

To Trust, or Not to Trust? A Study of Human Bias in Automated Video Interview Assessments

Chee Wee (Ben) Leong^{1†}, Katrina Roohr^{2†}, Vikram Ramanarayanan^{1*}, Michelle P. Martin-Raugh^{2†},
Harrison Kell^{2†}, Rutuja Ubale^{1*}, Yao Qian^{1*}, Zydrune Mladineo^{1†}, Laura McCulla^{2†}
NLP, Speech, Dialogic & Multimodal Research¹, Academic To Career Research²
Princeton[†], San Francisco^{*}, Educational Testing Service, USA

Abstract

Supervised systems require human labels for training. But, are humans themselves always impartial during the annotation process? We examine this question in the context of automated assessment of human behavioral tasks. Specifically, we investigate whether human ratings themselves can be trusted at their face value when scoring video-based structured interviews, and whether such ratings can impact machine learning models that use them as training data. We present preliminary empirical evidence that indicates there are biases in such annotations, most of which are visual in nature.

1. Introduction and Related Work

Structured interviews standardize the questions and/or evaluation methods used during the interview process (i.e., each interview has the same questions in the same order [12]). As demonstrated by extensive meta-analytic evidence [9, 23], structured interviews consistently outperform unstructured interviews (i.e., an interview where questions and evaluations are non-standardized). They are particularly effective hiring tools that account for significant incremental validity in predicting job performance over and above other popular hiring methods such as personality and intelligence tests [4]. Despite their pervasiveness and effectiveness, interviews are susceptible to *human biases*, a source of measurement error. Biases occur when interviewers collect or evaluate non-job-related information about applicants [12], such as sex, age, race, or attractiveness. Research shows that human interviewers tend to be less likely to hire individuals that are older [14], of a different race than themselves [13], perceived as unattractive [8], or of a sex not stereotypically associated with a particular job (e.g., male nurse; [5]), among many other biasing factors. In this work, we differentiate human bias from *human subjectivity*, where interviewers make different hiring recommendations for the same interviewee due to their different

weightings of job-related strengths and weakness of the applicant. In spite of mounting evidence of human bias in industrial and organizational (I-O) psychology research, supervised machine learning approaches to automated scoring of video-based performance tasks (e.g. interview, presentations, public speaking, etc.) have largely focused on mitigating *human subjectivity* by creating comprehensive rubrics [24] to guide precise scoring of performance constructs [1], enforcing the calibration of human raters prior to scoring [1, 16], encouraging inter-rater discussions and reviews [15], enlisting multiple raters [17, 21], including both behavioral experts and layman as raters [15], and averaging ratings [15, 21, 2]. The work in [18] attempted to collect demographic information such as age, gender and ethnicity as part of the data collection survey but did not explore their impact on automated scoring of the interviews. While the authors in [15] claimed additional Mechanical Turk raters were used as raters to remove bias from the actual interviewers due to their interaction with the interviewees, no effort was expended to further remove nor qualify the type of biases being investigated. In this paper, we present a study that motivates future efforts toward modeling fairness in automated video interview assessments. We first provide an overview of the video interview dataset used in our experiments. Next, we explain the annotation scheme used for generating the bias metadata vector, which is used to (1) construct a standalone model, and, (2) augment a multi-modal model for structured interview performance prediction. Finally, we discuss some thoughts for future directions.

2. Dataset

Our video interview dataset is a corpus of monologic, structured video interviews collected online through Amazon Mechanical Turk from the authors in Chen et al. [2]. To our knowledge, this is the largest collection of structured interview responses simulating an actual hiring scenario (i.e. hiring for an entry-level office position). It com-

