

PERSONALITY ALIGNMENT OF LARGE LANGUAGE MODELS

Minjun Zhu^{1,2}, Yixuan Weng², Linyi Yang³, Yue Zhang^{2*}

¹Zhejiang University ²School of Engineering, Westlake University

³University College London

zhuminjun, zhangyue@westlake.edu.cn, wengsyx@gmail.com

yanglinyiucd@gmail.com

ABSTRACT

Aligning large language models (LLMs) typically aim to reflect general human values and behaviors, but they often fail to capture the unique characteristics and preferences of individual users. To address this gap, we introduce the concept of Personality Alignment. This approach tailors LLMs' responses and decisions to match the specific preferences of individual users or closely related groups. Inspired by psychometrics, we created the Personality Alignment with Personality Inventories (PAPI) dataset, which includes data from over 320,000 real subjects across multiple personality assessments - including both the Big Five Personality Factors and Dark Triad traits. This comprehensive dataset enables quantitative evaluation of LLMs' alignment capabilities across both positive and potentially problematic personality dimensions. Recognizing the challenges of personality alignments—such as limited personal data, diverse preferences, and scalability requirements—we developed an activation intervention optimization method. This method enhances LLMs' ability to efficiently align with individual behavioral preferences using minimal data and computational resources. Remarkably, our method, PAS, achieves superior performance while requiring only 1/5 of the optimization time compared to DPO, offering practical value for personality alignment. Our work paves the way for future AI systems to make decisions and reason personally, enhancing the relevance and meaning of AI interactions for each user and advancing human-centered artificial intelligence. The dataset and code are released at <https://github.com/zhu-minjun/PAlign>.

1 INTRODUCTION

The alignment with human preferences has been a focal point of both theoretical and applied research for AI systems (Shen et al., 2023; Ji et al., 2023; Hendrycks et al., 2023). However, existing alignment approaches often treat society values as a monolithic construct, lacking the granularity needed to cater to individual differences. This broad approach may overlook the nuanced preference of individual users or closely related groups, leading to a one-size-fits-all model that fails to effectively serve diverse user needs (Fernandes et al., 2023; Carranza et al., 2023; van Wynsberghe, 2021). This gap is particularly critical as personality alignment can enhance user satisfaction, trust, and engagement by making interactions more meaningful. As shown in Figure 1, we propose "**Personality Alignment**":

Aligning AI behavior to match an individual's unique preferences and priorities of specific individuals, mirroring their behavior, thinking, and decision-making. It ensures AI responses not only are precise but also resonate with the user's traits, understanding their communication nuances and preferences.

As large language models (LLMs) advance (Solaiman et al., 2019; Farina & Lavazza, 2023), AI systems increasingly exhibit societal values and human-like personality traits (Ng et al., 2024; Bai et al., 2022b; Durmus et al., 2023; Zhang et al., 2024; tse Huang et al., 2024; Jang et al., 2024). The

*Correspondence to: Yue Zhang (zhangyue@westlake.edu.cn)

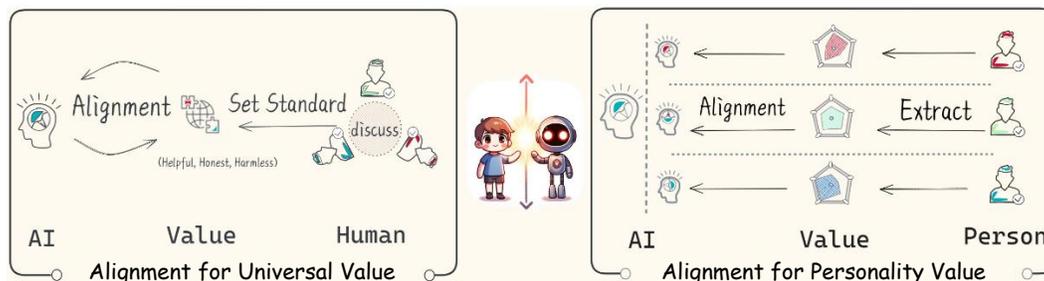


Figure 1: On the left, AI aligns to broad human values like helpfulness, honesty, and harmlessness using a standard set. On the right, the focus is on aligning AI behavior with individual users’ specific traits and preferences, using detailed profiles to reflect unique personal values.

