in Dyadic and Small Group Interactions Challenge

Fact sheet: Automatic self-reported personality recognition Track

This is the fact sheet’s template for the ICCV 2021 Understanding Social Behavior in Dyadic and Small Group Interactions Challenge [1], “Automatic self-reported personality recognition Track”. Please fill out the following sections carefully in a scientific writing style. Then, send the compressed project (in .zip format), i.e., the generated PDF, .tex, .bib and any additional files to juliojj@gmail.com, and put in the Subject of the email “Fact Sheets: ICCV 2021 (DYAD) Challenge”, following the schedule and instructions provided in the Challenge webpage [1] “*Winning solutions (post-challenge), Fact Sheets*”.

I. TEAM DETAILS

- Team leader name: Francisca Pessanha and Gizem Sogancioglu
- Username on Codalab: `fpessanha, gizemsogancioglu`
- Team leader affiliation: Utrecht University
- Team leader email: `f.pessanha@uu.nl, g.sogancioglu@uu.nl`
- Name of other team members (and affiliation): Metehan Doyran, Heysem Kaya, Ronald Poppe, Albert Ali Salah, Almila Akdag Salah
- Team website URL (if any): <https://www.uu.nl/en/research/interaction/social-and-affective-computing/people>

II. CONTRIBUTION DETAILS

A. Title of the contribution

We propose an informed baseline to help disentangle the various contextual factors of influence in this type of case studies. For this purpose, we analysed the correlation between the given metadata and the self-assigned personality trait scores and developed a model based solely on this information. Further, we compared the performance of this informed baseline with models based on state-of-the-art visual, linguistic and audio features. For the present dataset, a model trained solely on simple metadata features (age, gender and number of sessions) proved to have superior or similar performance when compared with simple audio, linguistic or visual features based systems.

B. Representative image / workflow diagram of the method

The pipelines for the informed (metadata based) and combined (acoustic-linguistic + metadata) systems are shown in Figures 1 and 2, respectively.




Fig. 1. Personality prediction informed baseline.

C. Detailed method description

The regressor used for all steps of this section was a multitask Random Forest regressor. We defined the hyperparameters of each model based on a 3-fold participant-independent cross-validation. A model was trained for each task to evaluate what were the most representative interactions. The best performing model was chosen based on the performance in the development set. We used both training and development sets for training in the final evaluation phase. Both per task model and fused models (using task as a feature) experimented.

Firstly, we defined different models based on gender, age, number of sessions, background and mood dimensions. A simple model using only gender, age and number of sessions as features outperformed the development baseline in all dimensions motivating us to explore this matter further (Figure 1). For this purpose, we explored solutions based on linguistic, audio and video features and compared them with this informed baseline.

In the linguistic domain, we concatenated the transcriptions of the participant in the session and extracted state-of-the-art features, namely Sentence-Bert [2], TFIDF [3], LIWC [4] and VAD [5]. Further, we handcrafted features to represent the interaction between the two participants: turn rate, speech rate and speech duration rate. In the audio domain, we sampled the audio using turn-wise chunks and extracted functionals from geMAPS for these audio segments [6]. For visual, we obtained i3D features for every 2.5 seconds of the participant camera video [7]. The final prediction for both audio and visual models was given by the median of all the predictions for the session in question.

The best performing models in the development set were then fused as illustrated in Figure 2.

D. Challenge results

The best performance on the test set was achieved using the simple pipeline described in Figure 1.




Fig. 2. Personality prediction model using linguistic, visual and metadata features. We defined the best performing models for each personality trait based on their performance on the development set and applied late fusion to get a final prediction. The best performing models for both linguistic and audio features were obtained in the ghost task.

TABLE I

RESULTS FROM LEADERBOARD (TEST PHASE) OBTAINED BY THE PROPOSED APPROACH.

Rank position	O	C	E	A	N	MSE
3	0.752	0.687	0.917	0.671	1.098	0.825

E. Final remarks

In the development phase, the combined model described in Figure 2 achieved the best performance (Table II). However, these results were not replicable in the final evaluation phase, possibly due to differences in trait balance amongst the splits and the small number of samples for both development and test sets, leading to overfitting to the development set characteristics. Therefore, the results presented in Section II were obtained with a metadata trained model.