prises 260 human interviewees with a total recording time of 3784 minutes. Due to the dataset being collected “in the wild”, there are videos where faces in dim lighting cannot be detected, or unexpected clips in audios. These represent around 10% of the dataset and unfortunately have to be discarded since we failed to extract reliable multimodal features from them. Eventually, we received 1887 (interviewee, video response) datapoints after pre-processing. All responses were collected from participants across the United States, and differed in gender, race, age, experience, etc. Recording conditions also varied in terms of devices, lighting, backgrounds, etc. All videos from interviewees are collected indoors. The authors developed a 7-point Likert rating scale (1 = Strongly Disagree, 7 = Strongly Agree) to score interview performance using overall hiring recommendation guidelines proposed in [10]. During annotation, 5 human raters were asked to score each response using the Likert scale, and their averaged ratings are used as the ground-truth for the video. The same raters were used for scoring all the video interviewees. Subsequently, the authors used the median score of all ground-truths as a threshold to separate all video responses into HIGH (above-average) vs. LOW (below-average), and frame their experiments as a binary classification task using multimodal features to predict the outcome of each interviewee’s performance. On the quality of human annotations, the authors reported an Intraclass correlation (ICC) of 0.79 (using the two-way random average measure of consistency), and R_{mean} of 0.74, where R is the correlation coefficient of individual raters’ scores to the averaged scores. Note that, for each interviewee, 8 different prompts (questions) were attempted. The response length is limited to 2 min/prompt.

3. Coding Bias Metadata

Research in I-O psychology has identified several categories of human biases relevant for our work, which we have listed here with the possible labels in parenthesis: SETTING (full room visible, partial room visible, only wall), WALL (blank, almost blank, with many items), GENDER (male, female), APPEARANCE (very unattractive, unattractive, average, attractive, very attractive), RACE (White, African American, Asian, Hispanic, Other), AGE (18-25, 26-35, 36-45, 46-55, 56-65, 65+), WEIGHT (very thin, thin, average, overweight, obese), FACIAL_STIGMA (no, yes), ACCENT (American Typical, International, Domestic (e.g. Southern, Bostonian, etc.)). Two human raters trained in education and assessment research annotated each video interviewee independently on each of the categories using the appropriate label after going through an initial calibration. The two raters differed in race, accent, age when self-rated. Note that this bias metadata annotation effort is independent from the hiring recommendation annotation in [2]. Additionally, we hypothesize that physical

backgrounds of interviewees *may* induce rater bias. Hence, we introduce two categories coding for the visible environment (SETTING) and the physical state of the wall (WALL) behind the interviewee. Cohen’s Kappa (κ) between the 2 raters are reported for each category: SETTING (.62), WALL (.76), GENDER (.99), APPEARANCE (.49), RACE (.72), AGE (.69), WEIGHT (.70), FACIAL_STIGMA (.13), ACCENT (.32). Due to ethical concerns (e.g. judging whether a person is attractive or not), we elected to do the annotations privately instead of crowdsourcing them. Where disagreement between the two raters occurs, we adopted a consistent rule for rounding off to the integer coding for the nearest label for all ordinal categories (e.g. for APPEARANCE, if two labels coded are average (3) and attractive (4), then the average is 3.5, but the rounded label is attractive (4)). For non-ordinal categories (i.e. GENDER, RACE and ACCENT), we always arbitrate to rater 1’s label where there is a disagreement for the sake of consistency. However, such disagreement cases account for only $\sim 2\%$ of the total non-ordinal label pairings. After the bias metadata vector for an interviewee is coded, the same vector is used for all 8 interview prompts answered by the same interviewee for experiments. We apply all labels in the bias metadata as a single feature set to construct models for predicting interviewee performance in a binary classification task (i.e. below-average or above-average). Since each of the 8 prompts are different, and responses across the 8 prompts are scored independently, each tuple (interviewee, video response) can be treated as single datapoint for experimentation. Of the 254 unique interviewees in the dataset, 152 ($\sim 60\%$) have attained different classification scores across the 8 prompts, suggesting a potentially significant variance in the ability to handle different prompts even within an individual. We use stratified sampling in a 10-fold cross-validation applied at the (interviewee,prompt)-level to maintain distribution of the classes while sampling, and employ learners with proven effectiveness on relatively small-to-medium sized datasets using the scikit-learn toolkit [20].

