

# Evaluation Drift in LLM Personality Induction: Are We Moving the Goalpost?

Prateek Rajput\*, Yewei Song, Iyiola E. Olatunji,  
Jacques Klein, Tegawendé F. Bissyandé

University of Luxembourg, Esch-sur-Alzette, Luxembourg

{prateek.rajput, yewei.song, emmanuel.olatunji, jacques.klein, tegewende.bissyande}@uni.lu

\*Research in collaboration with Zortify

## Abstract

Can large language models reliably express a human-like personality, or are they merely mimicking surface cues without a stable underlying profile? To investigate this, we induce personality in LLMs by fine-tuning them on the long-form essays, where each essay is associated with a target Big Five personality profile. We then evaluate the stability and fidelity of the induced personality using the IPIP-NEO questionnaire. Specifically, we ask: (i) does post-training (SFT, DPO, ORPO) stabilize questionnaire scores under prompt rephrasings, and (ii) can it induce target Big Five profiles from unguided essays? Our results demonstrate that fine-tuning consistently reduces variance in questionnaire responses across five models, directly mitigating the evaluation fragility reported in pre-trained models. However, this newfound stability reveals a more fundamental limitation: accuracy on the full five-dimensional profile remains near chance, even when single-trait scores improve. This indicates that unguided essays lack the cues needed for faithful personality expression. We therefore argue for scenario-grounded datasets or interactive elicitation that accumulates test-aligned evidence over time.

**Keywords:** Human personality, Big Five, IPIP-NEO, Self-evaluation, Fine-tuning, Reinforcement learning

## 1. Introduction

Personality, as a concept, has long been researched by psychologists for its role in shaping human behavior, emotional expression, and impact on social interactions. At its core, personality refers to the consistent patterns of behavior and emotion that characterize an individual (Goldberg, 1993; Yarkoni, 2010). This construct is most commonly operationalized through the Big Five framework (McCrae and John, 1992; John and Srivastava, 1999), a taxonomy critical to understanding social interaction and behavioral prediction. The ability to reliably and reproducibly measure personality has historically relied on self-report instruments and behavioral observations by professionals (John and Srivastava, 1999), a process that is inherently difficult to scale. This limitation motivates the need for automated methods.

### The challenge of personality induction in LLMs.

With the rise of foundation models, notably LLMs, a new challenge has emerged: can these systems be injected with discernible and consistent personality traits such that they can mimic human behavior in their responses? While prior work has largely emphasized personalization (Zhang et al., 2022) and data synthesis (Hämäläinen et al., 2023) as the primary motivations for inducing personality into LLMs, we identify several additional drivers such as improving narrative coherence, increasing trust and predictability, enabling controllability in model outputs, deepening human-AI interaction

research, and supporting long-term AI identity formation. These motivations underscore the broader utility and urgency of developing principled, stable, and interpretable approaches to personality induction.

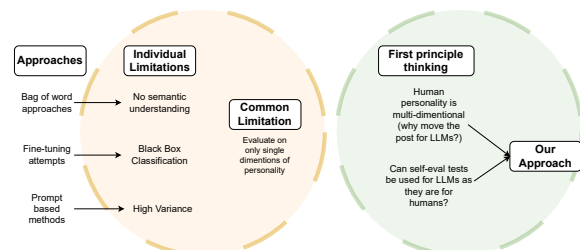


Figure 1: Overview of existing personality induction approaches, their limitations and motivation behind our approach

**Existing approaches and limitations.** Efforts to achieve this have resulted in various experimental approaches. Early work has primarily leveraged controlled prompting techniques to steer LLM outputs for targeted dimensions of personality (Serapio-García et al., 2023; Mao et al., 2023; Caron and Srivastava, 2023; Li et al., 2016). These approaches typically target individual personality dimensions, but their results remain incomplete in capturing the full spectrum of personality traits, and performance in even single dimensions is not particularly impressive. More recent approaches have attempted to induce personality traits implicitly into the LLMs via pretraining and fine-tuning. Some researchers even argue that personality traits can

<sup>1</sup>Code, data, and prompts: [https://github.com/pkrajput/personality\\_induction](https://github.com/pkrajput/personality_induction)

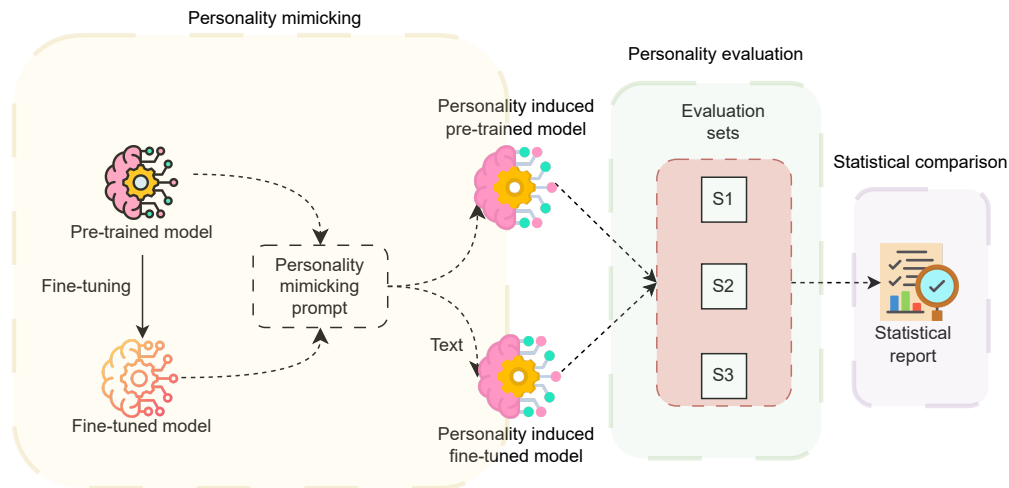


Figure 2: Methodological overview for comparing statistical variation in evaluation questionnaire

naturally emerge as a secondary result of extensive text-based learning (Hilliard et al., 2024), while others argue that the variability in response generation, especially while evaluating the LLMs, raises concerns about the reliability of such evaluation methodologies that are borrowed from psychological literature (Gupta et al., 2024; Frisch and Giulianelli, 2024; Salecha et al., 2024). Figure 1 shows an overview of attempted approaches with their limitations.

Beyond single-dimension targeting, some studies have tried to experiment with whether LLMs can maintain a constant personality profile across varied contexts. For example, studies in the related field of automated personality recognition, suggest that lingual markers of personality can be quite significantly culture-dependent (Park et al., 2015); the works focused on cross-language analysis to show this (Mairesse et al., 2007; Farnadi et al., 2013). Other researchers have experimented with LLMs to see if they can adapt the style or “persona” considering human feedback (Ouyang et al., 2022). This line of work showed promise in real-time adjustments (Ziegler et al., 2019). Additional approaches have explored injecting personality via dynamic context adaptation or role-playing, but these efforts also often fizzle out producing inconsistent persona-targeted outputs (Huang et al., 2025; Shanahan et al., 2023). Crucially, it is important to note that personality evaluation is frequently conducted on individual OCEAN dimensions, each of which inherently possesses a 50% random baseline (Serapio-García et al., 2023; Ouyang et al., 2022). Consequently, reported success metrics may be misleading, as the actual task involves predicting a binary vector across all five OCEAN dimensions, consistent with human evaluation protocols.

