



PersonaTAB: Predicting Personality Traits using Textual, Acoustic, and Behavioral Cues in Fully-Duplex Speech Dialogs

Sho Inoue^{1,3}, Shuai Wang^{2,1,3,*}, Haizhou Li^{1,3,4}

¹School of Data Science, The Chinese University of Hong Kong, Shenzhen, China
²School of Intelligence Science and Technology, Nanjing University, Suzhou, China
³Shenzhen Research Institute of Big Data, Shenzhen, China
⁴Department of ECE, National University of Singapore, Singapore

shoinoue@link.cuhk.edu.cn, shuaiwang@nju.edu.cn, haizhouli@cuhk.edu.cn

Abstract

Despite significant progress in neural spoken dialog systems, personality-aware conversation agents—capable of adapting behavior based on personalities—remain underexplored due to the absence of personality annotations in speech datasets. We propose a pipeline that preprocesses raw audio recordings to create a dialogue dataset annotated with timestamps, response types, and emotion/sentiment labels. We employ an automatic speech recognition (ASR) system to extract transcripts and timestamps, then generate conversation-level annotations. Leveraging these annotations, we design a system that employs large language models to predict conversational personality. Human evaluators were engaged to identify conversational characteristics and assign personality labels. Our analysis demonstrates that the proposed system achieves stronger alignment with human judgments compared to existing approaches.

Index Terms: personality prediction, dialog dataset, conversation agents, fully-duplex, large language models

1. Introduction

Conversational agents have been extensively studied and deployed in applications such as voice assistants and chatbots. These agents are generally into two models: the turn-taking (streaming) model and the fully-duplex (interaction) model [1]. In the turn-taking model, the agent responds only after the user completes speaking [2, 3, 4], whereas the fully-duplex model permits simultaneous speech [5, 6, 7], enabling richer interactions that include interjections and backchanneling. Notably, backchanneling can stimulate speaker engagement [8] and influence speech flow [9, 10]. We focus on the *fully-duplex setting* due to its closer resemblance to natural conversation.

Despite recent progress, personality-aware conversational agents remain underexplored. These agents may exhibit diverse behaviors and attitudes in identical scenarios due to personality traits, influencing both their responses [11, 12] and their use of interjections [13, 14, 15]. Current systems often fail to adapt behavior based on personality, instead relying on static or random personas. Conversation datasets with explicit personality labels are limited. To bridge this gap, we propose a method for predicting conversational personality using a fully-duplex speech dialog dataset.

Prior research has primarily employed text-based sources such as social media posts [16, 17, 18], essays [18, 19, 20, 21], and movie scripts [22] for personality prediction. Although [18] demonstrated that large language models (LLMs) can infer personality traits from text, they have not fully considered both chat history and conversational behaviors.

In this paper, we construct a dialog dataset with personality labels derived exclusively from speech data. We introduce a pipeline that preprocesses raw speech into a structured richly-labeled dialog format. Inspired by [18], we develop a prediction framework employing Large Language Models (LLMs) to assign personality labels to speakers. Our approach integrates various aspects of conversational behavior, including textual content (emotion and sentiment), acoustics (laughter), and additional characteristics (e.g. interjections and speaking dominance). This framework enables accurate, context-sensitive personality inference from speech-based dialogues. The contributions of this work are summarized as follows:

- We introduce a pipeline that preprocesses raw audio recordings to create a dialogue dataset annotated with timestamps, response types, and emotion/sentiment.
- We design a system to predict conversational personality from speech dialogues. This study is the first to integrate interjection behaviors, acoustics, and other conversational traits with textual data for personality prediction.
- We validate both the generated dataset and prediction model through human evaluations, offering valuable resources for developing conversational systems.

Supplementary materials are available on our project page¹.

