

Affective-NLI: Towards Accurate and Interpretable Personality Recognition in Conversation

Zhiyuan Wen

Department of Computing
The Hong Kong Polytechnic University
Hong Kong, China
cszwen@comp.polyu.edu.hk

Jiannong Cao

Department of Computing
The Hong Kong Polytechnic University
Hong Kong, China
jiannong.cao@polyu.edu.hk

Yu Yang

Department of Computing
The Hong Kong Polytechnic University
Hong Kong, China
cs-yu.yang@polyu.edu.hk

Ruosong Yang

Department of Computing
The Hong Kong Polytechnic University
Hong Kong, China
rsong.yang@polyu.edu.hk

Shuaiqi Liu

Department of Computing
The Hong Kong Polytechnic University
Hong Kong, China
shuaiqi.liu@connect.polyu.hk

Abstract—Personality Recognition in Conversation (PRC) aims to identify the personality traits of speakers through textual dialogue content. It is essential for providing personalized services in various applications of Human-Computer Interaction (HCI), such as AI-based mental therapy and companion robots for the elderly. Most recent studies analyze the dialog content for personality classification yet overlook two major concerns that hinder their performance. First, crucial implicit factors contained in conversation, such as emotions that reflect the speakers' personalities are ignored. Second, only focusing on the input dialog content disregards the semantic understanding of personality itself, which reduces the interpretability of the results. In this paper, we propose Affective Natural Language Inference (Affective-NLI) for accurate and interpretable PRC. To utilize affectivity within dialog content for accurate personality recognition, we fine-tuned a pre-trained language model specifically for emotion recognition in conversations, facilitating real-time affective annotations for utterances. For interpretability of recognition results, we formulate personality recognition as an NLI problem by determining whether the textual description of personality labels is entailed by the dialog content. Extensive experiments on two daily conversation datasets suggest that Affective-NLI significantly outperforms (by 6%-7%) state-of-the-art approaches. Additionally, our Flow experiment demonstrates that Affective-NLI can accurately recognize the speaker's personality in the early stages of conversations by surpassing state-of-the-art methods with 22%-34%¹.

Index Terms—human-computer interaction, personality recognition

I. INTRODUCTION

The advancement of large language models (e.g., ChatGPT) has led to the emergence of numerous applications within Human-Computer Interaction (HCI), including AI-based mental therapy [1], educational robots [2], and companion robots designed for the elderly [3]. In these application scenarios, fast and accurately recognizing the user's personality from the dialog content is crucial, as it enables a better understanding of the user's personalized needs and requirements, ultimately

¹Our source code and data is at <https://github.com/preke/Affective-NLI>.

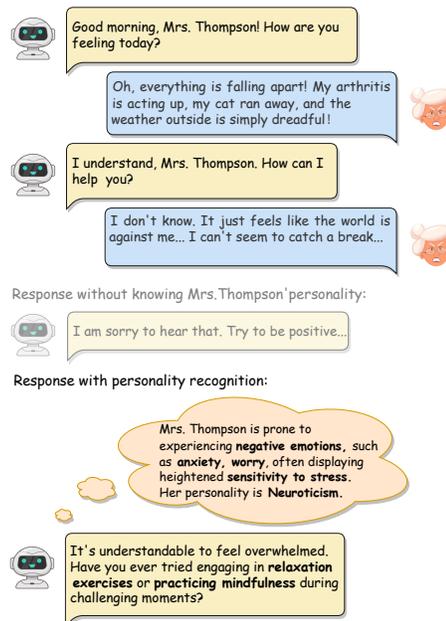


Fig. 1. When the robot lacks knowledge about Mrs. Thompson's personality, its comfort is generic and straightforward. However, if the robot deduces that Mrs. Thompson exhibits *Neuroticism* based on the dialog content, it can tailor its suggestions to offer more personalized and comforting advice.

leading to the provision of high-quality services [4]–[6], as shown in Figure 1.

Although important, addressing personality recognition in conversation (PRC) also faces a significant challenge. PRC strives to classify the specified speaker's personality into predefined personality classes (e.g., the big five personalities [7] shown in Table I) by analyzing the dialog content. However, personalities are relatively permanent traits that describe the consistency and individuality of a person's behaviors [8]. Accurate personality assessments typically involve obtaining

TABLE I
THE BIG FIVE PERSONALITY TRAITS (OCEAN) AND DESCRIPTION.

Traits	Descriptions
Openness	Openminded, imaginative, and sensitive.
Conscientiousness	Scrupulous, well-organized.
Extraversion	Tend to experience positive emotions.
Agreeableness	Trusting, sympathetic, and cooperative.
Neuroticism	Tend to experience psychological distress.

detailed personal information or long-term historical data, such as self-report essays, questionnaires, or longitudinal records from social media, which are always unavailable in daily conversation scenarios.

Recent studies solve the PRC problem utilizing the advanced semantic comprehension abilities of deep neural networks to analyze the dialog content. Specifically, researchers either design dialog encoding models [9] or fine-tune pre-trained language models [10]–[12] for PRC. However, these methods primarily concentrate on the explicit dialog content yet overlook two major concerns that hinder their recognition performance. First, semantic information contained in short utterances within conversations is inherently limited. These methods ignore crucial implicit factors contained in conversation, such as emotions and attitudes that reflect the speakers’ personalities, so that hinders their recognition accuracies. Second, only focusing on the input dialog content disregards the semantic understanding of personality class labels, which loses sight of additional discriminative cues and reduces the interpretability of the results.

To fill in the research gap, we propose Affective-Natural Language Inference (Affective-NLI), aiming for accurate and interpretable personality recognition in conversation with affective annotations (*i.e.*, discrete emotion category labels) of dialog content and the natural language descriptions of personality labels. Psychology findings suggest that there are strong correlations between personalities and affective expressions in conversations [13], [14]. However, accurately obtaining real-time affective annotations of dialog content for personality recognition is not a trivial task. We fine-tuned a pre-trained language model specifically for emotion recognition in conversations, facilitating real-time affective annotations for the utterances. Besides, the semantics behind personality labels contain descriptions of individual behavioral characteristics and tendencies in expressing affectivities, offering additional grounds for personality recognition. To utilize them for our problem, we adopt text descriptions of personality labels from psychology findings [15], [16] and formulate personality recognition as a natural language inference (NLI) problem. Specifically, the objective is to determine whether the personality description of the speaker (as the hypothesis) is correct or not, given the dialog content (as the premise). We solve this NLI problem by constructing a natural language prompt, including the premise and the hypothesis above, to fine-tune the pre-trained language models.

