

Identifying Multiple Personalities in Large Language Models with External Evaluation

Xiaoyang Song¹, Yuta Adachi^{2*}, Jessie Feng^{2*}, Mouwei Lin^{2*}, Linhao Yu^{2*}, Frank Li^{2*}, Akshat Gupta³, Gopala Anumanchipalli³, Simerjot Kaur⁴

¹University of Michigan, ²Columbia University, ³UC Berkeley, ⁴ AI Research, JPMorgan
xysong@umich.edu, akshat.gupta@berkeley.edu, simerjot.kaur@jpmchase.com

Abstract

As Large Language Models (LLMs) are integrated with human daily applications rapidly, many societal and ethical concerns are raised regarding the behavior of LLMs. One of the ways to comprehend LLMs' behavior is to analyze their personalities. Many recent studies quantify LLMs' personalities using self-assessment tests that are created for humans. Yet many critiques question the applicability and reliability of these self-assessment tests when applied to LLMs. In this paper, we investigate LLM personalities using an alternate personality measurement method, which we refer to as the *external evaluation* method, where instead of prompting LLMs with multiple-choice questions in the Likert scale, we evaluate LLMs' personalities by analyzing their responses toward open-ended situational questions using an external machine learning model. We first fine-tuned a Llama2-7B model as the MBTI personality predictor that outperforms the state-of-the-art models as the tool to analyze LLMs' responses. Then, we prompt the LLMs with situational questions and ask them to generate *Twitter posts* and *comments*, respectively, in order to assess their personalities when playing two different roles. Using the external personality evaluation method, we identify that the obtained personality types for LLMs are significantly different when generating posts versus comments, whereas humans show a consistent personality profile in these two different situations. This shows that LLMs can exhibit different personalities based on different scenarios, thus highlighting a fundamental difference between personality in LLMs and humans. With our work, we call for a re-evaluation of personality definition and measurement in LLMs.

1 Introduction

¹ The evolution of Large Language Models (LLM) has benefited humans in the past few years through

^{1*} equal contribution

their unprecedented capacities to understand and generate human-like languages (Radford et al., 2018, 2019; Brown et al., 2020; Ouyang et al., 2022; OpenAI, 2023, 2022; Zhang et al., 2022; Touvron et al., 2023a). For instance, LLMs are now being deployed as virtual assistants to provide mental health support (Lai et al., 2023), as online educators for common knowledge retrieval (OpenAI, 2022; Jeon and Lee, 2023), and even as helpers in symbolic music compositions (Agostinelli et al., 2023; Imasato et al., 2023). However, this growing integration of LLMs across different social sectors of human life raises important concerns about reliability, safety, and ethics. This dual nature of LLMs opens the need to study their behaviors, especially their behaviors when interacting with humans. Although most chat-based models including ChatGPT (OpenAI, 2022) and Llama (Touvron et al., 2023b) have undergone safety training to prevent delivering poisonous and biased information, there is still an urgent need to find a proper venue and metrics to understand their societal behaviors.

One common way to understand the behavior of LLMs is to measure their personalities through rigorous psychometric studies (Jiang et al., 2022; Miotto et al., 2022; Huang et al., 2023; Caron and Srivastava, 2022; Karra et al., 2022). According to the American Psychological Association (APA), personality for humans is defined as "*the enduring characteristic and behavior that comprise a person's unique adjustment to life*" (Association, 2023). On the other hand, the exact definition of LLM personalities remains an open yet mysterious question in the field. Nevertheless, many researchers made analogies to human personality and attempted to study LLM personalities by leveraging the psychometric tests used for humans. For instance, most recent literature prompted LLMs with standardized self-assessment personality test questions and then recorded and analyzed the results (Jiang et al., 2023; Karra et al., 2022; Miotto

et al., 2022; Bodroza et al., 2023; Safdari et al., 2023). However, while these tests are shown to be effective for human personality measurement (Digman, 1990), there is evidence that they can not be directly applied to reliably measure personality of both base LLMs including GPTs (Radford et al., 2018, 2019; Brown et al., 2020) and chat-based LLMs like ChatGPT (OpenAI, 2022) and Llama (Touvron et al., 2023b). For instance, Song et al. (2023) and Gupta et al. (2023) managed to show that the self-assessment personality test results of the same LLM differ significantly as the prompting templates are changed, which is not surprising as LLMs are known to be prone and sensitive to different prompts (Sclar et al., 2023; Chen et al., 2023). However, it has also been shown that even under the same prompt template, the psychometric test results are statistically different when the options of those multiple-choice questions are presented in different orders (Song et al., 2023; Gupta et al., 2023). These observations indicate that prompting LLMs with standardized self-assessment is not a reliable method to quantify LLMs’ personalities. Therefore, an alternative to self-assessment psychometric tests is desired to better understand and analyze personalities in LLMs.