4. Modeling structured interview through the exclusive use of bias metadata




Figure 1. Feature category importance weighting using Random-ForestClassifier (n=500), 10-fold cross-validated




Figure 2. Model interpretation using LIME applied to a prediction instance by RandomForestClassifier (n=500)

Prior to model building, each categorical label is converted into one-hot encodings fitted on the entire dataset of 1887 (interviewee, prompt) tuples. Consequently, each label is transformed into a numerical value to facilitate experiments across a range of standardized learners. A stratified 10-fold cross-validation predictor that respects the class distribution of the train fold (i.e. a chance predictor), as well as a majority vote baseline, are used as the baselines. The results of the 10-fold cross-validation is shown in Table 1. Additionally, Figure 1 shows the weight of each category (averaged over the 10-folds) used in the Random Forest model which scores the best performance ($F_1=0.765$) in our combination experiments. The category weighting scheme is based on permutation importance [19] instead of importance based on reduction in node impurity that is implemented in scikit-learn, which has reliability concerns [25]. Note that category *random* is a control with randomly generated values for ensuring validity in the feature importance computation: if a category has negative importance, removing it actually *improves* performance. Importance of the FACIAL_STIGMA category (0) is currently inconclusive, due to overly skewed distributions (i.e. rater 1 and 2 annotated yes on only 3.1% and 1.9% of the datapoints respectively). Otherwise, the most important label for modeling in each category are as follows: SETTING (*partial_room_visible*), WALL (*almost_blank*), GENDER (*female*), APPEARANCE (*unattractive*), RACE (*White*), AGE (26-35), WEIGHT (*overweight*), ACCENT (*Typical*).

Though encoded into numerical values for experimentation, the original labels (e.g. male, obese, unattractive, etc.) associated with each category are *human-interpretable*, hence lending themselves to an explanation on whether any automated scoring model is behaving reasonably as measured against established I-O psychology findings. We take advantage of this phenomenon to further validate our hypothesis through application of the Local Interpretable Model-Agnostic (LIME) [22] toolkit to our

dataset, by examining datapoints where our model predictions are confident. For a targeted datapoint and its prediction, LIME perturbs inputs to other neighboring datapoints around it to learn an interpretable and high-fidelity, local model to help explain the prediction made by the input model. For instance, Figure 2 shows our Random Forest prediction $p(class = "below - average") = 1$ for a specific (interviewee, prompt) datapoint, with the contribution of each categorical label accounting for each candidate class prediction. Here, we note the interviewee being labeled with *appearance=unattractive* adds a probability of 14% to an unfavorable class prediction outcome without regard for multimodal features extracted. Simultaneously, the *accent=typical_accent* (typical American accent) adds a probability of 17% to a favorable class prediction. This datapoint, and others we have examined, corroborate research findings in I-O psychology to a degree supported by the bias annotation agreement that biases *do* influence interview outcomes. To the best of our knowledge, this work is the first that provides empirical linkage between I-O psychology and ML/AI modeling for structured video interviews.

5. Multimodal Model Augmentation

We also experimented with multimodal model augmentation (i.e. an approach of augmenting a multimodal model with our bias metadata) to observe if we can further achieve model performance gains. For a fair comparison, we obtained the original feature set used in [2], evaluate on the same training/testing partitions, but opted to use a deep neural network (DNN) for modeling the task due to its proven effectiveness for modeling other similar constructs (e.g. engagement [6, 29] and emotions [7, 27]) of human behavior.

Given that we have only 1887 datapoints (1521/training, 366/testing), we train a deeper network but with dropouts to minimize overfitting. In our DNN model, a 2D convolution layer is applied to the time-series data i.e. video and audio to extract spatio-temporal features of interest within video/audio segments. Given the different sampling rates for video and audio feature extraction in [2], and their different feature dimensions, we use different kernel sizes and filters (i.e. audio: 30 filters, each with kernel and strides (10,2); video: 20 filters, each with kernel and strides (15,16)) in order to generate a somewhat balanced representation of the two time-series modalities at the input level before sending them deeper into the network for abstraction. GRUs are used as our recurrent units which are favored over LSTMs for faster convergence in our experiment, with a final attention layer mechanism applied similar to the one in [26]. Our DNN model, shown in Figure 3, is constructed and evaluated using Keras [3] with a TensorFlow backend, and its hyperparameters are tuned using talos [11]. After 30 epochs in the same training partition, we achieved a best-performing model of $F_1=0.70$ on the same test partition