Compared to existing approaches, Affective-NLI leverages auxiliary affective information to enhance the limited semantics in dialog content. Then, Affective-NLI correlates the language and behavior of speakers within conversations with descriptions of personality traits, yielding more interpretable recognition results. Moreover, Affective-NLI mainly manipulates the input dialog content, making it compatible with large pre-trained language models in various architectures.

We conducted extensive experiments on two daily conversation datasets annotated with personality labels. In the standard PRC setting, where the full dialog input is provided, Affective-NLI demonstrated a significant improvement (6%-7%) in the accuracy of personality recognition compared to state-of-the-art methods. Additionally, we designed a **Flow** experiment to assess whether Affective-NLI can instantly recognize the speaker’s personality during the early stages of a conversation. The results showed that with just one or two utterances from the speaker under analysis, Affective-NLI achieved an accuracy of around 0.5-0.6 in personality recognition, surpassing state-of-the-art methods by 22%-34%. These findings suggest that Affective-NLI can be effectively employed for personality recognition in real-world human-computer interaction (HCI) applications. To demonstrate the interpretability of Affective-NLI, we conduct a case study to illustrate how our methods obtain the personality recognition results in a conversation example.

To summarize, our contributions are listed as follows:

- We introduce Affective-Natural Language Inference (Affective-NLI) to facilitate personality recognition in conversation with affective annotations of dialog content and the text descriptions of personality labels.
- Affective-NLI is capable of providing interpretable recognition results by correlating the language and behavior of speakers within conversations with descriptions of personality traits.
- Extensive experiments conducted on two conversational datasets verify the effectiveness of Affective-NLI. Additionally, Affective-NLI has been validated for early-stage personality recognition during conversations.

II. RELATED WORKS

In this section, we review existing research related to personality recognition in conversation.

Accurately recognizing users’ personalities in conversation is crucial for various applications, like building user trust in conversational agents [17], providing considerate elderly companions [4], or delivering personalized customer service [5]. So, this task has long been a critical research issue and explored by many researchers.

Research in early stages focus on finding statistical features for recognizing personality. The statistical word usage patterns and social behavior habits are highly correlated to the personality traits [18]. Informed by these findings, [19] tracked 96 participants using the Electronically Activated Recorder (EAR) to examine the expression of personality in daily conversations. Besides, [20] firstly presented several non-linear

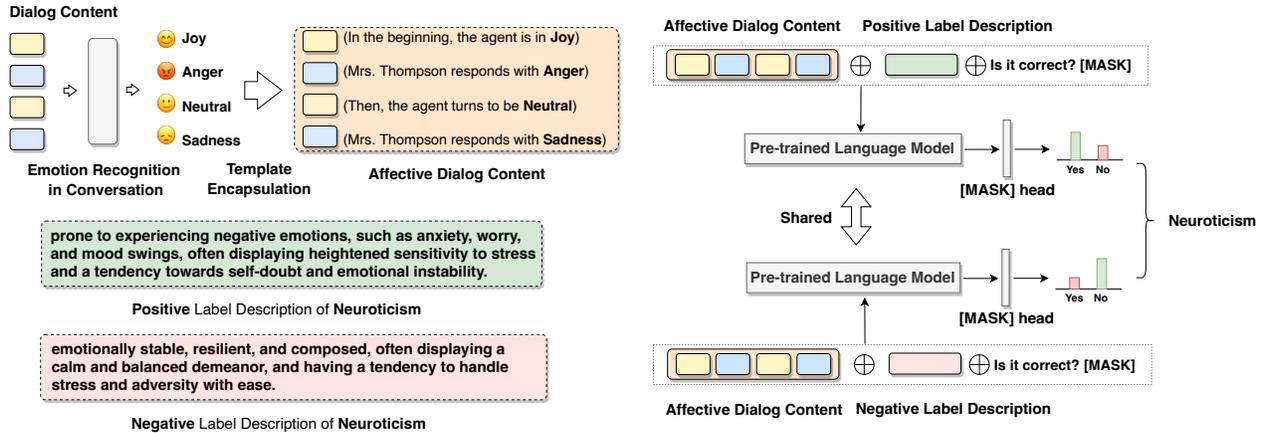


Fig. 2. An overview of Affective-NLI with the first four utterances in Figure 1. The upper-left part illustrates the affective dialog content construction. The lower-left part shows the positive and negative label descriptions of *Neuroticism*. The right part shows (1) how we construct NLI prompting samples with the affective dialog content and both positive and negative personality label descriptions and (2) how we combine both NLI results for the final personality recognition. Noting that the two different inputs to the same pre-trained language model.

statistical models to rank utterances on the Big Five personality traits with linguistic features. However, although the shallow features are efficient in providing statistical differences for personality recognition, they fail when the personalities are identified by deep understanding the dialog content.

With the development of deep learning, neural network models are widely applied to recognizing personality through understanding the semantics of the conversation. [9] proposed a classification model based on capsule neural networks to extract meaningful hidden patterns from conversations and use them to assess the personality of individuals. Subsequently, [10] used the pre-trained language models to encode the dialog content in different forms for personality classification. In [11], the authors inputted the full conversation text into BERT [21] and five independent SeNets [22] as the feature fusion layer of local personality features of different conversations. Some researchers also investigate prompt-tuning pre-trained language models (PLMs) for personality recognition. DesPrompt [12] utilizes a personality-descriptive prompt to encapsulate the input content, hence reformulates personality recognition as a cloze-style word-filling task to align with the pre-training Masked Language Modeling (MLM) task in PLMs.

However, in conversation scenarios, the utterances used for personality recognition are short sentences with limited quantity. Only focusing on semantic information in utterances fails to provide sufficient information to analyze personality, which hinders the recognition accuracy. Besides, concentrating on the input dialog content neglects the semantic comprehension of personality itself, thereby overlooking valuable discriminative cues and diminishing the interpretability of the results.

III. AFFECTIVE-NLI

In this section, we introduce the proposed Affective-NLI (as shown in Figure 2). We begin by presenting the formulation of the PRC problem in Section III-A. Subsequently, we delve into the construction of affective dialog content by Affective-

NLI in Section III-B, followed by the introduction of the personality label descriptions utilized in Section III-C. Lastly, we introduce how we construct NLI samples for fine-tuning and inference in pre-trained language models in Section III-D.

A. Problem Formulation

We study the Personality Recognition in Conversation (PRC) problem as stated below:

Given the conversation content $X = \{u_1, u_2, \dots, u_m\}$ with m text utterances u_i among a speaker s and other speakers, the objective is to recognize the personality traits of s , denoted as y . y is represented as a $5-d$ binary vector $[y_a, y_c, y_e, y_n, y_o]$ indicating the binary values of each trait referring to the big-five personality traits theory (shown in Table I, shortened as **AGR**, **CON**, **EXT**, **NEU**, and **OPN**, respectively). The big-five trait theory offers a categorical classification of personality derived from both the trait theory and the lexical hypothesis in psychology. It is widely used in analysis for social media content [23]–[26] and conversations [20].