2. Preliminaries

2.1. Personality Traits in Conversation

We adopt the Big Five Personality Model [23], a well-established framework characterizing individual cognition, affect, and behavior across five orthogonal dimensions. We detail each dimension below, presenting opposites in parentheses:

- *openness (closeness)*: intellectual, imaginative, independent-minded
- *conscientiousness (lack of direction)*: orderly, responsible, dependable
- *extraversion (introversion)*: talkative, assertive, energetic
- *agreeableness (antagonism)*: good-natured, cooperative, trustful
- *neuroticism (emotional stability)*: emotional instability, irritability, anxiety, self-doubt, depression

2.2. Terms in Fully-Duplex Setting

We define turn-taking events based on prior studies [24, 25, 26, 27]. These events encompass turn-taking behaviors, backchannels, and interjections. Fig. 1(1) illustrates the definitions:

¹**Project Page:** <https://github.com/shinshoji01/Personality-Prediction-for-Conversation-Agents>

* Corresponding author

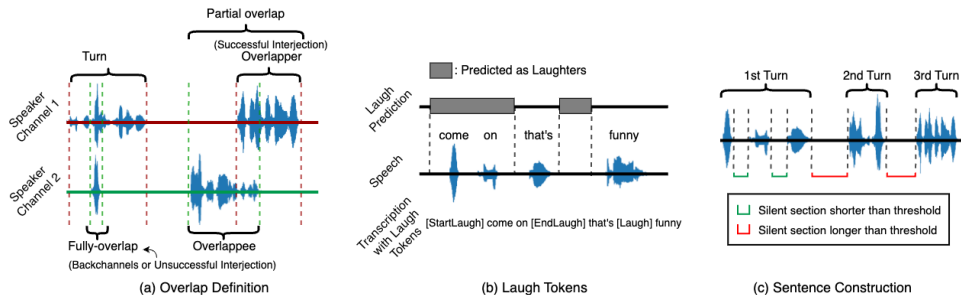


Figure 1: Visualization of examples of (a) Turns and Overlaps; (b) Laugh Token Integration; (c) Sentence Concatenation from Word-Level Time Stamps.

- **Turn:** A speaker’s contribution from the start of their speech until they stop or another speaker takes over.
- **Partial overlap:** An instance where the *overlapper*² stops speaking before the *overlapper*³ finishes speaking.
- **Fully overlap:** An overlap where the *overlapper*² continues speaking after the *overlapper*³ stops speaking.

When the overlapper interrupts the overlapped and successfully takes the turn, this is termed a “successful interjection”. Fully overlaps can result in either unsuccessful interjections or backchannels. These outcomes depend on the speaker’s intention to take the turn. If no such intention exists, the event is a backchannel. We categorize backchannels into two types:

- **Emotive backchannels:** Express the listener’s emotional valence, such as “wow” to convey surprise.
- **Cognitive backchannels:** Reflect the listener’s cognitive state or thought processes like “I see” to indicate comprehension.

Classifying backchannels requires contextual analysis. We also propose a pipeline for effective backchannel classification.

3. Methodology

3.1. Overall Pipeline

We designed a pipeline to construct a dialog dataset exclusively from two-channel speech-only datasets. The resulting dataset includes timestamps, laughter labels, response labels, emotions, sentiments for each response, and personality labels for each speaker. We began by preprocessing the speech data to construct dialog structures. Subsequently, we predicted personality traits via Large Language Models (LLMs). This prediction system integrates multiple features, including textual information (content, emotion, sentiment), acoustic features (laughter), and conversational traits (e.g., number of turns).

3.2. Dataset Preprocessing

We preprocess two-channel speech data to construct dialog datasets. Two-channel speech data comprises separate channels, with each channel containing speech only from one of two speakers. We detect and annotate overlapping speech segments, which naturally occur in conversational settings. The resulting dialog dataset includes timestamps, response labels (turns, backchannels, or interjections), and text-based annotations for emotion and sentiment for each response. The preprocessing procedure consists of six main steps.

²Overlapper: the speaker/response that gets overlapped.

³Overlapper: the speaker/response that overlaps the speech.

Step 1. Speech Transcription with Word-Level Timestamps

We transcribed the speech and obtained word-level timestamps using Whisper Turbo⁴ [28]. Subsequently, we removed silent intervals from each timestamp to refine the data.

Step 2. Laughter Event Annotation

Laughter is a key feature for determining personality traits. Therefore, we integrated laugh labels into the system. We obtained laughter timestamps using a laughter detector⁵ [29]. To enhance transcription, we added three laugh tokens based on their speech containment: [Laughter] for isolated vocal bursts, and [StartLaugh] and [EndLaugh] for speech-laughter co-occurrences. For example, in Fig. 1(2), the timestamp of the first laugh includes two words, “come” and “on”, resulting in “[StartLaugh] come on [EndLaugh]”. The second laugh contains no text and is labeled as “[Laughter]”.