In this paper, we investigate an alternate method to measure LLM personalities. To do so, we first develop a state-of-the-art personality prediction model. Specifically, we utilized the famous Myers-Briggs Type Indicator (MBTI) personality framework and fine-tuned a Llama2-7B (Touvron et al., 2023b) model on a human personality dataset. The dataset contains multiple posts written by a human subject and their MBTI personality type. This model is then used to analyze the personality of LLMs under the MBTI framework and compare the results with human counterparts. (Sec. 3.1)

To measure LLM personality, we have different LLMs write tweets which are used as input for our personality prediction model. While doing this, we have LLMs take two different roles. In the first role, the LLM is asked to write *tweets* about real-world events based on the event topics which were obtained by analyzing news articles. This is done to prevent data leakage and stop LLMs from repeating tweets seen previously in pre-training data. In the second role, we ask the LLM to write *replies* to existing tweets. The tweets are again collected in real time and are not part of the model pre-training corpus. We then evaluate the personality of different LLMs based on the tweets generated by them

using our personality prediction model. We perform this analysis for ChatGPT (OpenAI, 2022), Llama2-7B-chat, Llama2-13B-chat and Llama2-70B-chat models (Touvron et al., 2023b). At the same time, we also repeat the same procedure for human-written posts and comments to validate the proposed personality detection model. To our surprise, we find that the personality distribution of LLMs when writing tweets is completely different from the distribution when responding to tweets. As defined by APA, personality is supposed to be an *enduring characteristic* for humans. While this consistency is shown to be true for humans, we show that LLMs exhibit different personalities while playing different roles. (Sec. 3.2 & Sec. 3.3).

To summarize, our paper makes the following contributions:

- We fine-tune a Llama2-7B for MBTI personality detection that significantly outperforms the state-of-the-art models.
- We use the fine-tuned personality detection model on tweets made by humans and show that human personality remains consistent across different roles.
- We show that LLMs exhibit different personalities across different roles when using the external evaluation method.
- With our work, we call for a re-evaluation of personality definition and measurement in LLMs.

2 Related Work

2.1 Personality Theory

Human personality is defined as “*the enduring characteristic and behavior that comprise a person’s unique adjustment to life*” by APA (Association, 2023), which is typically measured across different effective trait dimensions (Cattell, 1943b,a). Central to human psychological profiling are the *Big Five* personality traits, also referred to as the OCEAN traits, which stands for **O**penness, **C**onscientiousness, **E**xtraversion, **A**greeableness, and **N**euroticism, respectively (Digman, 1990; Goldberg, 1990, 1993). The *Big Five* personality traits are typically measured through multiple choice self-assessment questions where a situational statement is presented and the test-takers are requested to choose from an option in Likert scales, typically 1 to 5, to reflect the degree of

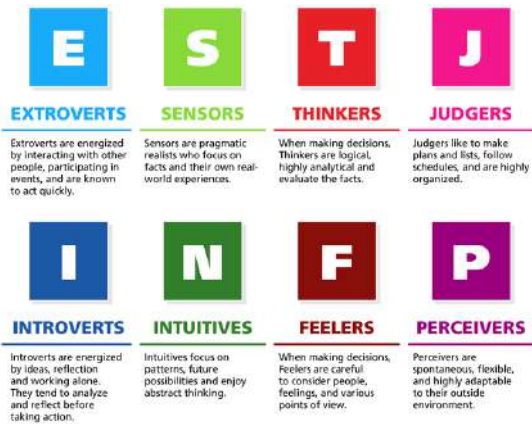


Figure 1: MBTI trait description (Faust, 2019).

fitness of the statement to themselves (Johnson, 2014). Many studies on LLM personality directly base their works on this straightforward test (Jiang et al., 2023; Song et al., 2023; Caron and Srivastava, 2022).

However, the result of this test is a distribution of scores among all five traits, making it difficult to analyze and compare. Instead, we build our paper on another famous categorical personality framework called the Myers-Briggs Type Indicator (MBTI), which details 16 distinct types based on trait combinations across four different dimensions: Extraverts vs. Introverts (E/I), Sensors vs. Intuitives (S/I), Thinkers vs. Feelers (T/F), and Judgers vs. Perceivers (J/P). Figure 1 above provides detailed information about the trait each dimension measures.

2.2 LLM Personality Measurement

Many recent works regarding LLM personality asked the LLM to perform Multiple-Choice Question-Answering (MCQA) on those well-known standardized self-assessment personality tests (Jiang et al., 2023; Miotto et al., 2022; Caron and Srivastava, 2022; Huang et al., 2023; Bodroza et al., 2023; Safdari et al., 2023; Pan and Zeng, 2023; Noever and Hyams, 2023). For instance, as one of the foundation works, Jiang et al. (2023) prompted the LLMs with the widely-used IPIP-120 dataset in psychology (Johnson, 2014) and evaluated their *Big Five* scores (i.e. OCEAN scores). However, although these works managed to elaborate on how personality can play a role in LLMs, the validity of these self-assessment tests remains unchecked, making their conclusions unreliable. As shown by Song et al. (2023) and Gupta et al. (2023), the same LLM tends to exhibit significantly

different personalities on different self-assessment tests, which fails to achieve the crucial *consistency* criterion in personality definition and highlights the invalidity of self-assessment tests. The origins of these concerns are related to both the difficulty of MCQA tasks and the fact that LLMs are sensitive to the change in prompts as well as the structure of the prompts. A well-known example of this is that the LLMs can be manipulated logically via Chain-of-Thoughts (CoT) (Wei et al., 2022).