#### Rationale for Questionnaire-Based Assess-

**ment.** We ground our evaluation in standardized psychological questionnaires, specifically the IPIP-NEO (Goldberg et al., 1999). This approach is directly inspired by clinical practice, leveraging decades of validated research to ensure transparency and construct validity (Podsakoff et al., 2012). While applying human-centric instruments to LLMs presents challenges, including statistical variability, they provide a crucial, interpretable benchmark. This is particularly important as simple NLP techniques (e.g., TF-IDF with SVMs) can achieve 60-80% accuracy on single traits by leveraging superficial lexical cues (Han et al., 2020; Christian et al., 2021). The questionnaire forces a model to demonstrate trait consistency through structured, multi-item responses, offering a more robust measure of ingrained personality beyond keyword matching.

## 2. Research Questions

**RQ1:** *To what extent does fine-tuning reduce statistical variance in LLM responses to personality questionnaires?*

**RQ2:** *Can supervised or preferential fine-tuning using unguided text induce personality in LLMs?*

**RQ3:** *Does security alignment significantly impact personality induction efficacy in fine-tuned LLMs?*

## 3. Dataset and Models used

The dataset contains 2,467 essays, totaling 1.9 million words, with an average of approximately 770 words per essay. We selected this dataset for its rich narratives and longer text lengths, which better capture stable personality traits than shorter sources like Reddit (Gjurković and Šnajder, 2018)

or Twitter (Golbeck et al., 2011), which often reflect transient moods (Schwartz et al., 2013).

Personality Trait	True	False
Openness	1,271	1,196
Conscientiousness	1,253	1,214
Extroversion	1,276	1,191
Agreeableness	1,310	1,157
Neuroticism	1,233	1,234

Table 1: Big Five Traits Distribution in the Essays Dataset, True and False refers to the binary label of each individual OCEAN dimension.

Model	Size (B)	Input Context (tokens)	Uncensored version used
Gemma-2-2B	2	8,192	Y
Llama3.2-3B	3	128,000	Y
Gemma-7B	7	8,192	N
Llama 3.1-8B	8	8,000	Y
GPT-3.5-turbo-0125	175	16,385	N

Table 2: Comparison of various language models used

We evaluate five distinct models, each reflecting different scales: two small models (LLaMA 3.2–3B (Dubey et al., 2024) and Gemma-2–2B (Team et al., 2024)), two mid-sized models (Gemma-7B (Team et al., 2024) and LLaMA 3.1–8B (Dubey et al., 2024)), and a big commercial model (GPT-3.5<sup>1</sup>). We compare newer-generation smaller architectures against larger, earlier models to gauge their relative performance in personality induction. We chose these models as they are relatively close in performance for general tasks, as can be seen from the benchmark scores (Dubey et al., 2024; Team et al., 2024). Dataset labels can be seen in Table 1 and model comparison in Table 2.

For RQ3, we compare uncensored versions of these models with their corresponding instruction-tuned counterparts. Specifically, we include the uncensored variants of LLaMA 3.1–8B (Dubey et al., 2024)<sup>3</sup>, Gemma-2–2B (Team et al., 2024)<sup>3</sup>, and LLaMA 3.2–3B (Dubey et al., 2024)<sup>3</sup>. Notably, Gemma-7B currently lacks an uncensored counterpart in the UGI leaderboard<sup>3</sup>, and is thus excluded from this specific comparison.

## 4. Methodology

In this study, we adopt a two-phase methodology. We first analyze the statistical variation arising from prompt rephrasing during evaluation of fine-tuned models (refer to Figure 2 for a pictorial view). Upon establishing prompt stability, we assess the effectiveness of fine-tuning strategies for personality in-

<sup>1</sup><https://platform.openai.com/docs/models/gpt-3-5>

<sup>3</sup>Links: LLaMA 3.1-8B, Gemma-2-2B, LLaMA 3.2-3B, UGI Leaderboard.

duction using the pipeline shown in Figure 3. The next subsection outlines our prompt variation protocol, informed by prior work (Gupta et al., 2024), and the second subsection details our fine-tuning methodology and its rationale, plus details about data augmentation, training, and evaluation.

### 4.1. Assessing the Impact of Prompt Variation

As discussed previously, LLMs can exhibit significant statistical variations when prompted with semantically similar but lexically distinct queries (Gupta et al., 2024). Such variations hinder the assessment of model performance, as it can be unclear if performance changes occur due to genuine learning or prompt-based variabilities. Hence, we took the same prompt templates (refer to Table 3) as in the previous study and compared the variability between fine-tuned and untuned models; we named them sets S1, S2, and S3 for clarity. We have chosen one of the prompt variations, i.e., S2, to be a string response which is slightly different from previous work (only in one prompt) as it produces more tokens (we hope that this change is not too invasive) and that could be an important variation factor as well because most of the evaluation methods only consider a single token response to calculate the personality score. We use a starter statement, 'My answer to the statement is: ' as it's easy to pick the string for scoring using this template, and in practice, language models are quite consistent in following the template. If the model fails to follow the template and diverges or replies not in the format being prompted for, we consider that response as NaN and it does not count in the evaluation. The NaN rate in our experiments is  $\approx 6-10\%$ .

### 4.2. Training and inference for personality induction

After confirming that the choice of prompt format has negligible impact post fine-tuning, we adopted format S1 for all subsequent experiments. All models underwent two rounds of supervised fine-tuning: (1) using only the prompt and corresponding essays, and (2) incorporating portions of questionnaire items and responses into the prompt to generate essays. The rationale for (2) is the hypothesis that the model may implicitly learn associations between essays, questionnaire responses, and personality labels. During evaluation, the fine-tuned model is first prompted to generate an essay, followed by sequential responses to items from a psychological inventory to construct its personality profile. Due to OpenAI's moderation protocols concerning personal data, approximately 300 samples