Step 3. Response Boundary Construction

We constructed sentences by concatenating words using word-level timestamps. For a single speaker, we merged successive words when the inter-word gap was below a predefined threshold (Fig.1(c)). We set this threshold at 700ms, following [30], which reported language generation latencies exceeding 600ms. We defined response boundaries using silent intervals rather than contextual cues. Prior studies [31, 32] indicate that speakers often signal turn transitions with silence, and we consider turn-taking frequency a critical feature in personality detection.

Step 4. Overlap Identification

After constructing sentences with timestamps for both speakers, we detected overlaps. We identified two types of overlaps: “Partial overlap” and “Fully overlap” (refer to Sec. 2.2 for definitions). Overlap labels were excluded if the duration of overlap was less than 700 ms. We categorized responses with “Fully-overlap” labels as potential backchannels. Among responses labeled as “Partial overlap”, we classified the latter response (overlapper) as a “successful interjection”.

Step 5. Backchannel Classification via LLMs

We classified the possible backchannels into emotive backchannels, cognitive backchannels, or unsuccessful interjections using Large Language Models (LLMs). Backchannels vary in nature and depend heavily on context [33, 34, 35], and they are not always limited to a predefined set of words, as they can be longer or more complex [36].

To address this, we employed classifiers that consider chat history. Specifically, we used GPT-4o⁶ to analyze and classify possible backchannels effectively. We included the following

⁴Whisper Turbo: <https://github.com/openai/whisper>

⁵<https://github.com/jrgillick/laughter-detection>

⁶gpt-4o-2024-11-20: <https://openai.com/api/>

features in the prompt for classification: (1) an explanation of the backchannel classes—“emotive”, “cognitive”, and “not backchannel”; (2) the chat history, including both past and future responses; and (3) the target text with its corresponding position in the overlappee’s speech. The full version of the prompt is available on our project page¹. Responses labeled as “not backchannel” were classified as “unsuccessful interjection.”

Step 6. Emotion and Sentiment Classification of Responses
Emotion and sentiment are critical features for personality prediction. We obtained those labels using text-based classifiers: the emotion classifier⁷ and the sentiment classifier⁸.

3.3. Predicting Personality Traits

We predict speaker personalities from dialogue data generated in Sec.3.2. We employ an LLM (GPT-4o⁶) to estimate the Big Five Personality Traits—*openness*, *conscientiousness*, *extraversion*, *agreeableness*, and *neuroticism* (see Sec.2.1). We instruct the LLM to classify each speaker using five labels per personality: “highly aligned”, “aligned”, “neutral”, “opposed”, and “highly opposed”. We assign the “opposed” label when a speaker exhibits the inverse of a trait (e.g., closeness instead of openness). We represent each speaker with attributes across four categories: *Emotion*, *Sentiment*, *Basics*, and *Samples*.

We convert those speaker’s attributes into textual format and incorporate them into the LLM prompt. The full prompt comprises the task definition, explanations of the personality dimensions, and the speaker’s attributes across the four categories. The complete prompt is available on our project page¹. The following section details our integration of the four categories of speaker attributes into the prompt.

- **Emotion/Sentiment**

We employ emotion and sentiment labels derived from the dialog dataset. We calculate the percentage of each class (e.g., “Anger”) within each category (e.g., “Emotion”) for all responses associated with each speaker. We then incorporate these percentages into the LLM prompt.

- **Basics**

We analyze the speaker’s conversational attributes based on timestamps, laughter, backchannel, and interjection labels. Specifically, we incorporate the following attributes:

- *Number of Turns*: The number of turns, defined as responses separated by 700 ms, excluding backchannels and unsuccessful interjections.
- *Speaking Duration*: The average speech duration per turn.
- *Average Laughters*: The number of laughter occurrences per minute of speech.
- *Average Emotive/Cognitive Backchannels*: The number of emotive or cognitive backchannels per minute during another speaker’s speech.
- *Number of Interjections*: The total count of successful and unsuccessful interjections within 12 minutes.