To the best of our knowledge, there is not much literature that attempts to evaluate LLM personality without leveraging standardized tests. Driven by this gap, in this work, we attempt to completely discard self-assessment tests and use an external personality assessment method, where a personality detection model is used to predict LLM personality. In addition, unlike previous works that conduct tests without validation, we also justify the personality prediction model used for evaluation by conducting a validation experiment on humans.

2.3 MBTI Personality Detection Models

Most MBTI personality detection models are built based on the public Kaggle dataset² (Tang et al., 2023; Yang et al., 2023; Mehta, 2023). The dataset consists of over 8600 entries, where each entry corresponds to an individual’s MBTI type and includes excerpts from the last 50 posts and comments made by the individual on the PersonalityCafe forum³. Detailed introduction to example entry, label distribution, and the tweet topic distributions can be found in Appendix A.1.

To predict MBTI personality types from texts, Mehta (2023) processed the texts using a BERT model (Devlin et al., 2018) and then trained an MLP to predict the MBTI personality type, while Yang et al. (2023) utilized a graph convolutional neural network to learn the connections between different posts made by an individual in order to make decent predictions. Furthermore, Tang et al. (2023) proposed an attention-based denoising framework (ADF) for MBTI personality detection, where they trained the model to effectively extract personality signals from noisy and verbose text data. However, a common weakness in these works is that their methods predict the personality type for each MBTI dimension separately. Therefore, although their methods achieve an average

²<https://www.kaggle.com/datasets/datasnaek/mbti-type>

³<https://www.personalitycafe.com/forums/myers-briggs-forum.49/>

of around 70% prediction accuracy for each trait dimension, the overall performance remains mysterious when the predictions for each dimension are aggregated together because the errors and uncertainties accumulate multiplicatively. In our method, we fine-tuned the popular Llama2-7B (Touvron et al., 2023b) model as the predictor, which outperformed their models significantly.

3 Experiments

In this paper, we refer to using personality prediction model to measure personality as an *external personality evaluation method*, juxtaposing it with the personality self-assessment methods. The experiments on the proposed external evaluation method are divided into three stages. Firstly, a Llama2-7B (Touvron et al., 2023b) model is fine-tuned on the public MBTI datasets as the personality detection model. Then, to perform external evaluations, we prompted the LLMs with pre-processed situational questions/scenarios and prompted them for open-ended generations based on the input. Finally, the collected responses from the LLMs were fed into the personality detection model to obtain the evaluation results. In this paper, we base our experiments on four popular chat-based LLMs: ChatGPT (OpenAI, 2022) and three versions of Llama2 (7B, 13B & 70B) (Touvron et al., 2023b). The overall experiment pipeline can be found in Figure 2. In addition, the computational resources used are introduced in Appendix A.2.

3.1 Stage I: Fine-tuning Llama2-7B-based Personality Detection Model

In this work, we utilized the widely used aforementioned Kaggle MBTI dataset to fine-tune the personality detection model. Although many previous studies build their models on the same dataset (Yang et al., 2023; Tang et al., 2023; Mehta, 2023), there is room for improvement in the overall performance of their models, which may affect the reliability of personality prediction.

In particular, we proposed two models: (1) a binary model and (2) a 16-class model. In terms of the binary model, we fine-tuned four pre-trained Llama2-7B models, and each of them is designed to be a binary classifier for each of the four MBTI dimensions. The predictions are then aggregated to form the final prediction, which is similar to most of the previous works (Tang et al., 2023; Yang et al., 2023; Mehta, 2023). On the other hand, in

the 16-classes model, a pre-trained Llama2-7B is directly fine-tuned to predict one of the sixteen classes. Aligning with the Kaggle dataset format, both models take 50 posts from an individual as inputs at a time to predict personality. The fine-tuning process was conducted under the LoRA (Hu et al., 2021) framework, where the target modules to fine-tune are the query (q) and the value (v) layers. We chose the rank to be $r = 16$ and fine-tuned it for only 5 epochs with a learning rate of 10^{-4} and batch size of 8. The training started with a warm-up phase for the first 100 iterations, followed by a linear learning rate decay.

Models	Accuracy	F1 score	Precision	Recall
BERT-Base + MLP (Mehta, 2023)	73.1	63.2	78.0	65.5
D-DGCN (Yang et al., 2023)	78.2	53.5	60.5	49.1
ADF (Tang et al., 2023)	61.0	38.8	47.5	56.8
Llama2-7B + FT (B) (ours)	93.3	91.1	92.3	90.1
Llama2-7B + FT (16) (ours)	93.5	91.4	92.0	90.8

Table 1: Average performance of personality detection models on each of the four MBTI dimensions. Each cell is a numerical average of the performance of each binary model in percentage.

Models	Accuracy	F1 score	Precision	Recall
BERT-Base + MLP (Mehta, 2023)	29.4	8.70	14.1	12.0
D-DGCN (Yang et al., 2023)	37.6	35.9	35.9	37.6
ADF (Tang et al., 2023)	14.0	4.10	3.28	5.74
Llama2-7B + FT (B) (ours)	81.0	74.3	76.9	73.1
Llama2-7B + FT (16) (ours)	81.7	76.9	79.8	75.2

Table 2: Performance of personality detection models on 16-class MBTI prediction tasks. The numbers in the cell are all percentages.