Role-playing tasks (Li et al., 2023; Salemi et al., 2024; Gupta et al., 2024b; Yu et al., 2024) highlight this trend. However, human preferences are complex, every individual has unique preferences and needs that must be considered to effectively meet user expectations in real-world applications (Pfeifer & Bongard, 2006). For example, in customer service, some users prioritize politeness, while others prefer efficiency. Therefore, personalization is crucial for aligning AI systems with user preferences. Traditional Role-Playing Agents simulate specific professions or roles, rather than deeply aligning with individual preferences (Chen et al., 2024; Kulveit, 2018; Hubinger et al., 2019; Stray, 2020). The lack of large-scale, theoretically supported datasets for aligning AI with diverse human personalities and preferences underscores the need for further research in personality and personality alignment (Jang et al., 2023; Lou et al., 2024; Li et al., 2024b).

There remains a practical challenge: aligning a model like Llama-2 typically requires around 3 million sentences (Touvron et al., 2023a). For personality alignment, collecting such an enormous set of behavioral data from individuals is both expensive and unrealistic. Instead, we propose a training-free personal-alignment paradigm, providing over 320,000 unique preference samples across two widely-used personality assessments. Specifically, we utilize the IPIP-NEO-120 questionnaire (Johnson, 2014), which quantifies individual behavior patterns across five dimensions: Openness, Conscientiousness, Extraversion, Agreeableness, and Neuroticism (De Raad, 2000; Roccas et al., 2002). This questionnaire enables subjects to provide detailed personality measures within 10-20 minutes (Aluja & Garcia, 2004). We collect preference data from over 300,000 subjects completing both IPIP-NEO-120 and IPIP-NEO-300 (Goldberg et al., 1999) questionnaires. Additionally, we incorporate 18,192 independent samples from the Short Dark Triad (Jones & Paulhus, 2014) assessment, which measures three negative personality traits (Machiavellianism, Narcissism, and Psychopathy) through 27 questions. We manipulate these datasets to build the Personality Alignment with Personality Inventories (PAPI) dataset. For evaluation, we selected 300 representative subjects from each assessment as independent test sets, establishing a comprehensive framework for assessing LLMs’ personality alignment capabilities across both positive and negative personality dimensions.

To manage the behaviors of millions, the system must be scalable, ensuring quality and efficient alignment. Traditional RLHF techniques, which lack data and computational efficiency (Yao et al., 2023), are unsuitable for this task. Similarly, Prompt Engineering, though efficiently, often varies too greatly between models to be reliable for precise alignment (Zhou et al., 2023b). While recent activation-based approaches have shown promise in steering model behaviors (Turner et al., 2023; Li et al., 2024a), they typically focus on maximizing a single objective (e.g., safety or truthfulness). In contrast, we propose a more nuanced approach that simultaneously searches for optimal activation value offset directions and distances across multiple personality dimensions, ensuring appropriate alignment levels. This balanced activation steering is crucial for maintaining authentic personality expression while achieving desired alignment. We propose a simple and effective personality alignment method: Personality Activation Search. PAS identifies the most effective attention heads in language models (LLMs) using established behavioral patterns that correlate with the five major personality traits. It then adjusts the activations in these specific directions and optimizes the magnitude of movement to achieve the best results. This approach is computationally efficient because it doesn’t rely on backpropagation to modify model parameters. Instead, PAS offers a minimally invasive intervention, resulting in greater data efficiency and more controlled personality alignment than traditional Prompt Engineering (Jiang et al., 2024).

We present a novel systematic and comprehensive personality alignment framework that enables LLMs to understand and align precisely with specific values, preferences, and behaviors required. Compared to aligning with universal values, personality alignment requires LLMs to align with individual values, preferences, and behaviors, which ensures that the LLMs can accurately match or mimic the specific individual characteristics that users need and pursue. To achieve this goal, we introduce the PAPI dataset, which covers over 320,000 subjects of different ages, genders, and periods using multiple-choice questions to comprehensively measure both positive and negative personality traits. Based on that, we propose the PAS method, which aims to address the challenges of scarce individual characteristic data and high scalability in personality alignment. Experiments have shown its high alignment efficiency compared to DPO and PPO (Schulman et al., 2017; Rafailov et al., 2023) where it requires only 1/6 of the time to achieve superior performance. It even outperforms the GPT-4o model using PAS based on the Llama-3-8B model.