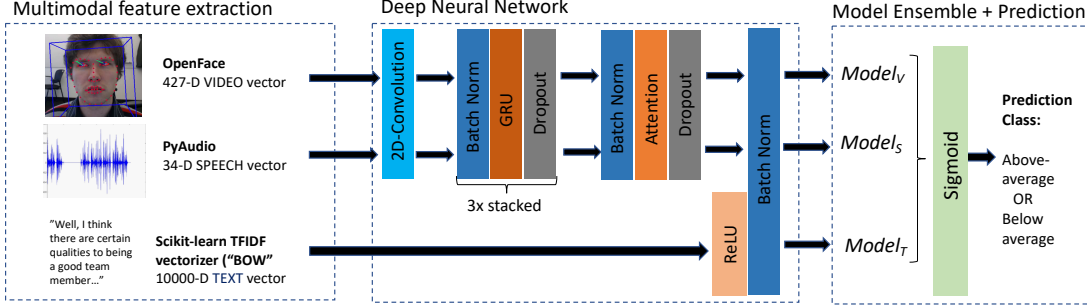


Figure 3. Multimodal DNN model for achieving the state-of-the-art in [2].

| | P | R | F ₁ |
|---------------------------------|--------------|-------------|----------------|
| INDIVIDUAL LEARNERS | | | |
| Logistic Regression | .581 | .567 | .572 |
| Nearest Neighbor | .720 | .707 | .713 |
| SVM (gamma=2, C=10) | .742 | .756 | .748 |
| Decision Tree | .692 | .732 | .711 |
| Random Forest (RF) (n=500) | .737 | .750 | .742 |
| Multilayer Perceptron (alpha=1) | .617 | .636 | .623 |
| Multimodal DNN | .606 | .915 | .727 |
| COMBINATIONS WITH DNN | | | |
| Multimodal DNN & RF (AND) | <u>.785*</u> | .701* | .739* |
| Multimodal DNN & RF (Stacking) | .746* | .787* | .765* |
| Multimodal DNN & SVM (AND) | .790* | .704* | .742* |
| Multimodal DNN & SVM (Stacking) | .755* | .773* | <u>.763*</u> |
| BASELINES | | | |
| Baseline (Stratified) | .512 | .552 | .531 |
| Baseline (Majority vote) | .510 | 1.00 | .675 |

Table 1. Mean 10-fold cross-validation of precision, recall and F₁ (true positive = above-average) of individual learners, selected combinations and baselines. All experiments are executed with the same random seed and random_state for replicability. Bold and underline indicates first and second ranked results per column, while * indicates statistical significance at $p < .001$ over the Multimodal DNN system, using the Wilcoxon signed-rank test.

using the same multimodal feature set, which is competitive to the model performance of $F_1=0.66$ achieved in [2] using SVM. With confirmation on the effectiveness of our DNN model, we retrained it with the same F_1 loss function and metric on each of the 10 train folds used in experiment 1. Next, we perform augmentation of the DNN model with the bias metadata using two approaches: (1) a simple, intuitive element-wise *AND* condition between predictions made by the DNN model and the other learners, and (2) a stacked generalization method [28] that uses the DNN prediction probabilities and combines it with the raw bias metadata vector before applying another learning algorithm to generate the final prediction. A justification for the latter approach is that the shorter but dense metadata bias vector may be masked by the much larger set of sparse, time-series modal features, hence combining them in the late fusion stage is more appropriate. Our results for the multimodal augmentation experiments are shown in Ta-