Although these attributes are interpretable, we observed that LLMs struggle to assess whether these values are high or low. To address this, we classified these values into five relative groups. We computed the mean and interquartile range (IQR) across all speakers, and defined the relative groups as follows:

- *Normal* if $|\text{value}| < \text{mean} + 0.8 \times \text{IQR}$

⁷<https://huggingface.co/j-hartmann/emotion-english-distilroberta-base>

⁸<https://huggingface.co/cardiffnlp/twitter-roberta-base-sentiment-latest>

- *Many/Few* if $|\text{value}| > \text{mean} + 0.8 \times \text{IQR}$
- *Very Many/Few* if $|\text{value}| > \text{mean} + 1.2 \times \text{IQR}$

We then incorporate these five relative groups of the analyzed features into the prompt (see Fig. 2).

```

Basic Statistics:
Number of turns: Normal
Talking time per turn: Very Short
:
Frequency of interjections: Frequent

```

Figure 2: Example Prompt for “Basics”

- **Samples**

We include sample responses in the LLM prompt to analyze each speaker’s linguistic behavior. We randomly select 20 turns exceeding 2 seconds in duration.

4. Experiment Setups

4.1. Dataset

We used the Fisher dataset [37], a large-scale collection of telephonic speech recordings. Each sample contains a 12-minute phone conversation between two speakers, with audio captured on separate channels. For this study, we evaluated our pipeline on a subset of the Fisher dataset [37] (folder 000), consisting of 95 conversations and 190 speakers.

4.2. Subjective Evaluation Experiments

We conducted two MUSHRA tests to assess the alignment between human judgments and prediction outcomes. In the first experiment, participants predicted trends in conversational characteristics based on personality traits (e.g., between “Frequent Laughters” and “Openness”). In the second experiment, evaluators rated alignment for each personality dimension using a 90-second speech dialogue. We recruited 20 English-speaking participants, each receiving comprehensive explanations of the personality traits. A sample page is available on our page¹.

4.3. Baselines

We employed three baseline systems to predict personality traits from text responses. The first baseline, “LM”⁹ [16], was trained on Reddit comments and posts and outputs a probability distribution over five personality dimensions. The remaining baselines, BERT (“BERT”)¹⁰ [38] and MiniLM (“MiniLM”)¹¹ [39], predict personality labels in a multi-label setting using distinct architectures. We computed personality probabilities for all responses and averaged them per speaker to derive alignment scores for each personality label.

5. Experiments and Results

5.1. Personality Prediction

We conducted experiments to compare our system with baseline methods. We classified speaker personality using GPT-4o⁶,

⁹https://huggingface.co/KevSun/Personality_LM

¹⁰<https://huggingface.co/Minej/bert-base-personality>

¹¹<https://huggingface.co/microsoft/MiniLM-L12-H384-uncased>

which predicts the alignment between speaker attributes and each of the five personalities. Classification labels are “highly aligned” (100), “aligned” (50), “neutral” (0), “opposed” (-50), and “highly opposed” (-100). We averaged five responses to yield reliable alignment scores, with the alignment values enclosed in curly braces.

We employed two subjective metrics to assess the alignment between human evaluations and model predictions. The first metric quantifies how well model predictions conformed to the expected trends in conversation characteristics for each personality. Our subjective evaluation confirmed these trends, as shown in Table 1 (with 100 denoting “very aligned” and -100 denoting “very opposed”). Consequently, we identified several attributes strongly associated with each personality. For instance, extraversion—a talkative or energetic personality—aligned with frequent turn-taking (74), while opposed to frequent fearful responses (-59).

Table 1: *The expected trend between speaker attributes and personalities. Blue and red denote positive and negative correlations, respectively. Strong correlations are highlighted in bold. Column names indicate abbreviated personality labels.*

	Op.	Con.	Ext.	Agr.	Neu.
Emotion — Anger	-44	-55	-30	-70	64
Emotion — Disgust	-63	-51	-47	-77	57
Emotion — Fear	-70	-61	-59	-66	64
Emotion — Joy	50	41	65	61	-16
Emotion — Neutral	25	37	17	31	-16
Emotion — Sadness	-55	-47	-50	-52	58
Emotion — Surprise	39	35	71	42	47
Sentiment — Positive	50	47	62	63	-49
Sentiment — Neutral	31	42	20	17	-17
Sentiment — Negative	-42	-54	-33	-64	56
Basics — Taking Turns	57	54	74	47	-15
Basics — Average Time per Turn	62	15	72	3	-5
Basics — Laughters	17	12	66	51	-9
Basics — Emotive Backchannels	43	48	73	68	1
Basics — Cognitive Backchannels	51	61	68	77	-15
Basics — Interjections	1	-30	58	-39	6