We propose Affective-Natural Language Inference (Affective-NLI) to formulate the PRC problem as the NLI problem between the affective dialog content (*premise*) and the descriptions of personality labels (*hypothesis*). Natural Language Inference (NLI) determines whether a hypothesis is true (entailment), false (contradiction), or undetermined (neutral) given a premise. In the scenario of PRC, Affective-NLI determines whether the *hypothesis* of the personality label description for the speaker s is correct or not, given the affective dialog content as the *premise*.

Formally, to recognize each specific personality trait $y_p \in [y_a, y_c, y_e, y_o, y_n]$ for the speaker s given input dialog content, Affective-NLI maximizes the sum of the two probabilities:

$$P_{pos} + P_{neg} = P(y_p | X_{af}; H_{pos}(p)) + P(1 - y_p | X_{af}; H_{neg}(p)) \quad (1)$$

where the $X_{af} = \{(u_1, a_1), (u_2, a_2), \dots, (u_m, a_m)\}$ is the affective dialog content obtained by attaching an affective annotation a_i to each utterance u_i . $H_{pos}(p)$ and $H_{neg}(p)$ are the natural language descriptions of people with and without the specific personality trait p , respectively.

Note that $H_{pos}(p)$ and $H_{neg}(p)$ are opposite descriptions to a same personality trait p . The principle of maximizing both P_{pos} and P_{neg} is to affirm the personality description that aligns with the speaker s and simultaneously negate the description that contradicts s 's personality.

B. Affective Dialog Content Construction

We first introduce how we construct the affective dialog content X_{af} from the original dialog content X .

It is been pointed out that there are strong relations between the *Extraversion* and *Conscientiousness* traits and the positive affects such as *Joy* and *Surprise*; and between *Neuroticism* and *disagreeableness* and various negative affects such as *Anger*, *Sadness*, or *Fear* [27]. By observing the dialogues in daily conversation, we also found that the utterance-level emotional interaction among the speakers (e.g., in what emotion the speaker s talks to others or how others respond to s) is essential in analyzing the personality traits of s .

Although affectivity is essential for personality recognition, in conversation scenarios, especially in real-time interaction between users and conversational agents, it is unrealistic to always have accurate affective annotations for utterances. Therefore, we fine-tuned a pre-trained language model M_{ERC} especially for **Emotion Recognition in Conversations** to obtain the emotion label e_i of each $u_i \in X$ in an offline manner:

$$e_1, e_2, \dots, e_m = M_{ERC}(u_1, u_2, \dots, u_m) \quad (2)$$

where each e_i is one of the six basic emotions [28] (i.e., *Anger*, *Disgust*, *Fear*, *Joy*, *Sadness*, *Surprise*) and *Neutral*. This emotion categories are widely employed in studies on personality analysis [29], [30].

Although existing research [31], [32] has validated the affective processing capabilities of language models for text, the interaction among speakers in conversational scenarios provides additional contextual cues for emotion recognition. Inspired by previous work on modeling dialogue flows [33]–[35], we incorporate the speaker's identity into the utterance sequence during the encoding process to achieve more accurate emotion recognition, where the speaker's identity refers to binary indicators showing whether the current utterance is from the speaker s .

After we get each emotion label e_i , we wrap it into a handcraft template to form the affective description a_i for the corresponding utterance u_i . The detailed process is depicted in Algorithm 1, where the input consists of the dialog content accompanied by an emotion label sequence, and the output yields the dialog content enveloped by affective descriptions. We tried various templates and found that different template contents had little impact on the final effect, so we adopted the current illustrated template content.

Algorithm 1: Affective Dialog Content Construction

Data: Dialog Content $X = \{u_1, \dots, u_m\}$, Speaker s , Other speakers s_{other} , Emotion labels $\{e_1, \dots, e_m\}$
Result: Affective Dialog Content X_{af}

```

1  $X_{af} \leftarrow \text{""}$ ;
2 for  $i \leftarrow 1$  to  $m$  do
3   if  $i = 1$  then
4     // the first utterance
5     if  $u_i$  is from  $s$  then
6        $a_i \leftarrow \text{"(s is initially } e_i \text{)"};$ 
7     else
8        $a_i \leftarrow \text{"(At the beginning, } s_{other} \text{ is } e_i \text{)"};$ 
9     end
10  else
11    if  $u_i$  is from  $s$  then
12       $a_i \leftarrow \text{"(s responds with } e_i \text{)"};$ 
13    else
14       $a_i \leftarrow \text{"(Then, } s_{other} \text{ turns to be } e_i \text{)"};$ 
15    end
16   $X_{af}$  append with  $(u_i, a_i)$ 
17 end

```

C. Personality Description Construction

The class labels of the big five personality traits (such as *Conscientiousness* and *Neuroticism*) are not commonly used in everyday language. So, language models trained on general corpora may struggle to understand these terms, making personality identification challenging in conversational contexts.

The lexical hypothesis of personality traits states that (1) the most distinctive, significant, and widespread phenotypic attributes tend to become encoded as words in the conceptual reservoir of language [36]. Therefore, we adopt and summarize descriptions of the big five personality traits in both positive and negative classes (H_{pos} , H_{neg}) from psychology findings [15], [16], as shown in Figure II. These descriptions contain rich semantic meanings of personality traits in daily used languages that outline affective and behavioral tendencies of people with different personality traits. In Affective-NLI, these descriptions establish a correlation between the language and behavior of speakers during conversations and the personality traits, resulting in more understandable recognition results.

One possible concern is that these label descriptions may not be comprehensive or subjective in their wording. However, considering that these descriptions will be encoded in pre-trained models, benefiting from the powerful semantic understanding capabilities of language models, the semantics of each word within the descriptions, including its synonyms and similar expressions, will also be captured and encoded, to provide the auxiliary basis for personality recognition.

D. NLI Training and Inference

To maximize the probability in Formula 1, we construct NLI samples for fine-tuning the pre-trained language model M in a prompt tuning manner.