ble 1. While using the logical AND augmentation generates a high-precision classifier, the stacked generalization approach is better at achieving overall model performance measured by F_1 . **Inter-feature correlation:** As further investigation, we compute Cohen’s Kappa between each of the bias metadata features and the multimodal DNN prediction output. The results indicate only very slight agreement between the biases and the prediction output, with most of the Kappas centering at zero (GENDER has the highest Kappa at .04). **Feature importance:** We again compute feature importances (Random Forest, n=500), this time using all available features. Consequently, the multimodal DNN prediction exclusively accounts for the largest importance weighting (.09), which is more than the weighting assumed by APPEARANCE (.06) that ranks second. These findings indicative that the multimodal DNN model by itself still accounts for a substantial variance in performance of a joint model. Bias metadata features have a non-negligible, combined importance weighting (.27), which is concerning if they are indeed construct-irrelevant. However, it could also be the case that some of these metadata features proxy for a construct-relevant latent trait that is not captured in the DNN features.

6. Conclusion

Despite having little correlation with construct-relevant multimodal features, the fact that demographic characteristics and other biasing variables have a non-negligible impact on modeling human behavioral tasks is worrisome, and poses implications for tasks such as the automated assessment of structured video interviews used in high-stakes employment settings where the decisions made using scoring models need to be both fair and valid. Future work will further explore whether there is a causal relationship between these biases and human scores. If so, we will focus on debiasing techniques, possibly using avatars during human scoring that mirror an interviewees facial expressions without carrying over the visual biases, or modeling an individual human raters biases so that they can be statistically controlled for.