2 RELATED WORK

Alignment with language models. The concept of integrating human values into AI systems, initially proposed by Dewey Dewey (2011), underscores the necessity of value learning, which entails the AI’s ability to process and prioritize a broad spectrum of human values and preferences. In the context of LLMs—functioning as high-performance assistants—aligning these models with human values is crucial, especially since they operate in complex reasoning tasks across various domains (Hadar-Shoval et al., 2024; Wei et al., 2022b;a; Weng et al., 2023). Methods such as RLHF (Stiennon et al., 2020; Bai et al., 2022a; Dong et al., 2023) and RLAIIF (Lee et al., 2023; Sharma et al., 2024) seek to align LLMs by employing strategies like PPO (Schulman et al., 2017) and DPO (Rafailov et al., 2023). These approaches focus on maximizing cumulative rewards that are closely aligned with predefined human values (Azar et al., 2023; Lin et al., 2023; Dai et al., 2024; Swamy et al., 2024; Dai et al., 2024). Contrary to these broad approaches, our research focuses on personality alignment, which tailors the LLMs’ behavior to each individual’s complex value set, thereby enhancing the model’s capacity to understand and reflect individual human values.

Preferences of language models. The community of LLMs has started using personality testing tools for both qualitative and quantitative evaluations (Hoover et al., 2019; Jiang et al., 2022; Qiu et al., 2022; Xu et al., 2023; Salemi et al., 2024; Wang et al., 2024c). For example, the Machine Personality Inventory (MPI) Jiang et al. (2024), utilizes the Big-Five Personality Factors to standardize the assessment of LLMs. This approach analyzes the behaviors of LLMs across the five dimensions of the Big-Five, providing a quantitative interpretation of the values and personality traits of LLMs. Such tools are grounded in widely accepted theories, including the Big-Five and the 16 Personality Factors (Cattell & Mead, 2008), helping to contextualize LLMs’ personality traits in terms of human individual differences (Gupta et al., 2024a; Wang et al., 2024b; Aher et al., 2023; Liu et al., 2024a; Li et al., 2023). These studies indicate that aligned LLMs share similar values with humans (Mehta et al., 2020; Pan & Zeng, 2023; Mao et al., 2023; Go et al., 2023; Chen et al., 2024). Jang et al. (2023) has also mentioned the similar term "personality alignment", which considers a limited scope of preference personalization (Wang et al., 2024a; Zhuang et al., 2024; Zeng et al., 2023), focusing on only 8 different conflicting user preferences derived from combinations of three dimensions: Expertise (elementary vs. expert), Informativeness (concise vs. informative), and Style (friendly vs. unfriendly). Our work extends these theories and methods by leveraging psychological measurement to assess the value alignment gap between aligned LLMs and individual users. We aim to understand how closely aligned LLMs are with individual values through personality and personalized adjustments.

Interventions Activate. Activation interventions in neural networks trigger specific changes in internal activations (Olsson et al., 2022; Wu et al., 2024a) and have emerged as a pivotal tool in model robustness (He et al., 2019), editing (Meng et al., 2022), circuit finding (Goldowsky-Dill et al., 2023), and knowledge tracing (Geva et al., 2023). These interventions are advantageous due to their adjustable nature and minimal invasiveness (Wu et al., 2024b). Prior research has demonstrated that activation interventions can effectively identify the crucial heads and directions in a model’s internal representations (Burns et al., 2023a; Liu et al., 2024b), which offers better data efficiency and faster optimization than traditional reinforcement learning approaches (Turner et al., 2023; Rimsky et al., 2023; Shi et al., 2024; Weng et al., 2024a). Prior activation steering work has focused on maximizing single objectives like safety (Wang & Shu, 2024; Zheng et al., 2024) or truthfulness (Li et al., 2024a). Building on these findings, our method enhances scalability by pinpointing value directions of

different individuals and calculating optimal offset distances, aiming to achieve appropriate alignment for each dimension while maintaining model capabilities.