This table is derived solely from a human survey

We quantified the alignment between speaker conversation attributes and predicted personality scores. For each personality label, we computed a weighted sum of correlation coefficients between personality traits and raw speaker attributes (prior to categorization), where higher values indicate stronger positive correlations. We derive the weights by normalizing the values in Table 1. Here, we multiplied coefficients corresponding to negative trends by -1 . Table 2 summarizes these scores. Our model outperforms baseline models in most metrics; however, we observed a lower score in “Conscientiousness”. Our analysis reveals that the correlation for the “Emotion” attribute deviates from expected trends, with values of -0.166 , 0.0176 , and 0.079 for “surprise”, “angry”, and “sadness”, respectively. We attribute these discrepancies to the model’s reliance on sample responses for determining responsibility, which emotional cues may not adequately capture.

Table 2: *Alignment of human judgments on conversation trait trends. Column names denote abbreviated personality labels.*

(†)	Open.	Cons.	Extr.	Agre.	Neur.	Avg.
Ours	0.167	0.061	0.174	0.210	0.315	0.186
BERT	0.053	-0.04	0.036	-0.015	-0.087	-0.011
LM	0.030	-0.011	-0.021	0.044	0.144	0.037
MiniLM	0.060	0.035	0.010	0.123	0.122	0.070

In our second evaluation, we computed the similarity between predicted personality scores and human-annotated labels obtained from a listening test. Table 3 reports the correlation values for each personality trait and the average cosine similarity across all samples. Our model consistently outperforms the other baselines, particularly when we interpret personality labels as vector dimensions (Cosine Similarity).

Table 3: *Similarities between human labels and prediction results. Column names denote abbreviated personality labels.*

	Correlation from Ground-Truth Labels (Corr.) (†)						Cosine
	Open.	Cons.	Extr.	Agre.	Neur.	Avg.	Similarity
Ours	0.169	0.198	0.285	0.144	0.117	0.183	0.503
BERT	0.095	0.044	0.194	0.058	0.089	0.096	-0.543
LM	0.023	-0.09	-0.018	0.030	-0.027	-0.016	-
MiniLM	-0.122	-0.064	0.014	-0.112	0.021	-0.053	0.148

5.2. Ablation Study

We conducted an ablation study on character prediction by varying speaker attributes to assess the impact of the attribute categories such as Emotion/Sentiment, Basics, and Samples. We collected five responses for each experimental condition. Table 4 shows that the full prompt consistently outperforms other conditions across all metrics, indicating improved alignment with human judgment and enhanced correlation with expected trends. In particular, compared with the “Samples”-only condition, the higher “Trend” score confirms that we effectively direct the prediction focus to the desired attributes.

Table 4: *An ablation study of speaker characteristics. “O” and “-” denote the inclusion and exclusion of features, respectively. Averaged scores in Table 2 and Table 3 are denoted as “Trend”, “Corr.”, and “Cosine”, respectively.*

Samples	Conditions				Scores		
	Basics	Emotion	Sentiment	Trend	Corr.	Cosine	
O	-	-	-	0.108	0.167	0.551	
-	O	-	-	0.051	0.009	0.524	
-	-	O	O	0.277	0.064	-0.203	
O	O	-	-	0.102	0.135	0.603	
O	-	O	O	0.173	0.170	0.407	
-	O	O	O	0.210	0.089	0.198	
O	O	O	O	0.186	0.183	0.503	

6. Conclusion

We propose a pipeline for constructing a dialogue dataset with extensive labels and for predicting personality labels from speech-only data. Our analysis shows that our prediction model aligns more closely with human judgments than existing models. Future work will investigate synthetic datasets with personality labels, motivated by evidences that LLMs can generate dialogue based on personality traits [40, 41], and will develop conversational agents conditioned on personality labels.

7. Acknowledgements

(1) Shenzhen Science and Technology Program (Shenzhen Key Laboratory, Grant No. ZDSYS20230626091302006), (2) Shenzhen Science and Technology Research Fund (Fundamental Research Key Project, Grant No. JCYJ20220818103001002), and (3) Program for Guangdong Introducing Innovative and Entrepreneurial Teams, Grant No. 2023ZT10X044. (4) Yangtze River Delta Science and Technology Innovation Community Joint Research Project (Grant No. 2024CSJGG01100).