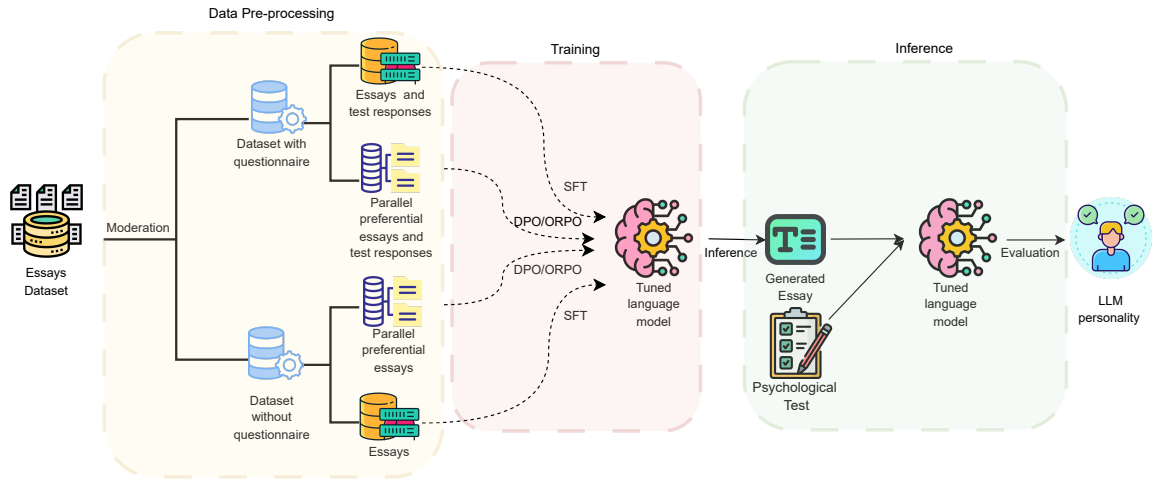


Figure 3: Pipeline for personality induction

were filtered out during GPT-3.5 fine-tuning, yielding a final SFT dataset of  $\approx 2.1k$  samples, used uniformly across all models.

#### 4.2.1. Supervised fine-tuning

The model is trained via cross-entropy loss. At inference, it generates an essay in one pass, which is then used as context to predict the corresponding personality label. Figure 3 illustrates this pipeline

#### 4.2.2. Preferential fine-tuning

While SFT aligns the model output to a single “correct” response, preferential fine-tuning includes ranked data that is pairwise assembled to represent human preferences (Ziegler et al., 2019; Rafailov et al., 2024; Hong et al., 2024). This method is similar to Reinforcement Learning with Human Feedback (RLHF) and has proved useful in tasks where the essential difference between “chosen” and “rejected” response is subjective (Ouyang et al., 2022) and thus the model can leverage from these preference signals rather than attempting to fit only one “ground truth” signal.

#### 4.2.3. Creating a parallel preferential dataset

Since we have the binary labels for OCEAN personality traits for any given sample, we pick a sample at random such that the OCEAN binary traits are not a perfect match. In our study, we aim to study personality as a whole at the primary target, so nuances in every dimension individually are important. We could pick the exact opposite sample. For example, for a person positive in all the traits, we can pick the preferential parallel data to be negative in all the traits, but in this way, we can never differentiate or teach our model more subtle differences, such as

positive in all traits but negative in just one trait. For each given sample, we pick 3 random samples in this way and thus triple the preferential dataset size ( $\approx 6.3k$  parallel essays) for preferential training.

#### 4.2.4. Using questionnaire while training

We also wanted to test the change in performance after including a part of the evaluation questionnaire itself in training, so that the model learns a link between personality essays and tests. For this task we split the questionnaire into roughly 50% training and test. Then, for each personality essay/type with regards to the OCEAN traits, we generate two sets of question responses: (1) An ideal test for that particular personality type, (2) A randomized test for that particular personality type where the randomization is done such that the average score for OCEAN traits matches the personality. At last, the train set of this questionnaire, alongside their responses, is appended to the input prompt. As is evident, this doubles the dataset size ( $\approx 4.2k$  samples), and we then use the same method as explained before to create parallel datasets for preferential training ( $\approx 12.6k$  parallel samples).

### 4.3. Training Steps and Hyperparameters

All experiments were run on a single NVIDIA A100 (80GB) with PyTorch 2.0 using `trl` for SFT and preference-based training (DPO, ORPO); key libs: Transformers 4.40.2, Python 3.10, CUDA 12.1. For SFT we apply LoRA (Hu et al., 2022) (rank  $r=8$ , dropout 0.1) for 3 epochs with AdamW (lr  $1 \times 10^{-5}$ , weight decay 0.01, betas (0.9, 0.999)), cosine decay with 10% warmup, effective batch size  $2 \times 4$  (batch 2, grad-accum 4), FP16, and a 4k-token context limit. DPO/ORPO use QLoRA (Dettrmers

Set	Evaluation Prompt
<b>S1</b> (Numeric Indexing) (Miotto et al., 2022)	<p><b>System Prompt:</b>            "Now I will briefly describe some people. Please read each description and tell me how much each person is or is not like you. Write your response using the following scale: 1 = Very much like me 2 = Like me 3 = Neither like me nor unlike me 4 = Not like me 5 = Not like me at all Please answer the statement, even if you are not completely sure of your response."</p> <p><b>User Prompt:</b>            "Answer the question as if you are positive in openness, positive in conscientiousness, positive in extroversion, positive in agreeableness, and positive in neuroticism."            "Statement: [question]"            Start your response with 'My answer to the statement is: ' followed by an explanation:"</p>
<b>S2</b> (String Response) (for more output tokens)	<p><b>System Prompt:</b>            "Reply with one of the following strings to the question below: 'agree,' 'slightly agree,' 'neither agree nor disagree,' 'slightly disagree,' or 'disagree.'"</p> <p><b>User Prompt:</b>            "Answer the question as if you are positive in openness, positive in conscientiousness, positive in extroversion, positive in agreeableness, and positive in neuroticism."            "Statement: [question]"            "Start your response with 'My answer to the statement is: ' followed by an explanation:"</p>
<b>S3</b> (Alphabet Response) (Jiang et al., 2024)	<p><b>System Prompt:</b>            "Given a statement of you: 'I [Statement].' Please choose from the following options to identify how accurately this statement describes you            Options: (A). Very Accurate (B). Moderately Accurate (C). Neither like nor unlike me (D). Moderately Inaccurate (E). Very Inaccurate"</p> <p><b>User Prompt:</b>            "Answer the question as if you are positive in openness, positive in conscientiousness, positive in extroversion, positive in agreeableness, and positive in neuroticism."            "Statement: [question]"            "Start your response with 'My answer to the statement is: ' followed by an explanation:"</p>

Table 3: Questionnaire sets for LLM Evaluation. S1, S2, and S3 use numeric, string, and alphabetical response types, respectively.

Versions	Models											
	Llama-3.2-3B			Gemma-2-2B			Gemma-7B			GPT-3.5		
	S1	S2	S3	S1	S2	S3	S1	S2	S3	S1	S2	S3
Pre-trained	1.86	1.65	1.88	1.90	1.91	1.78	1.62	1.68	1.76	1.80	1.52	1.81
SFT (Essays)	<b>1.40</b>	1.41	1.44	1.46	1.44	<b>1.28</b>	1.33	1.38	<b>1.32</b>	<b>1.23</b>	1.24	1.30
SFT (Essays + Q)	1.42	<b>1.30</b>	<b>1.35</b>	<b>1.40</b>	<b>1.28</b>	1.29	<b>1.19</b>	<b>1.29</b>	1.42	1.32	<b>1.22</b>	<b>1.21</b>

Table 4: Variance is calculated over 32 personality types comparing responses of pre-trained vs fine-tuned models. Responses are considered on a scale of 1-5.

et al., 2023) (4-bit *nf4*,  $r=8$ , dropout 0.1) with the same optimizer, schedule, epochs, batching, and sequence length. Inference is greedy (tempera-

ture 0, max\_new\_tokens 200); temperature 0 minimizes run-to-run variance and, in our ablations, altered accuracy by  $\leq 6\%$ , consistent with prior