We first wrap the original dialog content X into the following two NLI input samples when recognizing each single personality trait p :

TABLE II
POSITIVE AND NEGATIVE LABEL DESCRIPTIONS OF THE BIG FIVE
PERSONALITY TRAITS

Traits	Classes	Descriptions
AGR	H_{pos}	friendly, cooperative, empathetic, and compassionate, often prioritizing harmonious relationships and the well-being of others.
	H_{neg}	confrontational, uncooperative, lacking empathy, and often prioritizing their own needs and desires over the well-being of others.
CON	H_{pos}	organized, responsible, diligent, detail-oriented, and committed to achieving their goals with a strong sense of duty and self-discipline.
	H_{neg}	disorganized, careless, impulsive, lacking discipline, and often displaying a disregard for responsibilities and commitments.
EXT	H_{pos}	outgoing, sociable, energetic, and thriving in social interactions, often seeking stimulation and enjoying the company of others.
	H_{neg}	introverted, reserved, quiet, and often preferring solitude or smaller social settings, conserving energy and finding fulfillment in introspection and reflection.
NEU	H_{pos}	prone to experiencing negative emotions, such as anxiety, worry, and mood swings, often displaying heightened sensitivity to stress and a tendency towards self-doubt and emotional instability.
	H_{neg}	emotionally stable, resilient, and composed, often displaying a calm and balanced demeanor, and having a tendency to handle stress and adversity with ease.
OPN	H_{pos}	has curiosity, open-mindedness, creativity, tolerance, emotional expressiveness, and willingness to embrace new experiences and ideas.
	H_{neg}	closed-minded, resistant to change, lacking curiosity, intolerant of differences, emotionally guarded, and hesitant to explore new ideas or experiences.

$$\begin{aligned} \mathcal{T}_{pos}(X, p) &= X_{af}; H_{pos}(p); \text{Is it correct? [MASK]}. \\ \mathcal{T}_{neg}(X, p) &= X_{af}; H_{neg}(p); \text{Is it correct? [MASK]}. \end{aligned} \quad (3)$$

The model M is fine-tuned to fill in the [MASK] with “yes” and “no”.

During training, we construct the NLI samples as follows:

$$\begin{aligned} \{\mathcal{T}_{pos}(X, p), yes\}, \{\mathcal{T}_{neg}(X, p), no\}; \text{ if } y_p = 1 \\ \{\mathcal{T}_{pos}(X, p), no\}, \{\mathcal{T}_{neg}(X, p), yes\}; \text{ if } y_p = 0 \end{aligned} \quad (4)$$

where “yes” and “no” are the ground truth for NLI samples. According to the NLI samples, we denote the four probabilities as follows:

$$\begin{aligned} P_{pos}(yes) &= P([\text{MASK}] = yes | \mathcal{M}_\theta(\mathcal{T}_{pos}(X, p))) \\ P_{pos}(no) &= P([\text{MASK}] = no | \mathcal{M}_\theta(\mathcal{T}_{pos}(X, p))) \\ P_{neg}(yes) &= P([\text{MASK}] = yes | \mathcal{M}_\theta(\mathcal{T}_{neg}(X, p))) \\ P_{neg}(no) &= P([\text{MASK}] = no | \mathcal{M}_\theta(\mathcal{T}_{neg}(X, p))) \end{aligned} \quad (5)$$

The training objective is fine-tuning the parameters θ in \mathcal{M} to maximum the sum of the probabilities

$$\begin{aligned} \arg \max_{\theta} (P_{pos}(yes) + P_{neg}(no)) * y_p \\ + (P_{neg}(yes) + P_{pos}(no)) * (1 - y_p) \end{aligned} \quad (6)$$

In inference, the results are determined by the comparison of sums of the probabilities:

$$\arg \max [P_{neg}(yes) + P_{pos}(no), P_{pos}(yes) + P_{neg}(no)] \quad (7)$$

IV. EXPERIMENT SETTINGS

A. Datasets

Recording daily conversations for analysis, especially multi-person conversations, has potential privacy concerns. According to the book *Television Dialogue: The sitcom Friends vs natural conversation* [37], television conversations and natural conversations are basically the same in terms of linguistic features. Therefore, we use two datasets analyzing the scripts of TV Series: FriendsPersona constructed by [10] and the Chinese Personalized and Emotional Dialogue (CPED) from [11] to evaluate Affective-NLI.

FriendsPersona is a dialog script dataset with personality annotations in 711 different dialogues, including 8,157 utterances. In each dialogue, The personality of one analyzed speaker is represented as 5-d binary vectors for the big-five traits. Following the problem settings in similar works [9], [10] of personality analysis in conversation, we conduct binary classification tasks over the five personality traits respectively to evaluate our method.

CPED is a large-scale Chinese personalized and emotional dialogue dataset, which contains around 12K dialogues of 392 speakers from 40 popular Chinese TV shows. To be consistent with the affective prompts and the personality descriptions in constructing input samples, we translate the dialogues in CPED into English with Google Translation² and manual verification. The detailed statistics are shown in Table III.

B. Methods in Comparison

Affective-NLI mostly manipulates the input content, making it compatible with various large pre-trained language models (PLMs). Therefore, we implement it based on the three popular open-source PLMs in different architectures (*i.e.*, encoder-only: RoBERTa [38], encoder-decoder: T5 [39], and decoder-only: Llama [40]). We also compare Affective-NLI with two state-of-the-art methods [10], [12] in fine-tuning and prompt-tuning PLMs for personality recognition in conversation. Details of the methods are introduced as below.

²<https://translate.google.com/>

TABLE III
STATISTICS OF THE TWO DATASETS.

Statistics	FriendsPersona	CPED
# Dialogues	711	11,835
# Uttrs per dialogue	11.8	11.2
# Unique Uttrs	8,157	109,455
Uttr Length	16.3	28.9
Label Distribution (Positive:Negative)	AGR(.43:.57) CON(.46:.54) EXT(.44:.56) OPN(.35:.65) NEU(.47:.53)	AGR(.58:.42) CON(.67:.33) EXT(.65:.35) OPN(.50:.50) NEU(.59:.41)

Affective-NLI (RoBERTa): Here, we use the pre-trained Roberta-base model for Affective-NLI. RoBERTa [38] is a famous pre-trained language model whose pre-training task is Masked-Language-Modeling (filling the masked tokens in sentences). It’s natural to adopt it to conduct the mask-filling task in Affective-NLI.

Affective-NLI (T5): Here, we use the pre-trained T5-base model as the backbone model of Affective-NLI. T5 [39] is a unified pre-trained language model to convert all text-based language problems into a sequence-2-sequence format. In affective-NLI, we use the token generated by the t5-decoder to fill in the [MASK] in affective-NLI.

Affective-NLI (Llama): Llama [40] is one of the most popular large generative language models, comparable with ChatGPT³ in some tasks⁴. We adopt a parameter-efficient fine-tuning approach (*i.e.* P-tuning [41]) to only fine-tune a small number of extra parameters ($\sim 0.6\%$ of the whole model) while freezing most parameters of the pre-trained Llama-7b with around 7 billion parameters. Specific to Affective-NLI, we use the last token generated by Llama to fill in the [MASK].