8. References

- [1] S. Ji, Y. Chen, M. Fang, J. Zuo, J. Lu, H. Wang, Z. Jiang, L. Zhou, S. Liu, X. Cheng *et al.*, “Wavchat: A survey of spoken dialogue models,” *arXiv preprint arXiv:2411.13577*, 2024.
- [2] D. Zhang, S. Li, X. Zhang, J. Zhan, P. Wang, Y. Zhou, and X. Qiu, “Speechgpt: Empowering large language models with intrinsic cross-modal conversational abilities,” in *Conference on Empirical Methods in Natural Language Processing*, 2023.
- [3] E. Nachmani, A. Levkovitch, R. Hirsch, J. Salazar, C. Asawaroenchai, S. Mariooryad, E. Rivlin, R. Skerry-Ryan, and M. T. Ramanovich, “Spoken question answering and speech continuation using spectrogram-powered LLM,” in *The Twelfth International Conference on Learning Representations*, 2024.
- [4] Q. Fang, S. Guo, Y. Zhou, Z. Ma, S. Zhang, and Y. Feng, “Llama-omni: Seamless speech interaction with large language models,” 2024.
- [5] T. A. Nguyen, E. Kharitonov, J. Copet, Y. Adi, W.-N. Hsu, A. Elkahky, P. Tomasello, R. Algayres, B. Sagot, A. Mohamed, and E. Dupoux, “Generative spoken dialogue language modeling,” 2022.
- [6] A. D’efosse, L. Mazar’e, M. Orsini, A. Royer, P. P’erez, H. J’egou, E. Grave, and N. Zeghidour, “Moshi: a speech-text foundation model for real-time dialogue,” *ArXiv*, vol. abs/2410.00037, 2024.
- [7] P. Wang, S. Lu, Y. Tang, S. Yan, W. Xia, and Y. Xiong, “A full-duplex speech dialogue scheme based on large language model,” in *The Thirty-eighth Annual Conference on Neural Information Processing Systems*, 2024.
- [8] V. Pipek, “On backchannels in english conversation,” 2007.
- [9] J. Bavelas, L. Coates, and T. Johnson, “Listeners as co-narrators,” *Journal of personality and social psychology*, vol. 79, pp. 941–52, 12 2000.
- [10] J. Tolins and J. Fox Tree, “Overhearers use addressee backchannels in dialog comprehension,” *Cognitive science*, vol. 40, 09 2015.
- [11] M. Soto-Sanfiel, L. Aymerich-Franch, E. Romero, M. Soto, M. Soto, and M. Sanfiel, “Personality in interaction: how the big five relate to the reception of interactive narratives,” *COMMUNICATION & SOCIETY / COMUNICACIÓN Y SOCIEDAD*, 06 2014.
- [12] F. Mairesse, M. A. Walker, M. R. Mehl, and R. K. Moore, “Using linguistic cues for the automatic recognition of personality in conversation and text,” *Journal of Artificial Intelligence Research*, vol. 30, pp. 457–500, 2007.
- [13] Y.-H. Park, W. Liermann, Y.-S. Choi, S. Kim, J.-U. Bang, S. Yun, and K. Lee, “Backchannel prediction, based on who, when and what,” 09 2024, pp. 3570–3574.
- [14] H. Li, Y. Cui, and Z. Wang, “Backchannel responses and enjoyment of the conversation: The more does not necessarily mean the better,” *International Journal of Psychological Studies*, vol. 2, 05 2010.
- [15] S. K. MAYNARD, “On back-channel behavior in japanese and english casual conversation,” *Gengokenkyu*, vol. 1987, no. 91, pp. 129–130, 1987.
- [16] R. Wang and K. Sun, “Continuous output personality detection models via mixed strategy training,” 06 2024.
- [17] H. Christian, D. Suhartono, A. Chowanda, and K. Zamli, “Text based personality prediction from multiple social media data sources using pre-trained language model and model averaging,” *Journal of Big Data*, vol. 8, 05 2021.
- [18] Y. Ji, W. Wu, H. Zheng, Y. Hu, X. Chen, and L. He, “Is chatgpt a good personality recognizer? a preliminary study,” 07 2023.
- [19] M. Ramezani, M.-R. Feizi-Derakhshi, M. A. Balafar, M. Asgari-Chenaghlu, A. R. F. Derakhshi, N. Nikzad-Khasmakhi, M. Ranjbar-Khadivi, Z. Jahanbakhsh-Nagadeh, E. Zafarani-Moattar, and T. Rahkar-Farshi, “Automatic personality prediction: an enhanced method using ensemble modeling,” *Neural Computing and Applications*, vol. 34, pp. 18 369 – 18 389, 2020.