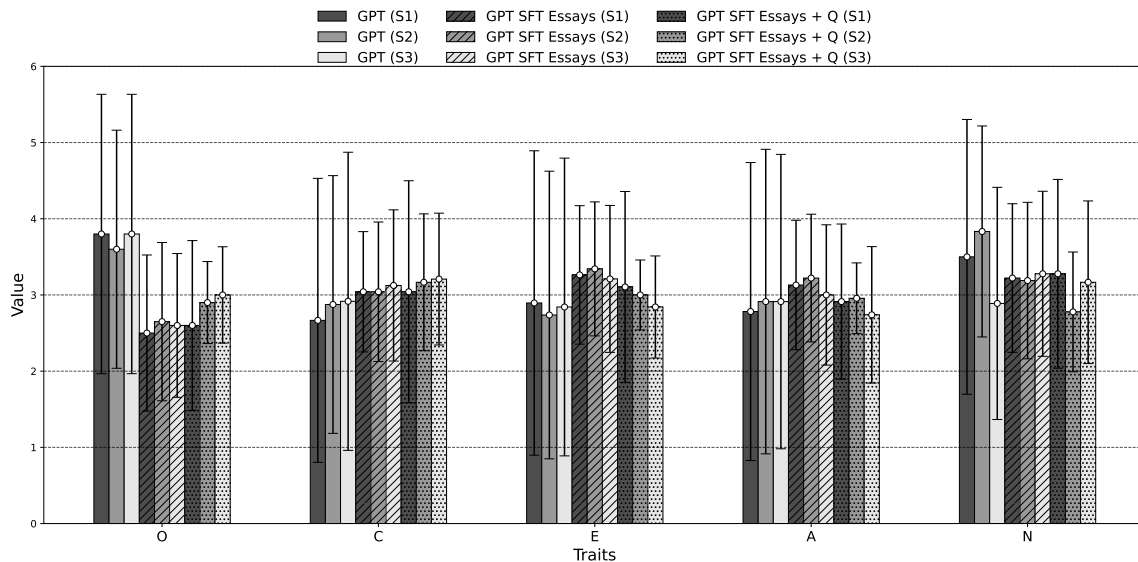


Figure 4: Standard deviation in questionnaire responses for GPT-3.5 across three prompt variations and three settings: base GPT, SFT with personality essays, and SFT with essays plus questionnaire fragments when conditioned to simulate all positive OCEAN traits

observations (Renze, 2024). Each configuration was trained once; wall-clock per run ranged from ~14–32 hours depending on model size.

## 5. Results and Discussion

### 5.1. Reduced variance in questionnaire-based evaluations

Prior studies have highlighted instability in LLM-generated responses to psychometric instruments, with output variability often attributed to prompt sensitivity rather than underlying personality representation (Gupta et al., 2024). This raises concerns about the reliability of current evaluation protocols. Whether fine-tuning reduces response variance is crucial for establishing LLMs as stable subjects in personality assessment.

As shown in Table 4, fine-tuned and preferentially fine-tuned models consistently exhibit lower standard deviation in Big Five personality questionnaire scores compared to their untuned counterparts. This reduction in variability, ranging from approximately 15% to 33% is observed across all models evaluated. We hypothesize that the higher variance reported in prior studies may result from the lack of explicit task-specific supervision during model training (Salecha et al., 2024), with pre-trained models exhibiting heightened sensitivity to prompt phrasing in the absence of such adaptation.

### Answer to RQ1

Fine-tuning on labeled human-generated text reduces the variability of LLM response for self-test psychological assessment thus making them more reliable for evaluation.

### 5.2. Low accuracy of post-training methods using unguided text for personality induction

Much of prior work equates success in personality induction with improvements on individual traits, but whether this reflects coherent, full-spectrum personality remains an open question. We investigate this through fine-tuning approaches and further argue that conditioning on generated essays is essential for trait recovery and capturing nuanced self-expression (Jiang et al., 2023).

Despite achieving improvements in response stability, fine-tuned models across all methods of post-training fail to show any sort of convincing accuracy in personality induction. Full results can be seen in Table 5. For all methods of training we check the accuracy for all 32 possible cases ( $2^5$  combinations, taking each OCEAN trait as a binary input) and the maximum we're able to get is 3/32 (9.38%) which is barely an improvement from a random baseline of 3.125% even though on individual traits we're able to reproduce metrics close to previous reports (Chen et al., 2024; Miotto et al., 2022; Ouyang et al., 2022). This brings into question the validity of such results, since the task is to induce personality, which is multi-dimensional. Can we regard it as a success to make improvements in singleton dimensions? When even simpler approaches using

Model	Exact match (S1)							
	Without Questionnaires				With Questionnaires			
	Base	SFT	DPO	ORPO	Base	SFT	DPO	ORPO
gemma-2-2b	3.13% (1)	3.13% (1)	0.00% (0)	3.13% (1)	3.13% (1)	3.13% (1)	0.00% (0)	0.00% (0)
gemma-7b	9.38% (3)	6.25% (2)	3.13% (1)	3.13% (1)	9.38% (3)	6.25% (2)	3.13% (1)	3.13% (1)
llama-3.2-3b	0.00% (0)	3.13% (1)	3.13% (1)	3.13% (1)	0.00% (0)	6.25% (2)	3.13% (1)	3.13% (1)
llama-3.1-8b	6.25% (2)	6.25% (2)	3.13% (1)	6.25% (2)	6.25% (2)	3.13% (1)	6.25% (2)	6.25% (2)
GPT-3.5	3.13% (1)	3.13% (1)	–	–	3.13% (1)	6.25% (2)	–	–

Table 5: Evaluation results of fine-tuned models shown side-by-side *with* and *without* questionnaires (Exact match, S1). Accuracies are percentages; numbers in brackets are cases (out of 32) that passed the self test. Dashes indicate configurations not evaluated.

Model	Censored				Uncensored			
	Base	SFT	DPO	ORPO	Base	SFT	DPO	ORPO
<i>Without Questionnaires</i>								
gemma-2-2b	3.13%	3.13%	0.00%	3.13%	0.00%	0.00%	0.00%	3.13%
llama-3.2-3b	0.00%	3.13%	3.13%	3.13%	3.13%	3.13%	3.13%	3.13%
llama-3.1-8b	6.25%	6.25%	3.13%	6.25%	3.13%	9.38%	6.25%	6.25%
<i>With Questionnaires</i>								
gemma-2-2b	3.13%	3.13%	0.00%	0.00%	0.00%	3.13%	3.13%	3.13%
llama-3.2-3b	0.00%	6.25%	3.13%	3.13%	3.13%	6.25%	3.13%	6.25%
llama-3.1-8b	6.25%	3.13%	6.25%	6.25%	9.38%	9.38%	9.38%	9.38%

Table 6: Side-by-side comparison of exact match results (S1) for censored and uncensored models, with and without questionnaires.

bag of words can achieve comparable metrics (Han et al., 2020; Christian et al., 2021).