The testing performance of the proposed personality detection model and the baselines are reported. With the capacities of LLMs to understand texts and extract key knowledge, both the fine-tuned binary model and 16-classes model outperform the state-of-the-art models significantly in all metrics. In Table 1, we report the average performance of the model for each trait dimension. To do this, the accuracy of the model in predicting each of the four dimensions of MBTI personality is calculated separately. Then the accuracy for these four dimensions is averaged to present the numbers in Table 1,

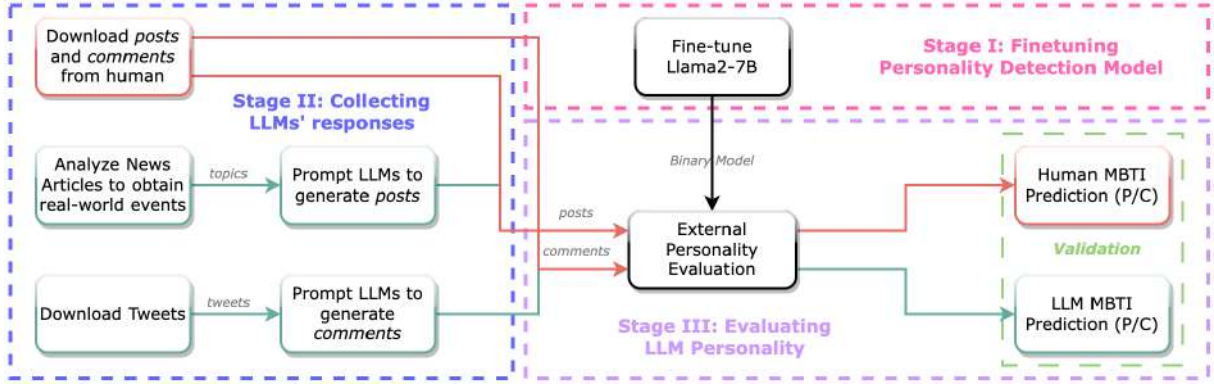


Figure 2: Methodology flowchart.

which is the most common form of evaluation used in prior work (Tang et al., 2023; Yang et al., 2023).

One drawback of such an evaluation is that it doesn’t take into account the overall accuracy of the model in predicting one of the 16 classes of personality. In Table 2, we take into account the overall prediction errors and present the aggregated 16-class classification performance. The difference in performance between Table 1 and Table 2 for the same models highlights the effectiveness of our model. For instance, while the baseline methods achieve around 70% accuracy when averaged over each trait, when those predictions are aggregated together to predict the actual 16-classes MBTI type, the accuracy drops significantly to around 30%. However, for our fine-tuned binary model, the prediction accuracy for each trait is 93.3% while the overall accuracy remains as high as 81.0%. This evaluation method helps highlight error accumulation during the personality evaluation procedure, rendering the experimental results more accurate. We observed that the binary model and 16-classes model yield comparable performance for nearly all the metrics. In later experiments, we chose to stick to the binary model as it is often used in previous works (Tang et al., 2023; Yang et al., 2023; Mehta, 2023), making the results comparable. Complete results for fine-tuned models can be found in Appendix A.3.

3.2 Stage II: Collecting LLMs Responses Toward Open-ended Situational Questions

In the second stage, we prompt LLMs with open-ended situational questions and collect their responses as inputs for external evaluation. This is done to collect social media post data generated by an LLM which can be used as input for our

personality prediction model.

Prompting Pipeline. To effectively analyze daily news, prompt LLMs, and process their responses, we designed a prompting pipeline. The pipeline starts by analyzing news from different topics and summarizing them into the latest news events before prompting the LLMs. We analyzed the latest news articles from November 2023 to make sure that the events were not present in the training data of the personality prediction models. These latest news events serve as the basis for prompting the LLMs to generate tweets, with the LLMs taking on the role of individuals posting tweets about these news events. Additionally, we also gather online tweets written by other individuals and prompt the LLMs to generate comments in response to these tweets. This allows the LLMs to assume different roles, specifically that of replying to tweets made by others. To prevent any potential data leakage, these tweets are collected exclusively from the month of November 2023.

Depending on whether we want the LLMs to write *posts* (i.e. *tweets*) about the summarized events or *comments* (i.e. *replies*) to the existing tweets, the prompting templates are chosen to be different accordingly. For instance, the template for prompting LLMs to write a post is the following: “As a user on Twitter. Write a tweet on the following contents: [summarized contents]”. It is noteworthy that we chose not to engineer the prompting templates in order to keep them as simple as possible, so as to not induce behavior of LLMs. The exact prompts used in this work are shown in Table 3. We use the simplest prompts to elicit generation in order to not bias the model through prompts for generations.

In total, we analyzed several thousand news arti-

Type	System Prompt Used	User Prompt Used
Posts	Generate a Twitter post	As a user on Twitter, write a tweet on the following contents: [summarized contents]
Comments	Generate a Twitter comment	As a user on Twitter, write a tweet to comment on this Tweet: [tweet contents]

Table 3: Prompting templates.