FT-RoBERTa: FT-RoBERTa [10] uses RoBERTa-base as the dialog encoder with three different kinds of input for personality recognition. FT-RoBERTa(S) only uses the utterances from the analyzed speaker s as the input; FT-RoBERTa (S+C) concatenates utterances from the analyzed speaker and other utterances as context together as the input, while FT-RoBERTa (F) inputs all the utterances within the whole dialog flow in their natural order for classification. To better illustrate the three inputs: if a dialogue flow is $[U_1, C_1, U_2, U_3, C_2, \dots]$, U_i means the speaker’s utterance and C_i means the utterance from other speakers in the context. The input of FT-RoBERTa (S) is $[U_1, U_2, U_3, \dots]$, the FT-RoBERTa (S+C) input is $[U_1, U_2, U_3, \dots, C_1, C_2, \dots]$, while the FT-RoBERTa(F) input is $[U_1, C_1, U_2, U_3, C_2, \dots]$ with the original sentence order.

DesPrompt: DesPrompt [12] generates personality-descriptive prompts and fine-tunes pre-trained language models (*i.e.*, here we use RoBERTa-base) for efficient personality recognition.

³<https://openai.com/blog/chatgpt>

⁴We didn’t implement Affective-NLI on ChatGPT as it is not open-sourced.

It reformulates personality recognition as a cloze-style word-filling task and aligns with the pre-training Masked Language Modeling (MLM) task in most PLMs, hence enhancing the utilization of the pre-trained parameters.

C. Evaluation Tasks and Metrics

We split both datasets into train, validation, and test sets with a portion around 8:1:1. Models are trained on the train and validation sets and the results on test sets are reported.

To comprehensively evaluate the performance of Affective-NLI and other models, we design two personality recognition tasks in different settings specific to the conversation scenario: **Overall** and **Flow**. The **Overall** is the regular task that takes the full length of dialog content for personality recognition. The **Flow** examines the personality recognition performance at first 25%, 50%, 75%, and the whole dialog flow. This is to show whether Affective-NLI can be applied to real conversation services to instantly recognize the personality of the user in the early stages of conversations.

For both tasks, we use the binary classification accuracy to evaluate the personality recognition performance on each trait and the averaged result.

D. Implementation Details

During implementation, we pad all the utterances with [PAD] to the MAX_LEN of 256. Each dialog flow is padded to a Dialogue MAX_LEN of 20 according to the dataset statistics. The dialog flows are fed into the models in batches with a size of 32. As for the pre-trained Emotion Recognition in Conversation model M_{ERC} , we classify 1000 randomly selected utterances in both datasets (in balanced emotion labels) by M_{ERC} . The accuracies are 0.792 for FriendsPersona and 0.815 for CPED, respectively, which is considerably accurate for the classification of seven classes. We use the Adam [42] as the optimization algorithm in training. The learning rate for each model is selected to refer to the best performance on the validation sets.

V. RESULTS ANALYSIS

Based on the aforementioned settings, we perform comprehensive experiments and conduct both quantitative and qualitative analyses of the obtained results. The quantitative analysis aims to address the following three Research Questions (**RQs**), while the qualitative analysis utilizes a case study to showcase how Affective-NLI can provide interpretable personality recognition results.

- **RQ 1:** Can Affective-NLI outperform state-of-the-art approaches on personality recognition in conversation?
- **RQ 2:** How much do different modules in Affective-NLI influence the performance in personality recognition?
- **RQ 3:** How early can Affective-NLI recognize the personality in a dialog flow?

TABLE IV
PERSONALITY RECOGNITION ACCURACIES ON FRIENDSPERSONA AND CPED.

Dataset	Method	AGR	CON	EXT	OPN	NEU	Avg
FriendsPersona	FT-RoBERTa (S)	0.579±0.056	0.564±0.062	0.567±0.065	0.621±0.042	0.525±0.051	0.571
	FT-RoBERTa (S+C)	0.579±0.033	0.568±0.065	0.568±0.065	0.621±0.042	0.531±0.049	0.573
	FT-RoBERTa (F)	0.579±0.056	0.565±0.065	0.568±0.065	0.621±0.042	0.529±0.048	0.572
	DesPrompt	0.521±0.076	0.535±0.048	0.549±0.045	0.649±0.015	0.495±0.039	0.550
	Affective-NLI (RoBERTa)	0.597±0.041	0.569±0.051	0.597±0.036	0.639±0.044	0.583±0.069	0.597
	Affective-NLI (T5)	0.625±0.034	0.592±0.039	0.614±0.056	0.655±0.076	0.581±0.073	0.613
	Affective-NLI (Llama)	0.595±0.021	0.566±0.004	0.612±0.011	0.645±0.030	0.568±0.008	0.597
CPED	FT-RoBERTa (S)	0.602±0.012	0.654±0.011	0.660±0.017	0.529±0.012	0.595±0.014	0.608
	FT-RoBERTa (S+C)	0.599±0.016	0.656±0.012	0.659±0.017	0.529±0.014	0.600±0.016	0.608
	FT-RoBERTa (F)	0.600±0.010	0.653±0.008	0.662±0.013	0.532±0.020	0.600±0.011	0.610
	DesPrompt	0.611±0.003	0.660±0.032	0.671±0.022	0.516±0.027	0.608±0.020	0.613
	Affective-NLI (RoBERTa)	0.637±0.007	0.662±0.021	0.676±0.016	0.534±0.014	0.612±0.041	0.624
	Affective-NLI (T5)	0.677±0.014	0.656±0.013	0.689±0.027	0.569±0.018	0.647±0.018	0.648
	Affective-NLI (Llama)	0.645±0.033	0.668±0.004	0.677±0.009	0.542±0.021	0.605±0.043	0.627



Fig. 3. Personality recognition accuracy comparison among the sub-models in the ablation study.

A. Can Affective-NLI outperform state-of-the-art approaches on personality recognition in conversation (PRC)?

To answer this research question, we analyze the PRC results of the methods introduced in Section IV-B on FriendsPersona and CPED are shown in Table IV. These results are obtained from the **Overall** experiment setting.

On both datasets, Affective NLI achieves the best accuracies (including single personality trait recognition and overall performances) compared with other methods, which validates its effectiveness in personality recognition in conversations. Specifically, the best Affective-NLI method (*i.e.*, Affective-NLI (T5)) outperforms the best baseline methods on both datasets by around 6-7%.

Due to different distributions, the models' performance in recognizing the same personality trait varies across the two datasets. However, we can still observe that Affective-NLI exhibits the highest improvement over other methods in terms of AGR and NEU. We hypothesize that this may be attributed to a greater disparity in affective information expressed by positive and negative speakers in conversations pertaining to these two personality types.