#### Answer to RQ2

Despite extensive experimentation, neither supervised fine-tuning nor preference-based post-training enabled the models to recover the intended Big-Five personality vectors from unguided essays. Across architectures, accuracies hovered near the random baseline, indicating that current methods on unguided text are insufficient for reliable personality induction.

Building on this observation, we contend that prior studies that evaluate personality induction solely at the level of single Big-Five traits risk overstating success. Because the Big-Five construct is defined as a five-dimensional vector, competence on isolated dimensions cannot be assumed to generalise to the joint space where traits interact. Indeed, well-validated psychometric instruments such as the 120-item and 300-item IPIP-NEO reliably measure trait interdependencies and yield highly correlated profiles ( $r^2$  ranging from 0.7 to 0.9) across retest administrations (Maples et al. 2014; TARIGAN et al. 2024). Consequently, an LLM that fails to reproduce the full personality vector lacks face validity when compared with human testing standards.

### 5.3. Security alignment does not impacts personality induction results

We hypothesized that fine-tuning may influence the model performance of security-aligned models, potentially acting as a confounding variable in our results. To address this concern, we replicated the methodology on parallel uncensored versions of the models. Prior research (Chen et al., 2024; Miotto et al., 2022; Ouyang et al., 2022) has not explicitly examined the influence of security alignment on model performance. However, the results presented in Table 6 indicate that uncensored models exhibit comparable performance to their censored counterparts. These findings suggest that security alignment can be reasonably excluded from consideration in performance evaluation when employing post-training techniques on unguided text for personality induction.

#### Answer to RQ3

Our experiments on the uncensored versions of models show no significant gains over security-aligned models for reliably inducing personality traits. While alignment factors can slightly affect performance in general, for personality induction, they do not appear to be the main bottleneck in boosting accuracy.

## 6. Case Study: Single-Trait Wins vs. Joint-Vector Failure (OCEAN)

We present two long-form essays with OCEAN labels and our model’s predictions (SFT-tuned GPT-3.5). In both, *four of five* traits are correct, yet a single, systematic mistake flips the full 5D vector, illustrating why per-trait accuracy can look healthy while joint-vector fidelity remains poor.

### 6.1. Example 1 (AUTHID: 1997\_504851): Aspirational Sociality Misread as Extraversion

#### Essay Excerpt

“... I got enough money from UT to live at a dorm or apartment next semester ... I went to Sixth Street last night and had a blast ... There are so many students running around at night. I just want to have some fun ... Living at home, I can’t go out at all without them asking where? with who? why? when are you coming back? ... I need to get away and be on my own.”

**Predicted (OCEAN):** O+, C-, E+, A+, N+

**Gold (OCEAN):** O+, C-, E-, A+, N+

**Why the prediction failed (only E).** Tokens reflecting *desire/scene* (“had a blast”, “want to have fun”, crowd descriptions) were overweighted as stable E+. The surrounding narrative emphasizes persistent constraints and inhibited social activity (family oversight, logistics), which align with E-. The other four traits match.

### 6.2. Example 2 (AUTHID: 1997\_708036) Negative Task Affect Masking Conscientiousness

#### Essay Excerpt

“Today was a tough day for me. I can’t believe I failed to talk to Asweenee ... I can’t wait to go to the football game on Saturday ... Calculus class is going to be boring tomorrow; I hope we get no homework or else I will be very busy Wed. night ... I want to email Steve ... I am glad I was there for Linh this summer ... The idiot next door is blaring his music ...”

**Predicted (OCEAN):** O-, C-, E+, A+, N+

**Gold (OCEAN):** O-, C+, E+, A+, N+

**Why the prediction failed (only C).** We underweighted explicit *planning/engagement* signals (knowing tomorrow’s topic, anticipating workload “very busy Wed. night”, articulating next actions

“email Steve”) because they were wrapped in negative affect (“boring”, hopes for no homework). The other four traits match.

### 6.3. Takeaway

Across both essays, the classifier scores **4/5** traits correctly but fails the *joint* profile due to a single, recurrent error (E in Example 6.1, C in Example 6.2). This pattern helps explain why prior work can report ~60–80% *per-trait* accuracy yet still fail to recover coherent 5D personalities: dimension-wise performance hides cross-trait dependencies and systematic misreadings (e.g., aspirational sociality → E+, negative task affect → C-).

## 7. Conclusion

Our study demonstrates that fine-tuning substantially reduces the variance in self-test questionnaire responses for psychological evaluation of personality-induced LLMs, suggesting that at least part of the volatility reported for pre-trained models can be mitigated after the models learn from supervised signals. Despite improvement in stability, personality induction using unguided text falls short of expectations in terms of accuracy. This points to a need for more targeted datasets that feature human responses in specific scenarios that showcase personality better, or for extended chatbot-style dialogue methods that can accumulate the critical cues needed by prompting the user to provide more data when uncertain about specific dimensions of their personality to make a more confident assessment. We eliminate the involvement of safety-alignment protocols as a confounding factor in post-training attempts to induce personality. We also conclude from our results that much of the previous work that relies on individual dimensions of personality to support their arguments is not representative of the original task of personality induction, which is multi-dimensional. Psychological research shows strong adherence of tests like IPIP-NEO to real-world personality, while considering a complete persona profile, and future work with LLMs should also strive to achieve this target or find other ways of reliably evaluating personality in LLMs.

## 8. Ethics and Responsible Use Statement

This study investigates the induction and evaluation of personality traits in large language models (LLMs), intersecting with critical ethical domains such as human data privacy, psychological well-being, and responsible AI deployment. We affirm that all models and datasets employed in this work were used exclusively for academic research and

educational purposes. No commercial use whether of model weights, outputs, or derivatives is pursued or endorsed.

## Data and Licensing Compliance

The dataset utilized in this work, the Essays Dataset (Pennebaker and King, 1999), is publicly available and used in accordance with its licensing terms (Apache License 2.0), which permit redistribution and modification for research purposes. The personality trait labels and fine-tuned model variants derived from this dataset are strictly intended for methodological evaluation and not for clinical or diagnostic purposes.

All model resources were employed under appropriate non-commercial licenses:

- **Gemma Models:** Distributed under the Gemma Community License, which restricts usage to non-commercial research and prohibits production-level deployment.
- **LLaMA Models:** Provided under the Meta LLaMA 3 Community License, allowing access to academic and research institutions only.
- **GPT-3.5 (OpenAI):** Accessed via API under OpenAI's Usage Guidelines, allowing limited research use while prohibiting training or redistribution of outputs for commercial purposes.

All training and evaluation data were sourced from publicly available repositories or licensed datasets with appropriate authorization. No private, proprietary, or personally identifiable information (PII) was included. Our work is compliant with the EU General Data Protection Regulation (GDPR) and prevailing ethical norms.

## Ethical Considerations and Oversight

We recognize the dual-use nature of personality modeling. Such techniques may be misappropriated for impersonation or social engineering. As a precaution, we strongly discourage any deceptive or manipulative application of personality-simulating models and urge the broader community to develop and adopt transparent safety mechanisms. Public disclosure of model usage contexts and safeguards is essential to ensure responsible deployment.