- [20] Y. Mehta, S. Fatehi, A. Kazameini, C. Stachl, E. Cambria, and S. Eetemadi, “Bottom-up and top-down: Predicting personality with psycholinguistic and language model features,” in *2020 IEEE International Conference on Data Mining (ICDM)*, 2020, pp. 1184–1189.
- [21] H. Lin, C. Wang, and Q. Hao, “A novel personality detection method based on high-dimensional psycholinguistic features and improved distributed gray wolf optimizer for feature selection,” *Information Processing & Management*, vol. 60, no. 2, p. 103217, 2023.
- [22] Q. Liu, “Modeling dyadic conversations for personality inference,” 09 2020.
- [23] J. M. Digman, “Personality structure: Emergence of the five-factor model,” *Annual Review of Psychology*, vol. 41, no. 1, pp. 417–440, 1990.
- [24] Anonymous, “Talking turns: Benchmarking audio foundation models on turn-taking dynamics,” in *Submitted to The Thirteenth International Conference on Learning Representations*, 2024, under review.
- [25] S. Duncan, “Some signals and rules for taking speaking turns in conversations,” *Journal of Personality and Social Psychology*, vol. 23, no. 2, pp. 283–292, 1972.
- [26] F. K. Ameka, “Interjections,” in *Encyclopedia of Language & Linguistics*, K. Brown, Ed. Elsevier, 2006, pp. 743–746.
- [27] G. Skantze, “Turn-taking in conversational systems and human-robot interaction: A review,” *Computer Speech & Language*, vol. 67, p. 101178, 2021.
- [28] A. Radford, J. W. Kim, T. Xu, G. Brockman, C. McLeavey, and I. Sutskever, “Robust speech recognition via large-scale weak supervision,” in *International conference on machine learning*. PMLR, 2023, pp. 28 492–28 518.
- [29] J. Gillick, W. Deng, K. Ryokai, and D. Bamman, “Robust laughter detection in noisy environments,” 08 2021, pp. 2481–2485.
- [30] S. C. Levinson and F. Torreira, “Timing in turn-taking and its implications for processing models of language,” *Frontiers in Psychology*, vol. 6, p. 731, 2015.
- [31] T. Stivers, N. J. Enfield, P. Brown, C. Englert, M. Hayashi, T. Heinemann, G. Hoymann, F. Rossano, J. P. D. Ruiter, K.-E. Yoon, S. C. Levinson, P. Kay, and K. Y., “Universals and cultural variation in turn-taking in conversation,” *Proceedings of the National Academy of Sciences*, vol. 106, pp. 10 587 – 10 592, 2009.
- [32] E. Schegloff, “Sequence organization in interaction: A primer in conversation analysis,” *Sequence Organization in Interaction: A Primer in Conversation Analysis I*, vol. 1, pp. 1–300, 01 2007.
- [33] —, *Discourse as an interactional achievement: Some uses of ‘uh huh’ and other things that come between sentences*, 01 1982, pp. 71–93.
- [34] R. Gardner, *When listeners talk: Response tokens and listener stance*. John Benjamins Publishing Company, 2001.
- [35] D. Fujimoto, “Listener responses in interaction: A case for abandoning the term, backchannel,” *Journal of Osaka Jogakuin 2year College*, vol. 37, 01 2007.
- [36] D. Mereu, F. Cangemi, and M. Grice, “Backchannels are not always very short utterances. the case of italian multi-unit backchannels,” *Journal of Pragmatics*, 2024.
- [37] C. Cieri, D. Miller, and K. Walker, “Fisher english training speech parts 1 and 2,” 2004, accessed: 2025-01-09.
- [38] J. Devlin, M. Chang, K. Lee, and K. Toutanova, “BERT: pre-training of deep bidirectional transformers for language understanding,” *CoRR*, vol. abs/1810.04805, 2018.
- [39] W. Wang, F. Wei, L. Dong, H. Bao, N. Yang, and M. Zhou, “Minilm: Deep self-attention distillation for task-agnostic compression of pre-trained transformers,” 2020.
- [40] W. Wanqi and T. Sakai, “Response generation based on the big five personality traits,” 2020.
- [41] E. Chen, “Generate labeled training data using prompt programming and gpt-3. an example of big five personality classification,” 03 2023.