Within the three implementations, Affective-NLI (T5) achieved the highest accuracy in recognizing seven out of the ten personality traits. This resulted in Affective-NLI (T5) having the best average results on both datasets. This suggests that compared to other types of pre-trained language models, the encoder-decoder architecture of the T5 model performs better in this form of NLI task. In addition, Affective-NLI (RoBERTa) and Affective-NLI (Llama) performed similarly, indicating that the parameter size of the pre-trained models does not play a decisive role in the results of the current PRC task. Therefore, Affective-NLI can be implemented on slightly smaller pre-trained language models without sacrificing accuracy, which validates the usability of Affective-NLI.

Although our work has achieved state-of-the-art performance compared to most existing models, we are aware that the current performance still has some distance from real-world applications. Addressing this issue more effectively will require further efforts from interdisciplinary work in the future.

B. How much do different modules in Affective-NLI influence the performance in personality recognition?

To figure out the answer, we design an ablation study to show the results of sub-models without different modules in Affective-NLI, respectively. We selected the Affective-NLI (T5) model with the best overall performance and designed the following sub-models:

w.o. affective: we remove the affective annotations, replacing X_{af} with the original X . However, we retained the personality description in the input. This is to evaluate the effect of affective information on personality recognition.

w.o. personality: we keep the affective annotations and the NLI task format in the input but replace the personality label descriptions as directly the personality label name. This is to investigate the influence of semantic information contained in the personality descriptions on the PRC task.

TABLE V
DIALOGUE SAMPLES FOR PERSONALITY RECOGNITION.

Dialog Content	Probability of Personality				
	AGR	CON	EXT	OPN	NEU
Agent: Good morning, Mrs. Thompson! How are you feeling today? (In the beginning, the agent is in Joy)			-		
Mrs. Thompson: Oh, everything is falling apart! My arthritis is acting up, my cat ran away, and the weather outside is simply dreadful! (Mrs. Thompson responds with Anger)	37%	42%	45%	36%	54%
Agent: I understand, Mrs. Thompson. How can I help you? (Then, the agent turns to Neutral)			-		
Mrs. Thompson: I don't know. It just feels like the world is against me... I can't seem to catch a break... (Mrs. Thompson responds with Sadness)	21%	41%	47%	28%	68%
Agent: It's understandable to feel overwhelmed. Have you ever tried engaging in relaxation exercises or practicing mindfulness during challenging moments? ...			-		

only Pos: in this sub-model, we maintained the input of Affective-NLI but only used positive personality descriptions during training and inference. This is to examine whether combining two different descriptions of the same personality enhances performance.

We conduct the personality recognition in conversation on FriendsPersona and CPED with all the sub-models, their accuracies are shown in Figure 3. These results are also obtained from the **Overall** experiment setting.

In general, removing any module in Affective-NLI results in a decrease in model performance, which validates our model design. Specifically, when removing the personality description (**w.o. personality**), the performance of Affective-NLI drops significantly on both datasets, highlighting the importance of enabling the model to understand the semantics of personality for personality recognition. Besides, removing the affective annotations (**w.o. affective**), also weakens the model to a considerable extent, particularly evident in the FriendsPersona dataset. By referring to Table III, it can be observed that utterance length in FriendsPersona is apparently shorter. This suggests that providing affective information can be more helpful for personality recognition when there is less conversation content available. When using only the positive descriptions of each personality trait in Affective-NLI (**only Pos**), the model also shows a certain degree of performance decline, with a more pronounced decrease in FriendsPersona. Therefore, supervising the model to learn semantic information from different categories of personality traits can provide auxiliary benefits.

C. How early can Affective-NLI recognize the personality in a dialog flow?

To evaluate whether Affective-NLI can be applied to conversation scenarios for instant personality recognition, we design an experiment **Flow** by feeding the first 25%, 50%, 75%, and 100% of each dialog flow in Test sets to the personality recognition models. According to the different lengths of

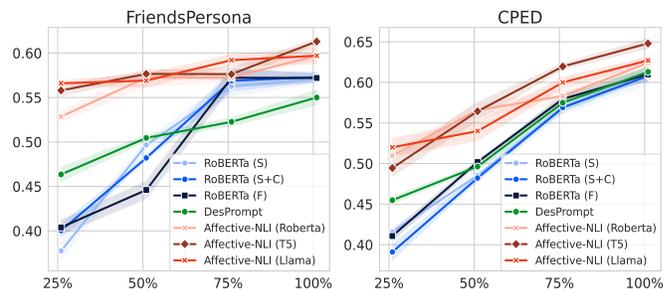


Fig. 4. The personality recognition accuracy comparison in the **Flow** experiment. The transparent error bars indicate the lower and upper bounds of the current accuracies.

dialog flows, we calculate the average numbers of utterances from the analyzed speaker s are 0.52, 1.48, 2.63, and 4.09, respectively. So, we make an additional rule that each dialog flow fed into the models must contain at least one utterance from the analyzed speaker s . Based on the aforementioned data settings, the average accuracy of all models for personality recognition of the five personality traits on both datasets is shown in Figure 4.

When the dialog flow is at 25%, Affective-NLI models significantly outperform other methods. It suggests that in the early stage of a dialogue (only one of two sentences from the analyzed speaker), the affective information and label semantics provide additional evidence to analyze the personality effectively. DesPrompt also performs relatively well as it provides an ample amount of descriptive label words.

When feeding 50%-75% of the whole dialog flow, Affective-NLI models also outperform the rest of the models. However, fine-tuning RoBERTa models (FT-RoBERTa (S), FT-RoBERTa (S+C), and FT-RoBERTa (F)) improve dramatically, and even outperform DesPrompt on FriendsPersona. This demonstrates the importance of rich semantic content in dialogues for traditional classification methods. Therefore, Affective-NLI effectively addresses the challenge of insuffi-

cient semantic information in early-stage conversations, which is a limitation of traditional classification methods.

The improvement rate of the model’s performance slows down from feeding 75% to the entire dialog flow. This suggests that the relationship between semantic information and effectiveness in personality recognition during a conversation is not simply linearly correlated.

To summarize, the strength of integrating affective information and label semantics in Affective-NLI is reflected in the early stage of conversations. With only one or two utterances from the analyzed speaker, Affective-NLI can instantly recognize the personality with around 0.5-0.6 in accuracy.

D. Case Study

We use a case study to illustrate the usage of Affective-NLI in real HCI application scenarios, as shown in Table V.

Affective-NLI is designed for personality recognition in conversations. Therefore, the main experiments intend to verify the personality recognition performance with existing dialog pieces. However, in real HCI applications, we still need to design conversation agents to select appropriate, emotional, and personalized dialogue strategies based on the recognized personality. The case study here demonstrates how Affective-NLI supports such dialogue strategies during interactions. In this example, we pretend to be Mrs. Thompson and engage in a conversation with the Agent. The probability of personality in the case study is obtained from Affective-NLI, and the Agent’s responses are generated by ChatGPT based on the identification results from Affective-NLI.