A correlation between personality, gender and age has been reported in previous studies. According to Marsh et al. [8] women have higher latent scores on all OCEAN traits, except for openness. Further, with age, individuals become happier, more self-content and self-centred and more conformed, which leads to less openness, neuroticism, extroversion and conscientiousness and more agreeability. Further, we hypothesise that, considering the time investment and social energy required for participation in a session, the number of sessions for each participant is correlated to their personality traits, namely openness. For this reason, using metadata for personality prediction is a relevant experiment.

Despite this information, the three metadata factors referred should not be sufficient for a strong personality prediction, particularly when compared with audio-visual-metadata informed systems like the challenge baseline. This fact may indicate that the dataset includes unwanted bias. Considering the small size of the dataset and the similarities in backgrounds of the participants, with a mainly highly educated population, we expect to find patterns in personality that won't generalize for a more diverse population. Further, "apparent personality" may deviate from the "real personality", particularly for traits low in visibility like emotional stability, which can then influence the performance of appearance-based models' performance like the baseline. Further, "apparent personality" may deviate from the "real personality", particularly for traits low in visibility like emotional stability, which can then influence the performance of appearance-

based models' performance like the baseline. In this domain, recent work on "apparent personality" prediction also points out the influence of inherent characteristics of the subject such as age, gender, attractiveness and ethnicity, on the way their personality is perceived [9]. The same is true for the way the participants interact with each other, expecting different interactions when the participant is paired with similar or dissimilar personality types [10]. So, we consider the dataset would benefit from third-person annotation to analyze the correlations between self-reported and perceived personality traits.

In sum, our goal with the present submission was to propose an informed baseline to assess the influence of metadata in the performance of models for multimodal tasks. We argue that understandable baseline models are fundamental to access the contributions and shortcomings of complex systems in the affective computing field and contributes to the early detection of dataset related bias.

III. ADDITIONAL METHOD DETAILS

Please, reply if your challenge entry considered (or not) the following strategies and provide a brief explanation. For the non-binary questions, you can mark multiple options.

- **Mark with an X the modalities you have exploited.** () Visual, () Acoustic, () Transcripts, (x) Metadata, () Landmark annotations, () Eye-gaze vectors.

- **In case you used metadata, mark with an X the types of metadata you have exploited.** (x) Age, (x) Gender, () Country of origin, () Max. level of education, () Pre-session mood, () Post-session mood, () Pre-session fatigue, () Post-session fatigue, () Relationship among interactants, () Task type, () Task order, () Task difficulty, () Language, (x) Other.

If "other", or if you have used just a subset of info for a given type of metadata (e.g., just a subset of mood values), please detail: Number of sessions

- **Mark with an X the tasks you used for training.** () Talk, () Lego, () Animals, () Ghost.

The model proposed uses solely age, gender and number of sessions as features so, no task information is needed.

TABLE II

RESULTS OBTAINED BY THE APPROACHES REPRESENTED IN FIGURE 1 AND FIGURE 2 ON THE DEVELOPMENT SET.

Method	O	C	E	A	N	MSE
Challenge baseline	0.6589	1.3541	1.1893	0.9651	1.1487	1.0632
Metadata model	0.6074	1.3285	1.2082	0.9554	1.1029	1.0405
Combined model	0.6074	1.3285	1.0864	0.8476	1.0955	0.9931

In the training phase, we explored approaches using all four tasks.

- **Mark with an X the tasks you used for evaluation.**
() Talk, () Lego, () Animals, () Ghost.

- **Did you use the provided validation set as part of your training set?** (x) Yes, () No

If yes, please detail:

The final model was trained in the combined validation and training set.

- **Did you use any fusion strategy of modalities?** () Yes, (x) No

If yes, please detail:

- **Did you use ensemble models?** () Yes, (x) No

If yes, please detail:

- **Did you follow a multi-task approach or trained each trait individually?** (x) Multi-task, () Trained each trait individually.

- **Did you use information from the other interlocutor (e.g., their visual info) to predict the personality of the target interlocutor?** () Yes, (x) No.