Figure 3: MBTI distribution of ChatGPT and Llama2 models using external evaluation method for 100 times. In the figure, the first row is the assessment results on the generated posts dataset (P), whereas the second row provides the results on the comments dataset (C). In addition, personality types that appear less than 3 times are merged together to form the class “Others”.

cles to extract relevant news events and collected 5000 tweets from 10 popular topics. Subsequently, we tasked the LLMs with generating 4500 posts based on the news events and 5000 comments in response to the collected tweets. These generated texts were then utilized to study the personalities of the LLMs using our personality detection model. It is worth noting that the roles assigned to the LLMs varied depending on whether they were asked to generate posts or comments, allowing for a diverse range of responses. A detailed description of this pipeline and downloaded contents can be found in Appendix A.4.

3.3 Stage III: External LLMs Personality Detection & Validation

After obtaining the posts and comments datasets for each LLM of interest, we evaluate their personalities using our fine-tuned personality detection model. The personality detection model is trained to take 50 social media posts of a user as input and output their personality based on the 16-class MBTI personality framework. To minimize errors

due to sampling and model inaccuracy, we created 100 sets of 50 generations for each type of tweet (*posts vs. comments*) generated by the LLM. To do so, we sampled 50 responses with replacements from our generation dataset 100 times. We then report the predicted MBTI personality distribution for the 100 samples for each role in Figure 3 and the most frequent personality type for each LLM is reported as the predicted personality in Table 4.

Model	External Evaluation	
	Posts	Comments
ChatGPT	INFJ	INFJ
Llama2-7B	ESTJ	INFP
Llama2-13B	ESTJ	INFP
Llama2-70B	ESTJ	INFJ

Table 4: MBTI personality type obtained by external evaluation method. The most frequent MBTI personality type is reported here.

We can observe from Table 4 that for the same LLM, the external evaluation results on the generated posts and comments datasets are very different. For instance, for Llama2-7B, the most frequent




Figure 4: MBTI distribution of 4 celebrities using the external evaluation method for 100 times. In the figure, the first row is the assessment results on the generated posts dataset (P), whereas the second row provides the results on the comments dataset (C). In addition, personality types that appear less than 3 times are merged together to form the class “Others”. Complete results for all 8 selected celebrities can be found in Appendix A.5.

MBTI type is ESTJ when it posts tweets, whereas it is INFP on the generated comments. In addition, the predicted personality distributions for the 100 trials in Figure 3 are also significantly different between posts and comments. For example, in ChatGPT personality evaluation, although the most frequent MBTI type is the same in Table 4 between posts and comments, the predicted personality distributions are extremely different, as can be seen in the first column of pie charts in Figure 3. We can see that INFJ and INFP are nearly equally distributed when writing posts, whereas INFJ dominates over all other MBTI types when writing comments. In plain words, this indicates ChatGPT behaves very differently when generating posts and comments. We can see this very clearly for the other Llama2 models, where the personality distribution while writing posts (first row of pie charts in Figure 3) is very different from the personality distribution we get when writing comments (second row in Figure 3). **These experiments clearly show that LLMs exhibit different personalities when playing different roles.**

Validation with Human Counterpart. Personality is defined as an *enduring characteristic* in humans by the American Psychological Association. This means that for humans, the personality distribution obtained for humans when writing posts versus comments is expected to be consistent. As we employed a personality detection model to

evaluate LLM personality, it becomes important to check if the detection model is able to produce similar personality distributions when humans take on the two roles of writing posts versus comments. If this is not the case, then the inconsistencies in LLM personalities in the two roles can be attributed to the personality detection model. Furthermore, this is crucial to justify the correctness of our later analysis, which is completely based on this personality detection model.

To check this, we conduct the same study with human counterparts. We randomly downloaded the *Twitter posts* and *comments* from 8 celebrities⁴ from different domains and applied the personality detection model to these two different datasets. We again created 100 samples of 50 tweets each and found the personality distribution for these celebrities in the two different roles - while writing posts versus comments. The personality distribution results for four celebrities can be seen in Figure 4, while the remaining four can be seen in Figure 5.

Table 5 presents the results for the eight celebrities. We can see that the most frequent MBTI type, the second most frequent MBTI type, and even the predicted distributions (Figure 4) are very akin to each other for nearly all the celebrities that we investigated. This empirically confirms that humans tend to exhibit *enduring* MBTI personality types

⁴Due to proprietary reasons, we have masked the identities of the selected celebrities.

Celebrity	External Evaluation	
	P	C
Celebrity I	INTJ (INFJ)	INTJ (INFJ)
Celebrity II	INFJ (ENFJ)	INFJ (ENFJ)
Celebrity III	INFJ (ENFJ)	INFJ (ENFJ)
Celebrity IV	INFJ (ENFJ)	INFJ (ENFJ)
Celebrity V	INFJ (INFP)	INFJ (INFP)
Celebrity VI	INFJ (INTJ)	ENTJ (ENFJ)
Celebrity VII	INFJ (INFP)	INFJ (INFP)
Celebrity VIII	INFJ (ISFJ)	INFJ (ENFJ)

Table 5: The most frequent MBTI personality type for 8 celebrities under external evaluation method using *posts* and *comments*. The second most frequent MBTI personality type is reported in the parenthesis.

even when playing different roles, which echoes the observations in most of the psychology literature (Green et al., 2019; Association, 2023) as well as the personality definition (American Psychological Association, 1983). A detailed introduction to this validation experiment can be found in Appendix A.5. Simultaneously, this confirms that our fine-tuned model is able to attain similar personality distributions for humans and rules out its effects in the observed *inconsistency* in LLMs.