In Table V, the dialog content with affective annotations is sequentially inputted into Affective-NLI to recognize the personality trait of Mrs. Thompson. Specifically, when Mrs. Thompson responds to the Agent, Affective-NLI calculates the probabilities (confidences of recognition results) of Mrs. Thompson exhibiting all the big five personality traits displayed on the right side of the Table. To facilitate decision-making, we can set a probability threshold, denoted as δ (e.g., 60%). When the probability of Mrs. Thompson possessing a certain personality trait (1) exceeds δ and (2) is the highest among the five personality traits, the Agent can generate a response with a specific dialogue policy tailored to that personality trait.

VI. DISCUSSION

In this section, we discuss three practical issues of utilizing Affective-NLI in real applications.

Firstly, the recognition results of Affective-NLI cannot be directly used as diagnostic evidence for psychological therapy. Personality is a multifaceted concept that is investigated and discussed in various disciplines, including physiology, linguistics, and cognitive science. Therefore, rigorous personality assessment still relies on professional approaches. The purpose and advantage of Affective-NLI lie in its ability to provide preliminary evaluations of users’ personalities in conversational settings, facilitating a better conversational experience.

Secondly, Affective-NLI can only recognize the personality manifested during conversations, which may not reflect the users’ actual personalities. The content of conversations is constrained by different contexts, and individuals may not fully express their personalities due to politeness, cultural factors, and other considerations in specific situations. However, this limitation is inherent to the problem of personality recognition in conversation itself.

Thirdly, both datasets used to evaluate Affective-NLI consist of scripts from TV series rather than real-life conversations. Collecting conversations for personality analysis should be cautiously approached due to potential ethical and privacy concerns. To provide a stronger rationale for future data collection and usage, we initially validate our concept using dialog scripts from TV series. However, since the scripts we use are derived from TV series in various cultures and contexts, our model can still be either directly applied or fine-tuned with a small amount of real data to adapt to real application scenarios.

VII. CONCLUSION, LIMITATIONS AND FUTURE WORK

This paper presents Affective-NLI, a simple yet effective approach for enhancing dialog content with affective information and personality label descriptions, enabling accurate personality recognition in conversations. Affective-NLI provides interpretable recognition results and is easy to implement with various architectures of pre-trained language models. Through extensive experiments, we validate the effectiveness of Affective-NLI in personality recognition. Furthermore, Affective-NLI demonstrates its suitability for early-stage personality recognition in conversations.

Although the effectiveness of Affective-NLI has been validated, we identify some limitations that can be improved and expanded upon in future work. First, Affective-NLI follows the setting in existing studies on personality recognition in conversation, where personality recognition is performed through independent binary classification of each trait. However, this setting overlooks the correlations among different personality traits, which have been explored in psychology findings [43], [44]. In future research, we will also investigate how to incorporate these correlations into Affective-NLI for more accurate personality recognition. Second, Affective-NLI focuses solely on the textual modality. Personality can be manifested in non-verbal factors such as facial expressions [45] and face gestures [46], or physiological signals like Electrocardiogram (ECG) and heart rate [47]. In future work, we will also investigate transforming signals collected from pervasive systems into textual descriptions and achieving multimodal personality recognition.

ACKNOWLEDGEMENT

This work is conducted at the Research Institute for Artificial Intelligence of Things (RIAIoT) at PolyU and supported by the Hong Kong Jockey Club Charities Trust fund (No. 2021-0369), and PolyU Research and Innovation Office (No. BD4A).