If yes, please detail:

- **Did you use pre-trained models?** () Yes, (x) No

If yes, please detail:

- **Did you use external data?** () Yes, (x) No

If yes, please detail:

- **Did you use any regularization strategies/terms?** () Yes, (x) No

If yes, please detail:

- **Did you use handcrafted features?** () Yes, (x) No

If yes, please detail:

- **Did you use any pose estimation method?** () Yes, (x) No

If yes, please detail:

- **Did you use any face / hand / body detection, alignment or segmentation strategy?** () Yes, (x) No

If yes, please detail:

- **At what level of granularity did your method perform personality inference?** () Frame-level,

() Audio/video chunk-level (i.e., short audio/video snippet), () Task-level, () Session-level, () Other.

If “other”, please detail. If you selected “chunk-level”, please comment on the chunk length and why you selected it: Participant-level

- **Did you use any aggregation method to compute a single personality prediction per participant?**

() Yes, (x) No

If yes, please detail:

- **Did you use any spatio-temporal feature extraction strategy?** () Yes, (x) No

If yes, please detail:

- **Did you perform any data augmentation?**

() Yes, (x) No

If yes, please detail:

- **Did you use any bias mitigation technique (e.g., rebalancing training data)?**

() Yes, (x) No

If yes, please detail:

IV. CODE REPOSITORY

Link to a code repository with complete and detailed instructions so that the results obtained on Codalab can be reproduced locally. This includes a list of requirements, pre-trained models, and so on. Note, training code with instructions is also required. This is recommended for all participants and mandatory for winners to claim their prize. **Organizers strongly encourage the use of docker to facilitate reproducibility.**

Code repository: https://github.com/gizemsogancioglu/FGM_Utrecht

REFERENCES

- [1] ChaLearn Looking at People. ICCV’2021 Understanding Social Behavior in Dyadic and Small Group Interactions Challenge. [Online]. Available: <https://chalearnlap.cvc.uab.cat/challenge/45/description/>
- [2] J. Devlin, M.-W. Chang, K. Lee, and K. Toutanova, “BERT: Pre-training of deep bidirectional transformers for language understanding,” *arXiv preprint arXiv:1810.04805*, 2018.
- [3] L.-P. Jing, H.-K. Huang, and H.-B. Shi, “Improved feature selection approach TFIDF in text mining,” in *Proc. International Conference on Machine Learning and Cybernetics*, vol. 2. IEEE, 2002, pp. 944–946.
- [4] J. W. Pennebaker, R. L. Boyd, K. Jordan, and K. Blackburn, “The development and psychometric properties of LIWC2015,” Tech. Rep., 2015.
- [5] S. M. Mohammad, “Obtaining reliable human ratings of valence, arousal, and dominance for 20,000 english words,” in *Proc. ACL*, 2018.

- [6] F. Eyben, K. R. Scherer, B. W. Schuller, J. Sundberg, E. André, C. Busso, L. Y. Devillers, J. Epps, P. Laukka, S. S. Narayanan *et al.*, “The geneva minimalistic acoustic parameter set (gemaps) for voice research and affective computing,” *IEEE transactions on affective computing*, vol. 7, no. 2, pp. 190–202, 2015.
- [7] J. Carreira and A. Zisserman, “Quo vadis, action recognition? a new model and the kinetics dataset,” 2018.
- [8] H. W. Marsh, B. Nagengast, and A. J. Morin, “Measurement invariance of big-five factors over the life span: Esem tests of gender, age, plasticity, maturity, and la dolce vita effects.” *Developmental psychology*, vol. 49, no. 6, p. 1194, 2013.
- [9] R. D. P. Principi, C. Palmero, J. C. Junior, and S. Escalera, “On the effect of observed subject biases in apparent personality analysis from audio-visual signals,” *IEEE Transactions on Affective Computing*, 2019.
- [10] R. Cuperman and W. Ickes, “Big five predictors of behavior and perceptions in initial dyadic interactions: Personality similarity helps extraverts and introverts, but hurts “disagreeables”.” *Journal of personality and social psychology*, vol. 97, no. 4, p. 667, 2009.