The above experiments show that LLMs clearly exhibit different personalities in different roles, which is not true for humans. Our experiments also show that the definition of personality as defined for humans may not be applicable to LLMs, as personality in LLMs no longer seems to be an enduring characteristic.

4 Conclusions

In this paper, we investigate the personality exhibited by LLMs using an *external personality evaluation* method, which is an alternative to self-assessment tests for personality measurement. To perform external personality evaluation, we fine-tuned a Llama2-7B model as the third-party agent for personality detection and collected LLMs’ responses and attitudes toward open-ended situational questions. We also developed a prompting pipeline to automate the prompts for LLMs with proper content. Our experiments show that LLMs exhibit different personalities in different roles, which is not true for humans.

These observations also shed light on the validity of quantifying and defining LLM personality using the same standard for humans. In our work, we see that LLM personality is not an enduring characteristic as it is seen in humans. Additionally, previous

works have shown that self-assessment tests are not the appropriate tool to measure personality in LLMs. When we combine our observations with these results, we can conclude that not only do we not have accurate tools to measure LLM personality, but we also lack the appropriate definition of personality for LLMs. With our work, we caution against a naive transfer of definition and methods used to analyze human personality on LLMs and call for more fundamental work on evaluating LLM personality and behavior, taking into account the specific characteristics of LLMs.

5 Limitation

This paper investigates the validity of applying external evaluation to measuring LLM personalities and calls for a re-evaluation of the definition and measurement of LLM personalities. As we employed a fine-tuned model as an external agent to evaluate personalities and based our analysis on it, one of the limitations of our work is the accuracy of the personality detection model. Although our model attains state-of-the-art performance and we reduce the chances of errors by using 100 samples of generations instead of 1, the uncertainties introduced due to error rates are still inevitable. Furthermore, although our work identifies that the current definition of LLM personality should be reevaluated, this paper does not provide an alternative for defining LLM personality and the measurement method, which are left for future work.

6 Acknowledgement

This paper was prepared for information purposes by the Artificial Intelligence Research group of JPMorgan Chase & Co and its affiliates (“JP Morgan”), and is not a product of the Research Department of JP Morgan. J.P. Morgan makes no representation and warranty whatsoever and disclaims all liability for the completeness, accuracy or reliability of the information contained herein. This document is not intended as investment research or investment advice, or a recommendation, offer or solicitation for the purchase or sale of any security, financial instrument, financial product or service, or to be used in any way for evaluating the merits of participating in any transaction, and shall not constitute a solicitation under any jurisdiction or to any person, if such solicitation under such jurisdiction or to such person would be unlawful. © 2023 JP Morgan Chase & Co. All rights reserved.

References

- Andrea Agostinelli, Timo I Denk, Zalán Borsos, Jesse Engel, Mauro Verzetti, Antoine Caillon, Qingqing Huang, Aren Jansen, Adam Roberts, Marco Tagliasacchi, et al. 2023. Musiclm: Generating music from text. *arXiv preprint arXiv:2301.11325*.
- American Psychological Association. 1983. *Publications Manual*. American Psychological Association, Washington, DC.
- American Psychological Association. 2023. *Definition of Personality* - <https://www.apa.org/topics/personality>.
- Bojana Bodroza, Bojana M Dinic, and Ljubisa Bojic. 2023. Personality testing of gpt-3: Limited temporal reliability, but highlighted social desirability of gpt-3's personality instruments results. *arXiv preprint arXiv:2306.04308*.
- Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. 2020. Language models are few-shot learners. *Advances in neural information processing systems*, 33:1877–1901.
- Graham Caron and Shashank Srivastava. 2022. Identifying and manipulating the personality traits of language models. *arXiv preprint arXiv:2212.10276*.
- Raymond B Cattell. 1943a. The description of personality: Basic traits resolved into clusters. *The journal of abnormal and social psychology*, 38(4):476.
- Raymond B Cattell. 1943b. The description of personality. i. foundations of trait measurement. *Psychological review*, 50(6):559.
- Banghao Chen, Zhaofeng Zhang, Nicolas Langrené, and Shengxin Zhu. 2023. Unleashing the potential of prompt engineering in large language models: a comprehensive review. *arXiv preprint arXiv:2310.14735*.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2018. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*.
- John M Digman. 1990. Personality structure: Emergence of the five-factor model. *Annual review of psychology*, 41(1):417–440.
- Katrina Faust. 2019. Myers-briggs type indicator (mbti) overview - <https://leadx.org/articles/mbti-myers-briggs-type-indicator-overview/>.
- Lewis R Goldberg. 1990. An alternative" description of personality": the big-five factor structure. *Journal of personality and social psychology*, 59(6):1216.
- Lewis R Goldberg. 1993. The structure of phenotypic personality traits. *American psychologist*, 48(1):26.
- Jennifer P Green, Reeshad S Dalal, Kristen L Swigart, Melissa A Bleiberg, David M Wallace, and Amber K Hargrove. 2019. Personality consistency and situational influences on behavior. *Journal of Management*, 45(8):3204–3234.
- Akshat Gupta, Xiaoyang Song, and Gopala Anumanchipalli. 2023. Investigating the applicability of self-assessment tests for personality measurement of large language models. *arXiv preprint arXiv:2309.08163*.
- Ari Holtzman, Jan Buys, Li Du, Maxwell Forbes, and Yejin Choi. 2019. The curious case of neural text degeneration. *arXiv preprint arXiv:1904.09751*.
- Edward J Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, and Weizhu Chen. 2021. Lora: Low-rank adaptation of large language models. *arXiv preprint arXiv:2106.09685*.
- Jen-tse Huang, Wenxuan Wang, Man Ho Lam, Eric John Li, Wenxiang Jiao, and Michael R Lyu. 2023. Chatgpt an enfj, bard an istj: Empirical study on personalities of large language models. *arXiv preprint arXiv:2305.19926*.
- Naomi Imasato, Kazuki Miyazawa, Caitlin Duncan, and Takayuki Nagai. 2023. Using a language model to generate music in its symbolic domain while controlling its perceived emotion. *IEEE Access*.
- Jaeho Jeon and Seongyong Lee. 2023. Large language models in education: A focus on the complementary relationship between human teachers and chatgpt. *Education and Information Technologies*, pages 1–20.
- Guangyuan Jiang, Manjie Xu, Song-Chun Zhu, Wenjuan Han, Chi Zhang, and Yixin Zhu. 2022. Mpi: Evaluating and inducing personality in pre-trained language models. *arXiv preprint arXiv:2206.07550*.
- Guangyuan Jiang, Manjie Xu, Song-Chun Zhu, Wenjuan Han, Chi Zhang, and Yixin Zhu. 2023. [Evaluating and inducing personality in pre-trained language models](#).
- John A Johnson. 2014. Measuring thirty facets of the five factor model with a 120-item public domain inventory: Development of the ipip-neo-120. *Journal of research in personality*, 51:78–89.
- Saketh Reddy Karra, Theja Tulabandhula, et al. 2022. Estimating the personality of white-box language models. *arXiv e-prints*, pages arXiv–2204.
- Tin Lai, Yukun Shi, Zicong Du, Jiajie Wu, Ken Fu, Yichao Dou, and Ziqi Wang. 2023. Psy-llm: Scaling up global mental health psychological services with ai-based large language models. *arXiv preprint arXiv:2307.11991*.