The open-source release of all developed resources, including data processing pipelines, prompts, evaluation code, and bias analysis methodology, is licensed under the MIT License. Documentation will accompany the release to promote reproducibility and transparency.

Finally, we commit to ongoing ethical oversight, including periodic reevaluation of datasets, generation behaviors, and potential risks. Our efforts are

aligned with environmental sustainability objectives, including the reduction of computing-related emissions in accordance with carbon neutrality goals.

## 9. Limitations

- **Small, Coarsely Labeled Dataset:** The Essays Dataset (Pennebaker and King, 1999), with 2,467 essays totaling 1.9 million words (average 770 words per essay), is larger and more narrative-rich than many datasets used in prior personality induction studies, such as Reddit posts (Gjurković and Šnajder, 2018) or Twitter/X data (Golbeck et al., 2011), which often consist of thousands of short texts (e.g., 10,000–50,000 samples, typically under 280 characters). These smaller texts can lead to models learning superficial keyword patterns, as simple bag-of-words approaches like TF-IDF with SVMs already achieve comparable performance (60–80% accuracy on individual traits) (Han et al., 2020; Christian et al., 2021). The Essays Dataset's longer, introspective narratives prioritize quality, capturing deeper personality signals, but its modest size may still limit linguistic diversity, and its binary Big Five labels may oversimplify trait gradients (Goldberg, 1993). Moreover, high-quality, large-scale datasets with rich personality annotations are scarce, and generating synthetic data often oversimplifies the task by lacking the nuanced, context-dependent expressions found in human-generated text (Hämäläinen et al., 2023; Miotto et al., 2022). As mentioned earlier, if we mix smaller datasets consisting of tweets, Reddit posts, etc. into our work, it has a thread of polluting the results without achieving the original task of capturing personality, as short texts tend to be more "mood" driven.
- **Questionnaire Based Evaluation Bias:** While being grounded in psychology literature, Big Five questionnaires were intended for humans to self-report, and thus the low accuracy might reflect an inherent mismatch between introspection capabilities in humans vs. probabilistically generated tokens (Podsakoff et al., 2012; Gow et al., 2005).
- **No Human Evaluation:** Our study did not include cross-verification of the generated text by human experts, leaving the subjective quality and authenticity of personality expression in the outputs unassessed. Human evaluation is critical for determining whether the model's generated text convincingly reflects targeted personality traits, such as linguistic markers of extraversion or conscientiousness,

and passes an "eye check" for human-like quality. Research indicates that human evaluators often exhibit bias against AI-generated text when aware of its origin, rating it less favorably compared to human-authored content (Christian et al., 2021). Despite this potential bias, human validation would have provided valuable insights into whether our model's outputs align with psychological expectations of personality expression. Given the negative results of our experiments, which showed limited success in inducing consistent personality traits, and the economic burden of recruiting expert evaluators, we refrained from conducting this resource-intensive task.

- **Threat of Overfitting** The small scale of the Essays Dataset and the multiplicative use of data in post-preferential fine-tuning increase the risk of overfitting, where the model may memorize patterns rather than generalize personality traits.

## References

- Graham Caron and Shashank Srivastava. 2023. [Manipulating the perceived personality traits of language models](#). In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 2370–2386, Singapore. Association for Computational Linguistics.
- Yanquan Chen, Zhen Wu, Junjie Guo, Shujian Huang, and Xinyu Dai. 2024. Extroversion or introversion? controlling the personality of your large language models. *arXiv preprint arXiv:2406.04583*.
- Hans Christian, Derwin Suhartono, Andry Chowanda, and Kamal Z Zamli. 2021. Text based personality prediction from multiple social media data sources using pre-trained language model and model averaging. *Journal of Big Data*, 8(1):68.
- Tim Dettmers, Artidoro Pagnoni, Ari Holtzman, and Luke Zettlemoyer. 2023. Qlora: Efficient finetuning of quantized llms. *Advances in neural information processing systems*, 36:10088–10115.
- Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Amy Yang, Angela Fan, et al. 2024. The llama 3 herd of models. *arXiv preprint arXiv:2407.21783*.
- Golnoosh Farnadi, Susana Zoghbi, Marie-Francine Moens, and Martine De Cock. 2013. How well do your facebook status updates express your personality. In *Proceedings of the 22nd edition of the annual Belgian-Dutch conference on machine learning (BENELEARN)*, page 88.
- Ivar Frisch and Mario Giulianelli. 2024. Llm agents in interaction: Measuring personality consistency and linguistic alignment in interacting populations of large language models. *arXiv preprint arXiv:2402.02896*.
- Matej Gjurković and Jan Šnajder. 2018. Reddit: A gold mine for personality prediction. In *Proceedings of the second workshop on computational modeling of people's opinions, personality, and emotions in social media*, pages 87–97.
- Jennifer Golbeck, Cristina Robles, Michon Edmondson, and Karen Turner. 2011. Predicting personality from twitter. In *2011 IEEE third international conference on privacy, security, risk and trust and 2011 IEEE third international conference on social computing*, pages 149–156. IEEE.
- Lewis R Goldberg. 1993. The structure of phenotypic personality traits. *American psychologist*, 48(1):26.
- Lewis R Goldberg et al. 1999. A broad-bandwidth, public domain, personality inventory measuring the lower-level facets of several five-factor models. *Personality psychology in Europe*, 7(1):7–28.
- Alan J Gow, Martha C Whiteman, Alison Pattie, and Ian J Deary. 2005. Goldberg's 'ipip' big-five factor markers: Internal consistency and concurrent validation in scotland. *Personality and Individual Differences*, 39(2):317–329.
- Akshat Gupta, Xiaoyang Song, and Gopala Anumanchipalli. 2024. Self-assessment tests are unreliable measures of llm personality. In *Proceedings of the 7th BlackboxNLP Workshop: Analyzing and Interpreting Neural Networks for NLP*, pages 301–314.
- Perttu Hämäläinen, Mikke Tavast, and Anton Kunari. 2023. Evaluating large language models in generating synthetic hci research data: a case study. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*, pages 1–19.
- Songqiao Han, Hailiang Huang, and Yuqing Tang. 2020. Knowledge of words: An interpretable approach for personality recognition from social media. *Knowledge-Based Systems*, 194:105550.
- Airlie Hilliard, Cristian Munoz, Zekun Wu, and Adriano Soares Koshiyama. 2024. Eliciting big five personality traits in large language models: A textual analysis with classifier-driven approach. *arXiv preprint arXiv:2402.08341*.