References

- [1] L. Chen, G. Feng, J. Joe, C. W. Leong, C. Kitchen, and C. M. Lee. Towards automated assessment of public speaking skills using multimodal cues. In *Proceedings of the 16th International Conference on Multimodal Interaction*, pages 200–203. ACM, 2014. 1
- [2] L. Chen, R. Zhao, C. W. Leong, B. Lehman, G. Feng, and M. E. Hoque. Automated video interview judgment on a large-sized corpus collected online. In *2017 Seventh International Conference on Affective Computing and Intelligent Interaction (ACII)*, pages 504–509. IEEE, 2017. 1, 2, 3, 4
- [3] F. Chollet et al. Keras. <https://keras.io>, 2015. 3
- [4] J. M. Cortina, N. B. Goldstein, S. C. Payne, H. K. Davison, and S. W. Gilliland. The incremental validity of interview scores over and above cognitive ability and conscientiousness scores. *Personnel Psychology*, 53(2):325–351, 2000. 1
- [5] H. K. Davison and M. J. Burke. Sex discrimination in simulated employment contexts: A meta-analytic investigation. *Journal of Vocational Behavior*, 56(2):225–248, 2000. 1
- [6] A. Dhall, A. Kaur, R. Goecke, and T. Gedeon. EmotiW 2018: Audio-video, student engagement and group-level affect prediction. In *Proceedings of the 2018 on International Conference on Multimodal Interaction*, pages 653–656. ACM, 2018. 3
- [7] D. Hazarika, S. Poria, R. Mihalcea, E. Cambria, and R. Zimmermann. Icon: Interactive conversational memory network for multimodal emotion detection. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 2594–2604, 2018. 3
- [8] M. Hosoda, E. F. Stone-Romero, and G. Coats. The effects of physical attractiveness on job-related outcomes: A meta-analysis of experimental studies. *Personnel psychology*, 56(2):431–462, 2003. 1
- [9] A. I. Huffcutt, J. M. Conway, P. L. Roth, and U.-C. Klehe. The impact of job complexity and study design on situational and behavior description interview validity. *International Journal of Selection and Assessment*, 12(3):262–273, 2004. 1
- [10] A. I. Huffcutt, J. M. Conway, P. L. Roth, and N. J. Stone. Identification and meta-analytic assessment of psychological constructs measured in employment interviews. *Journal of Applied Psychology*, 86(5):897, 2001. 2
- [11] M. Kotila. talos – hyperparameter optimization for keras, 2018. 3
- [12] J. Levashina, C. J. Hartwell, F. P. Morgeson, and M. A. Campion. The structured employment interview: Narrative and quantitative review of the research literature. *Personnel Psychology*, 67(1):241–293, 2014. 1
- [13] T.-R. Lin, G. H. Dobbins, and J.-L. Farh. A field study of race and age similarity effects on interview ratings in conventional and situational interviews. *Journal of Applied Psychology*, 77(3):363, 1992. 1
- [14] F. P. Morgeson, M. H. Reider, M. A. Campion, and R. A. Bull. Review of research on age discrimination in the employment interview. *Journal of Business and Psychology*, 22(3):223–232, 2008. 1
- [15] I. Naim, M. I. Tanveer, D. Gildea, and M. E. Hoque. Automated prediction and analysis of job interview performance: The role of what you say and how you say it. In *2015 11th IEEE International Conference and Workshops on Automatic Face and Gesture Recognition (FG)*, volume 1, pages 1–6. IEEE, 2015. 1
- [16] L. S. Nguyen, D. Fraundorfer, M. S. Mast, and D. Gatica-Perez. Hire me: Computational inference of hirability in employment interviews based on nonverbal behavior. *IEEE transactions on multimedia*, 16(4):1018–1031, 2014. 1
- [17] L. S. Nguyen and D. Gatica-Perez. I would hire you in a minute: Thin slices of nonverbal behavior in job interviews. In *Proceedings of the 2015 ACM on International Conference on Multimodal Interaction*, pages 51–58. ACM, 2015. 1
- [18] L. S. Nguyen and D. Gatica-Perez. Hirability in the wild: Analysis of online conversational video resumes. *IEEE Transactions on Multimedia*, 18(7):1422–1437, 2016. 1
- [19] T. Parr. Feature importances for scikit random forests, 2018. 3
- [20] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, et al. Scikit-learn: Machine learning in python. *Journal of machine learning research*, 12(Oct):2825–2830, 2011. 2
- [21] V. Ramanarayanan, C. W. Leong, L. Chen, G. Feng, and D. Suendermann-Oeft. Evaluating speech, face, emotion and body movement time-series features for automated multimodal presentation scoring. In *Proceedings of the 2015 ACM on International Conference on Multimodal Interaction*, pages 23–30. ACM, 2015. 1
- [22] M. T. Ribeiro, S. Singh, and C. Guestrin. Why should i trust you?: Explaining the predictions of any classifier. In *Proceedings of the 22nd ACM SIGKDD international conference on knowledge discovery and data mining*, pages 1135–1144. ACM, 2016. 3
- [23] F. L. Schmidt and R. D. Zimmerman. A counterintuitive hypothesis about employment interview validity and some supporting evidence. *Journal of Applied Psychology*, 89(3):553, 2004. 1
- [24] L. M. Schreiber, G. D. Paul, and L. R. Shibley. The development and test of the public speaking competence rubric. *Communication Education*, 61(3):205–233, 2012. 1
- [25] C. Strobl, A.-L. Boulesteix, A. Zeileis, and T. Hothorn. Bias in random forest variable importance measures: Illustrations, sources and a solution. *BMC bioinformatics*, 8(1):25, 2007. 3
- [26] R. Ubale, Y. Qian, and K. Evanini. Exploring end-to-end attention-based neural networks for native language identification. In *2018 IEEE Spoken Language Technology Workshop (SLT)*, pages 84–91. IEEE, 2018. 3
- [27] V. Vielzeuf, C. Kervadec, S. Pateux, A. Lechervy, and F. Jurie. An occam’s razor view on learning audiovisual emotion recognition with small training sets. In *Proceedings of the 2018 on International Conference on Multimodal Interaction*, pages 589–593. ACM, 2018. 3
- [28] D. H. Wolpert. Stacked generalization. *Neural networks*, 5(2):241–259, 1992. 4

- [29] J. Yang, K. Wang, X. Peng, and Y. Qiao. Deep recurrent multi-instance learning with spatio-temporal features for engagement intensity prediction. In *Proceedings of the 2018 on International Conference on Multimodal Interaction*, pages 594–598. ACM, 2018. [3](#)