- Yash Mehta. 2023. [Personality prediction](https://github.com/yashmehta/personality-prediction). <https://github.com/yashmehta/personality-prediction.git>. GitHub repository.
- Yash Mehta, Samin Fatehi, Amirmohammad Kazameini, Clemens Stachl, Erik Cambria, and Sauleh Eetemadi. 2020. Bottom-up and top-down: Predicting personality with psycholinguistic and language model features. In *2020 IEEE International Conference on Data Mining (ICDM)*, pages 1184–1189. IEEE.
- Mariù Miotto, Nicola Rossberg, and Bennett Kleinberg. 2022. Who is gpt-3? an exploration of personality, values and demographics. *arXiv preprint arXiv:2209.14338*.
- David Noever and Sam Hyams. 2023. Ai text-to-behavior: A study in steerability. *arXiv preprint arXiv:2308.07326*.
- OpenAI. 2022. Chatgpt - <https://openai.com/blog/chatgpt#OpenAI>.
- OpenAI. 2023. Gpt-4 technical report - <https://cdn.openai.com/papers/gpt-4.pdf>.
- Long Ouyang, Jeffrey Wu, Xu Jiang, Diogo Almeida, Carroll Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, et al. 2022. Training language models to follow instructions with human feedback. *Advances in Neural Information Processing Systems*, 35:27730–27744.
- Keyu Pan and Yawen Zeng. 2023. Do llms possess a personality? making the mbti test an amazing evaluation for large language models. *arXiv preprint arXiv:2307.16180*.
- Alec Radford, Karthik Narasimhan, Tim Salimans, Ilya Sutskever, et al. 2018. Improving language understanding by generative pre-training.
- Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, Ilya Sutskever, et al. 2019. Language models are unsupervised multitask learners. *OpenAI blog*, 1(8):9.
- Mustafa Safdari, Greg Serapio-García, Clément Crepy, Stephen Fitz, Peter Romero, Luning Sun, Marwa Abdulhai, Aleksandra Faust, and Maja Matarić. 2023. Personality traits in large language models. *arXiv preprint arXiv:2307.00184*.
- Melanie Sclar, Yejin Choi, Yulia Tsvetkov, and Alane Suhr. 2023. Quantifying language models’ sensitivity to spurious features in prompt design or: How i learned to start worrying about prompt formatting. *arXiv preprint arXiv:2310.11324*.
- Xiaoyang Song, Akshat Gupta, Kiyan Mohebbizadeh, Shujie Hu, and Anant Singh. 2023. Have large language models developed a personality?: Applicability of self-assessment tests in measuring personality in llms. *arXiv preprint arXiv:2305.14693*.
- Qirui Tang, Wenkang Jiang, Yihua Du, and Lei Lin. 2023. An attention-based denoising framework for personality detection in social media texts. *arXiv preprint arXiv:2311.09945*.
- Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajjwal Bhargava, Shruti Bhosale, et al. 2023a. Llama 2: Open foundation and fine-tuned chat models. *arXiv preprint arXiv:2307.09288*.
- Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajjwal Bhargava, Shruti Bhosale, et al. 2023b. Llama 2: Open foundation and fine-tuned chat models. *arXiv preprint arXiv:2307.09288*.
- Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Fei Xia, Ed Chi, Quoc V Le, Denny Zhou, et al. 2022. Chain-of-thought prompting elicits reasoning in large language models. *Advances in Neural Information Processing Systems*, 35:24824–24837.
- T. Yang, J. Deng, X. Quan, and Q. Wang. 2023. Orders are unwanted: Dynamic deep graph convolutional network for personality detection. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 37, pages 13896–13904.
- Susan Zhang, Stephen Roller, Naman Goyal, Mikel Artetxe, Moya Chen, Shuohui Chen, Christopher Dewan, Mona Diab, Xian Li, Xi Victoria Lin, et al. 2022. Opt: Open pre-trained transformer language models. *arXiv preprint arXiv:2205.01068*.