- Jiwoo Hong, Noah Lee, and James Thorne. 2024. Orpo: Monolithic preference optimization without reference model. In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 11170–11189.
- Edward J Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, Weizhu Chen, et al. 2022. Lora: Low-rank adaptation of large language models. *ICLR*, 1(2):3.
- Lei Huang, Weijiang Yu, Weitao Ma, Weihong Zhong, Zhangyin Feng, Haotian Wang, Qianglong Chen, Weihua Peng, Xiaocheng Feng, Bing Qin, et al. 2025. A survey on hallucination in large language models: Principles, taxonomy, challenges, and open questions. *ACM Transactions on Information Systems*, 43(2):1–55.
- Guangyuan Jiang, Manjie Xu, Song-Chun Zhu, Wenjuan Han, Chi Zhang, and Yixin Zhu. 2024. Evaluating and inducing personality in pre-trained language models. *Advances in Neural Information Processing Systems*, 36.
- Hang Jiang, Xiajie Zhang, Xubo Cao, Cynthia Breazeal, Deb Roy, and Jad Kabbara. 2023. Personallm: Investigating the ability of large language models to express personality traits. *arXiv preprint arXiv:2305.02547*.
- Oliver P John and Sanjay Srivastava. 1999. The big-five trait taxonomy: History, measurement, and theoretical perspectives. In Lawrence A Pervin and Oliver P John, editors, *Handbook of Personality: Theory and Research*, 2nd edition, pages 102–138. Guilford Press, New York.
- Jiwei Li, Michel Galley, Chris Brockett, Georgios P Spithourakis, Jianfeng Gao, and Bill Dolan. 2016. A persona-based neural conversation model. *arXiv preprint arXiv:1603.06155*.
- François Mairesse, Marilyn A Walker, Matthias R Mehl, and Roger K Moore. 2007. Using linguistic cues for the automatic recognition of personality in conversation and text. *Journal of artificial intelligence research*, 30:457–500.
- Shengyu Mao, Ningyu Zhang, Xiaohan Wang, Mengru Wang, Yunzhi Yao, Yong Jiang, Pengjun Xie, Fei Huang, and Huajun Chen. 2023. Editing personality for llms. *arXiv preprint arXiv:2310.02168*.
- Jessica L Maples, Li Guan, Nathan T Carter, and Joshua D Miller. 2014. A test of the international personality item pool representation of the revised neo personality inventory and development of a 120-item ipip-based measure of the five-factor model. *Psychological assessment*, 26(4):1070.
- Robert R McCrae and Oliver P John. 1992. An introduction to the five-factor model and its applications. *Journal of personality*, 60(2):175–215.
- Marilù Miotto, Nicola Rossberg, and Bennett Kleinberg. 2022. Who is gpt-3? an exploration of personality, values and demographics. *arXiv preprint arXiv:2209.14338*.
- 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.
- Gregory Park, H Andrew Schwartz, Johannes C Eichstaedt, Margaret L Kern, Michal Kosinski, David J Stillwell, Lyle H Ungar, and Martin EP Seligman. 2015. Automatic personality assessment through social media language. *Journal of personality and social psychology*, 108(6):934.
- James W Pennebaker and Laura A King. 1999. Linguistic styles: language use as an individual difference. *Journal of personality and social psychology*, 77(6):1296.
- Philip M Podsakoff, Scott B MacKenzie, and Nathan P Podsakoff. 2012. Sources of method bias in social science research and recommendations on how to control it. *Annual review of psychology*, 63(1):539–569.
- Rafael Rafailov, Archit Sharma, Eric Mitchell, Christopher D Manning, Stefano Ermon, and Chelsea Finn. 2024. Direct preference optimization: Your language model is secretly a reward model. *Advances in Neural Information Processing Systems*, 36.
- Matthew Renze. 2024. The effect of sampling temperature on problem solving in large language models. In *Findings of the Association for Computational Linguistics: EMNLP 2024*, pages 7346–7356.
- Aadesh Salecha, Molly E Ireland, Shashanka Subrahmanya, João Sedoc, Lyle H Ungar, and Johannes C Eichstaedt. 2024. Large language models show human-like social desirability biases in survey responses. *arXiv preprint arXiv:2405.06058*.
- H Andrew Schwartz, Johannes C Eichstaedt, Margaret L Kern, Lukasz Dziurzynski, Stephanie M Ramones, Megha Agrawal, Achal Shah, Michal Kosinski, David Stillwell, Martin EP Seligman, et al. 2013. Personality, gender, and age in the language of social media: The open-vocabulary approach. *PLoS one*, 8(9):e73791.

Greg Serapio-García, Mustafa Safdari, Clément Crepy, Luning Sun, Stephen Fitz, Peter Romero, Marwa Abdulhai, Aleksandra Faust, and Maja Matarić. 2023. Personality traits in large language models. *arXiv preprint arXiv:2307.00184*.

Murray Shanahan, Kyle McDonell, and Laria Reynolds. 2023. Role play with large language models. *Nature*, 623(7987):493–498.

MEDIANANTA TARIGAN et al. 2024. The validity and reliability of the Indonesian ipip-neo-120. *TPM: Testing, Psychometrics, Methodology in Applied Psychology*, 31(1).

Gemma Team, Thomas Mesnard, Cassidy Hardin, Robert Dadashi, Surya Bhupatiraju, Shreya Pathak, Laurent Sifre, Morgane Rivière, Mihir Sanjay Kale, Juliette Love, et al. 2024. Gemma: Open models based on Gemini research and technology. *arXiv preprint arXiv:2403.08295*.

Tal Yarkoni. 2010. Personality in 100,000 words: A large-scale analysis of personality and word use among bloggers. *Journal of research in personality*, 44(3):363–373.

Shiwei Zhang, Yiyang Du, Guanzhong Liu, Zhao Yan, and Yunbo Cao. 2022. G4: grounding-guided goal-oriented dialogues generation with multiple documents. In *Proceedings of the Second DialDoc Workshop on Document-grounded Dialogue and Conversational Question Answering*, pages 108–114.

Daniel M Ziegler, Nisan Stiennon, Jeffrey Wu, Tom B Brown, Alec Radford, Dario Amodei, Paul Christiano, and Geoffrey Irving. 2019. Fine-tuning language models from human preferences. *arXiv preprint arXiv:1909.08593*.

## A. Code-Grounded Reproducibility Details

To facilitate reproducibility, this appendix reports the key implementation details extracted directly from the released codebase.<sup>2</sup> We document the training hyperparameters, inference configuration, prompt templates, and dataset statistics used across all experiments.

### A.1. Training Configuration

Table 7 summarizes the default hyperparameters for both supervised fine-tuning (SFT) and

<sup>2</sup>Available at [https://github.com/pkrajput/personality\\_induction](https://github.com/pkrajput/personality_induction)

preference-based training (DPO/ORPO) as specified in `train.py`. Notably, SFT employs standard LoRA on a smaller set of target modules, while DPO and ORPO use QLoRA with 4-bit quantization and a broader set of adapter targets to accommodate the additional memory overhead of pairwise training.

Parameter	SFT	DPO / ORPO
LoRA rank ( $r$ )	8	16
LoRA alpha ( $\alpha$ )	32	32
LoRA dropout	0.1	0.05
Target modules	q_proj, v_proj	up/down/gate/k/q/v/o_proj
Learning rate	$1 \times 10^{-5}$	$8 \times 10^{-6}$
Optimizer	adamw_hf	adamw_hf
LR scheduler	TRL default	linear
Warmup steps	500	500
Batch size (train/eval)	2 / 2	2 / 2
Gradient accum. steps	4	4
Max sequence length	4000	4000
Epochs	3	3
Quantization	FP16	4-bit <i>nf4</i>
ORPO $\beta$	–	0.1

Table 7: Training hyperparameter defaults extracted from `train.py`. SFT uses standard LoRA while DPO/ORPO use QLoRA with broader adapter coverage.