REFERENCES

- [1] S. D'Alfonso, "Ai in mental health," *Current Opinion in Psychology*, vol. 36, pp. 112–117, 2020.
- [2] D. Yang, E.-S. Oh, and Y. Wang, "Hybrid physical education teaching and curriculum design based on a voice interactive artificial intelligence educational robot," *Sustainability*, vol. 12, no. 19, p. 8000, 2020.
- [3] J. G. Kahambing, "Chatgpt, public health communication and 'intelligent patient companionship'," *Journal of public health*, p. fdad028, 2023.
- [4] R. C.-S. Chang, H.-P. Lu, and P. Yang, "Stereotypes or golden rules? exploring likable voice traits of social robots as active aging companions for tech-savvy baby boomers in taiwan," *Computers in Human Behavior*, vol. 84, pp. 194–210, 2018.
- [5] N. Srinarong and J. Mongkolnavin, "A development of personality recognition model from conversation voice in call center context," in *The 12th International Conference on Advances in Information Technology*, 2021, pp. 1–5.
- [6] S. Fatahi, H. Moradi, and L. Kashani-Vahid, "A survey of personality and learning styles models applied in virtual environments with emphasis on e-learning environments," *Artificial Intelligence Review*, vol. 46, no. 3, pp. 413–429, 2016.
- [7] P. T. Costa and R. R. McCrae, "Normal personality assessment in clinical practice: The neo personality inventory," *Psychological assessment*, vol. 4, no. 1, p. 5, 1992.
- [8] J. Feist and G. Feist, "J.(2008). theories of personality," 2012.
- [9] E. A. Rissola, S. A. Bahrainian, and F. Crestani, "Personality recognition in conversations using capsule neural networks," in *IEEE/WIC/ACM International Conference on Web Intelligence*, 2019, pp. 180–187.
- [10] H. Jiang, X. Zhang, and J. D. Choi, "Automatic text-based personality recognition on monologues and multiparty dialogues using attentive networks and contextual embeddings (student abstract)," in *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 34, no. 10, 2020, pp. 13 821–13 822.
- [11] Y. Chen, W. Fan, X. Xing, J. Pang, M. Huang, W. Han, Q. Tie, and X. Xu, "Cped: A large-scale chinese personalized and emotional dialogue dataset for conversational ai," *arXiv preprint arXiv:2205.14727*, 2022.
- [12] Z. Wen, J. Cao, Y. Yang, H. Wang, R. Yang, and S. Liu, "Desprompt: Personality-descriptive prompt tuning for few-shot personality recognition," *Information Processing & Management*, vol. 60, no. 5, p. 103422, 2023.
- [13] A. Mehrabian, "Relationships among three general approaches to personality description," *The journal of Psychology*, vol. 129, no. 5, pp. 565–581, 1995.
- [14] —, "Analysis of the big-five personality factors in terms of the pad temperament model," *Australian journal of Psychology*, vol. 48, no. 2, pp. 86–92, 1996.
- [15] G. Saucier and L. R. Goldberg, "Evidence for the big five in analyses of familiar english personality adjectives," *European journal of Personality*, vol. 10, no. 1, pp. 61–77, 1996.
- [16] S. Roccas, L. Sagiv, S. H. Schwartz, and A. Knafo, "The big five personality factors and personal values," *Personality and social psychology bulletin*, vol. 28, no. 6, pp. 789–801, 2002.
- [17] T. Bickmore and J. Cassell, "Relational agents: a model and implementation of building user trust," in *Proceedings of the SIGCHI conference on Human factors in computing systems*, 2001, pp. 396–403.
- [18] J. W. Pennebaker, M. E. Francis, and R. J. Booth, "Linguistic inquiry and word count: Liwc 2001," *Mahway: Lawrence Erlbaum Associates*, vol. 71, no. 2001, p. 2001, 2001.
- [19] M. R. Mehl, S. D. Gosling, and J. W. Pennebaker, "Personality in its natural habitat: Manifestations and implicit folk theories of personality in daily life," *Journal of personality and social psychology*, vol. 90, no. 5, p. 862, 2006.
- [20] F. Mairesse and M. Walker, "Automatic recognition of personality in conversation," in *Proceedings of the Human Language Technology Conference of the NAACL, Companion Volume: Short Papers*, 2006, pp. 85–88.
- [21] 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.
- [22] J. Hu, L. Shen, and G. Sun, "Squeeze-and-excitation networks," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2018, pp. 7132–7141.
- [23] F. Iacobelli, A. J. Gill, S. Nowson, and J. Oberlander, "Large scale personality classification of bloggers," in *international conference on affective computing and intelligent interaction*. Springer, 2011, pp. 568–577.
- [24] A. Soufi, S. Hosseinpour, and A. M. Rahmani, "Personality classification based on profiles of social networks' users and the five-factor model of personality," *Human-centric Computing and Information Sciences*, vol. 8, no. 1, pp. 1–15, 2018.
- [25] Y. Ji, W. Wu, H. Zheng, Y. Hu, X. Chen, and L. He, "Is chatgpt a good personality recognizer? a preliminary study," 2023.
- [26] H. Rao, C. Leung, and C. Miao, "Can chatgpt assess human personalities? a general evaluation framework," 2023.
- [27] D. Watson and L. A. Clark, "On traits and temperament: General and specific factors of emotional experience and their relation to the five-factor model," *Journal of personality*, vol. 60, no. 2, pp. 441–476, 1992.
- [28] P. E. Ekman and R. J. Davidson, *The nature of emotion: Fundamental questions*. Oxford University Press, 1994.
- [29] M.-J. Han, C.-H. Lin, and K.-T. Song, "Robotic emotional expression generation based on mood transition and personality model," *IEEE Transactions on Cybernetics*, vol. 43, no. 4, pp. 1290–1303, 2013.
- [30] Z. Wen, J. Cao, R. Yang, S. Liu, and J. Shen, "Automatically select emotion for response via personality-affected emotion transition," in *Findings of the Association for Computational Linguistics: ACL-IJCNLP 2021*, 2021, pp. 5010–5020.
- [31] J. Broekens, B. Hilpert, S. Verberne, K. Baraka, P. Gebhard, and A. Plaat, "Fine-grained affective processing capabilities emerging from large language models," in *2023 11th International Conference on Affective Computing and Intelligent Interaction (ACII)*. IEEE, 2023, pp. 1–8.
- [32] M. M. Amin, R. Mao, E. Cambria, and B. W. Schuller, "A wide evaluation of chatgpt on affective computing tasks," 2023.
- [33] X. Gu, K. M. Yoo, and J.-W. Ha, "Dialogbert: Discourse-aware response generation via learning to recover and rank utterances," *arXiv preprint arXiv:2012.01775*, 2020.
- [34] T. Kim and P. Vossen, "Emoberta: Speaker-aware emotion recognition in conversation with roberta," *arXiv preprint arXiv:2108.12009*, 2021.
- [35] T. Wolf, L. Debut, V. Sanh, J. Chaumond, C. Delangue, A. Moi, P. Cistac, T. Rault, R. Louf, M. Funtowicz *et al.*, "Huggingface's transformers: State-of-the-art natural language processing," *arXiv preprint arXiv:1910.03771*, 2019.
- [36] L. R. Goldberg and G. Saucier, "So what do you propose we use instead? a reply to block," 1995.
- [37] P. Quaglio, "Television dialogue: The sitcom friends vs. natural conversation. amsterdam, nl," 2009.
- [38] Y. Liu, M. Ott, N. Goyal, J. Du, M. Joshi, D. Chen, O. Levy, M. Lewis, L. Zettlemoyer, and V. Stoyanov, "Roberta: A robustly optimized bert pretraining approach," *arXiv preprint arXiv:1907.11692*, 2019.
- [39] C. Raffel, N. Shazeer, A. Roberts, K. Lee, S. Narang, M. Matena, Y. Zhou, W. Li, and P. J. Liu, "Exploring the limits of transfer learning with a unified text-to-text transformer," *The Journal of Machine Learning Research*, vol. 21, no. 1, pp. 5485–5551, 2020.
- [40] H. Touvron, T. Lavril, G. Izacard, X. Martinet, M.-A. Lachaux, T. Lacroix, B. Rozière, N. Goyal, E. Hambro, F. Azhar *et al.*, "Llama: Open and efficient foundation language models," *arXiv preprint arXiv:2302.13971*, 2023.
- [41] X. Liu, K. Ji, Y. Fu, W. Tam, Z. Du, Z. Yang, and J. Tang, "P-tuning: Prompt tuning can be comparable to fine-tuning across scales and tasks," in *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, 2022, pp. 61–68.
- [42] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," *arXiv preprint arXiv:1412.6980*, 2014.
- [43] R. R. McCrae and P. T. Costa, "Validation of the five-factor model of personality across instruments and observers," *Journal of personality and social psychology*, vol. 52, no. 1, p. 81, 1987.
- [44] L. R. Goldberg, "The structure of phenotypic personality traits," *American psychologist*, vol. 48, no. 1, p. 26, 1993.
- [45] L. P. Naumann, S. Vazire, P. J. Rentfrow, and S. D. Gosling, "Personality judgments based on physical appearance," *Personality and social psychology bulletin*, vol. 35, no. 12, pp. 1661–1671, 2009.
- [46] M. M. LATHA, V. L. CHAITANYA, and M. S. DEVI, "By face gesture validating the personality of a person,"
- [47] X. Shui, Y. Chen, X. Hu, F. Wang, and D. Zhang, "Personality in daily life: Multi-situational physiological signals reflect big-five personality traits," *IEEE Journal of Biomedical and Health Informatics*, 2023.