A Appendix

A.1 Kaggle MBTI Dataset

The public Kaggle MBTI dataset is comprised of 8675 entries, each corresponding to an individual’s MBTI type. In particular, every data entry includes excerpts from the last 50 posts or comments made by the individual on the PersonalityCafe forum, separated by the special character “||”. Two example data entries are provided in Table 8. Note that this dataset is anonymous upon creation.

A.2 Computational Resources

In this work, all the experiments are done with Nvidia RTX A6000 GPU. In particular, we hosted Llama2-7B, Llama2-13B, and Llama2-70B on 1, 2, and 4 GPU cards, respectively, where Llama2-70B was hosted in mixed precisions. For the ChatGPT experiment, we directly use OpenAI APIs⁵.

A.3 Finetuning Experiments Details

In this section, we report the experimental details and complete results for both binary and 16-class models to reproduce Table 1. For baseline models, as the models are trained on the same dataset, we kindly refer the audience to the original papers for the performance of their models (Yang et al., 2023; Mehta et al., 2020; Tang et al., 2023). In particular, in the fine-tuning experiments, the original MBTI dataset is divided into training, evaluation, and heldout testing subsets in an 81:9:10 ratio. Furthermore, for the binary model, the model for each trait shared the same set of hyperparameters.

Metrics	E/I	N/S	T/F	P/J	Average
Accuracy (%)	94.01	96.08	92.97	89.98	93.26
F1 (%)	90.57	91.48	92.91	89.38	91.09
Precision (%)	93.16	93.44	93.13	89.41	92.29
Recall (%)	88.57	89.76	92.78	89.36	90.10

Table 6: Performance of personality detection models on each of the four MBTI dimensions for the binary model.

A.4 Prompting Pipeline & Downloaded Topics

In principle, we want to feed LLMs with real-world topics, and tweets and collect their responses and attitudes toward them. We analyzed daily news to obtain a list of real-world topics. In terms of the topics for tweets, we downloaded 5000 tweets containing the following 10 different topics: Bitcoin, NFL, Music, Oscars, Travel, Fashion, Food,

⁵<https://platform.openai.com/>

Metrics	E/I	N/S	T/F	P/J	Average
Accuracy (%)	93.09	96.31	93.55	90.90	93.46
F1 (%)	89.49	92.10	93.50	90.29	91.35
Precision (%)	90.33	93.37	93.65	90.64	92.00
Recall (%)	88.71	90.93	93.40	89.99	90.76

Table 7: Performance of personality detection models on each of the four MBTI dimensions for the 16-class model. Note that the numbers are calculated by extracting each dimension from the predictions.

Fitness, Gaming, and Technology, where they are nearly evenly distributed.

The templates that we used to generate Twitter posts and comments are provided in Table 3. As for the parameters for generation, we utilized a nucleus sampling (Holtzman et al., 2019) with the temperature and top_p set to 0.2 and 0.95, respectively. In addition, the maximum length of generated contents is set to 200. The same configuration is used for all LLMs that we investigated.

A.5 Validation Experiment on Human Counterpart

In the validation experiment, we downloaded the public Twitter posts and comments made by 8 celebrities from different domains and use the same personality evaluation procedure as what we have done for LLMs. Specifically, we created 100 samples with 50 tweets in each for every celebrity and repeated the external evaluation method for 100 times. The obtained MBTI distributions and the most frequent MBTI types are reported in Figure 5 and Table 5, respectively, where the results echo the fact in psychology literature where personality is a consistent characteristic for humans (Association, 2023; Green et al., 2019).

Type	Contents (50 posts or comments on PersonalityCafe)
ENTP	I am finding the lack of me in these posts very alarming This + Lack of Balance and Hand-Eye Coordination. ...
INTJ	2% still means about 1/50 people. I've probably seen 1-2 others today. I never understood fascination by virtue of rarity. I collect shoes. I do so because I like status and nothing communicates such a thing as much as a pair of Jordans...
...	...

Table 8: Two example data entries of the public Kaggle MBTI dataset.



Figure 5: MBTI distribution of 8 selected celebrities using external evaluation method for 100 times. In the figure, the first row is the assessment results on the generated posts dataset (P), whereas the second row provides the results on the comments dataset (C). In addition, personality types that appear less than 3 times are merged together to form the class “Others”.