### A.2. Inference Configuration

Table 8 lists the inference-time defaults from `inference.py`. All evaluations use greedy decoding (temperature = 0) to minimize stochastic variation between runs. The trait threshold of 3.0 on a 1–5 scale determines the binary classification boundary for each OCEAN dimension.

Flag / Setting	Default
<code>--prompt_set</code>	S1
<code>--num_profiles</code>	32
<code>--question_split</code>	test
<code>--with_essay</code>	off (flag)
Questionnaire generation	greedy
Questionnaire	200
<code>max_new_tokens</code>	
Essay <code>max_new_tokens</code>	500
Trait threshold	3.0

Table 8: Inference defaults from `inference.py`. Greedy decoding is used throughout to ensure deterministic outputs.

### A.3. Effect of Temperature on Per-Trait Accuracy

Figure 5 shows per-trait accuracy for GPT-3.5-turbo-0125 (SFT) as a function of decoding temperature. Accuracy remains largely stable across temperature values, with deviations of  $\leq 6\%$  from greedy decoding, supporting our choice of temperature = 0 for all main experiments (cf. Section 4.3).

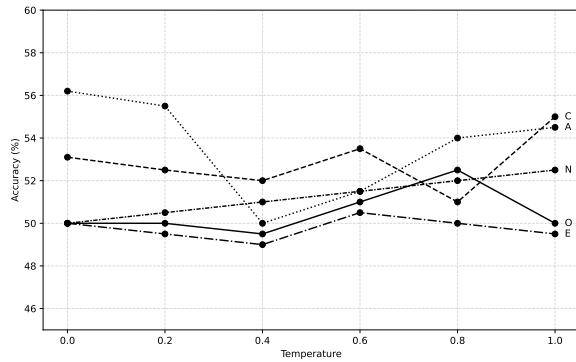


Figure 5: Per-trait OCEAN accuracy for GPT-3.5-turbo-0125 (SFT) across decoding temperatures. Accuracy varies by  $\leq 6\%$ , confirming that temperature has a minor effect on trait-level scores.

#### A.4. Prompt Templates

The codebase uses format-specific prefixes to parse model responses reliably. For the numeric prompt set (S1), the response prefix is: `My score for the statement is:`, while for the string and alphabetical sets (S2/S3), the prefix is: `My answer to the statement is:`. These prefixes enable straightforward extraction of structured responses from free-form model output.

For essay generation during evaluation, the following template is used:

```
Write a free-hand essay about your
feelings in the moment as if you are a
person who is {trait_desc}.
```

Here, `{trait_desc}` is dynamically populated with the target personality description (e.g., “positive in openness, negative in conscientiousness, ...”) corresponding to one of the 32 possible Big Five binary combinations.

#### A.5. Dataset Statistics

Table 9 reports the sample counts for every data artifact in the repository, grouped by training paradigm. The base SFT split contains 1,652 training samples derived from the Essays Dataset after OpenAI moderation filtering ( $\approx 300$  samples removed), following an 80/10/10 train/validation/test ratio. The *with questionnaires* variants double these counts because each essay appears twice: once with deterministic ideal questionnaire responses and once with randomized responses whose trait-level means match the target personality. The DPO/ORPO preference datasets are approximately three times the base training size, reflecting the triplet sampling strategy (3 rejected alternatives per chosen essay) described in Section 4.2.3. Preference data is constructed for both

the base essays and the questionnaire-augmented variants.

Training Paradigm	Split	Samples
SFT (essays only)	Train	1,652
	Validation	207
	Test	207
SFT (essays + questionnaires)	Train	3,304
	Validation	414
	Test	414
DPO / ORPO (preference pairs)	Train	5,577
	Test	621
DPO / ORPO (with questionnaires)	Train	11,154
	Test	1,242

Table 9: Dataset statistics grouped by training paradigm. The base SFT split uses an 80/10/10 ratio from the filtered Essays Dataset. The questionnaire-augmented SFT variant doubles each split (ideal + randomized responses). DPO/ORPO triples the datasets via 3 rejected alternatives per chosen essay, including both base and questionnaire-augmented variants.

#### A.6. Training Input Structure by Variant

Table 10 shows how the training input is assembled for each experimental variant. SFT variants use a chat-style format (system/user/assistant), while DPO/ORPO uses a flat prompt/chosen/rejected structure consumed directly by the TRL preference trainers.

Variant	Role	Content
SFT (essays only)	System	You are writing an essay that mimics the personality of real humans. You will be given a binary requirement for their Big Five traits.
	User	Write an essay as a person {O} in openness, {C} in conscientiousness, {E} in extroversion, {A} in agreeableness, and {N} in neuroticism.
	Assistant	<i>&lt;human-authored essay&gt;</i>
SFT (essays + questionnaires)	System	You are writing an essay that mimics the personality of real humans. You will be given a binary requirement for their Big Five traits along with their questionnaire responses.
	User	Write an essay as a person {O} in openness, ..., and {N} in neuroticism.  The following are the person's responses to personality questionnaire items: [Agreeableness Questionnaire Responses] Q: Accept people as they are. -> Score: 1 Q: Am annoyed by others' mistakes. -> Score: 5 ... [Extraversion Questionnaire Responses] ... <i>(repeated for all five OCEAN dimensions using the train-split items)</i>
	Assistant	<i>&lt;human-authored essay&gt;</i>
DPO / ORPO (preference pairs)	Prompt	Write a free-hand essay about your feelings in the moment as if you are a person who is {E} in extraversion, {N} in neuroticism, {A} in agreeableness, {C} in conscientiousness, and {O} in openness.
	Chosen	<i>&lt;essay matching the target OCEAN profile&gt;</i>
	Rejected	<i>&lt;essay with a different, non-matching OCEAN profile&gt;</i>
DPO / ORPO (with questionnaires)	Prompt	Write a free-hand essay about your feelings in the moment as if you are a person who is {E} in extraversion, {N} in neuroticism, {A} in agreeableness, {C} in conscientiousness, and {O} in openness.  The following are the person's responses to personality questionnaire items: <i>(same per-trait Q&amp;A format as SFT + questionnaires)</i>
	Chosen	<i>&lt;essay matching the target OCEAN profile&gt;</i>
	Rejected	<i>&lt;essay with a different, non-matching OCEAN profile&gt;</i>

Table 10: Training input structure for each experimental variant. {O}, {C}, {E}, {A}, {N} denote the binary trait descriptors (positive/negative). SFT uses chat-style messages consumed by the tokenizer's chat template; DPO/ORPO uses flat prompt/chosen/rejected fields processed by TRL's preference trainers. The questionnaire-augmented variants append per-trait Q&A items from the train split of the IPIP-NEO inventory to the prompt.