3 PERSONALITY ALIGNMENT DATASET CONSTRUCTION

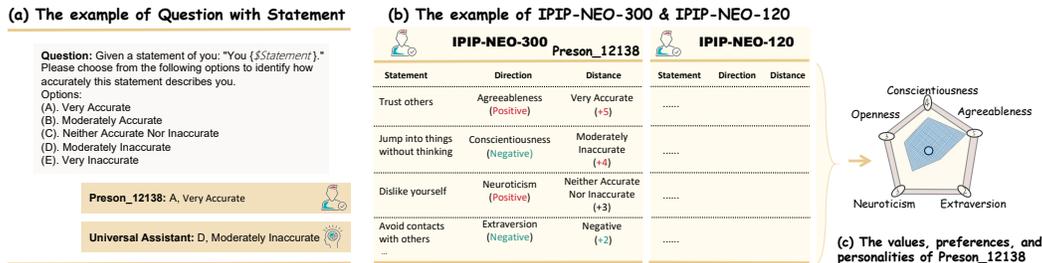


Figure 2: Overview of the PAPI dataset. (a) Illustrates the comparison between the subject’s self-assessment and the AI’s assessment of a specific question. (b) Shows an example of the IPIP-NEO-120 and IPIP-NEO-300 questionnaire responses. (c) Depicts the Big Five personality traits profile.

Motivation. To systematically evaluate the consistency between Personality-aligned LLMs and human preferences, we create the large-scale Personality Alignment with Personality Inventories (PAPI) dataset, comprising 325,505 samples across multiple-choice questionnaires. The dataset uniquely combines two complementary personality assessment systems: 307,313 samples from the Big Five personality traits (IPIP-NEO assessments) and 18,192 samples from the Dark Triad inventory. While the Big Five (IPIP) samples measure adaptive personality traits through IPIP-NEO-120 and IPIP-NEO-300 questionnaires, the Dark Triad samples assess subclinical maladaptive traits (Machiavellianism, Narcissism, and Psychopathy) through 27 targeted questions. This comprehensive approach captures both healthy personality dimensions and potentially problematic traits. As shown in Figure 2, subjects choose behaviors similar to their own (a), yielding questionnaire responses (b) that collectively form detailed personality profiles (c).

Question Type. Our dataset contains two independent questionnaire systems, each using a five-point scale from "Very Accurate" to "Very Inaccurate." The first system comprises IPIP questions associated with the Big Five personality dimensions. We use each subject’s IPIP-NEO-120 questionnaire responses to predict their IPIP-NEO-300 answers, testing the model’s ability to generalize from known personality patterns to broader behavioral descriptions. The second system consists of Dark Triad assessments, where we use 18 questions (6 each for Machiavellianism, Narcissism, and Psychopathy) to align with and predict responses to the remaining 9 questions (3 per trait). Both systems evaluate the stability of individual personality tendencies and the model’s capacity to capture and generalize personality patterns, though they operate independently on distinct trait dimensions.

Data Collection. We have collected 307,313 samples based on the IPIP-NEO-120 and IPIP-NEO-300 from the International Personality Item Pool (IPIP)¹. This includes approximately 60% (185,149) female and 40% (122,164) male participants, covering ages from 10 to 99 years (average age 25.19 years). The data were collected through online surveys on a website, requiring participants to provide genuine feedback on 120 & 300 behavioral questions. Participants are volunteers from various industries worldwide, including the US, UK, France, India, and China, with tests conducted between 1998 and 2019. Each participant had to actively acknowledge that careless responses would invalidate the usefulness. Each survey took between 30 to 50 minutes, was anonymous, and collected no traceable or socioeconomic status-related data. We have permission from the administrators to

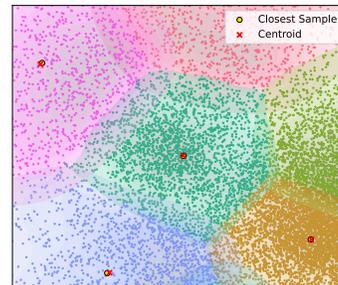


Figure 3: Visualization of the K-Means Clustering on the PAPI Dataset. The centroids are marked in red, demonstrating the central point of each cluster, while the closest samples to these centroids are highlighted in gold.

¹<https://ipip.ori.org/index.htm>

use IPIP items, scales, and inventories for any purpose, commercial or non-commercial. Additionally, we collected 18,192 Dark Triad samples² using the same collection pipeline and quality control procedures.

Data Processing. We apply K-Means clustering separately to the IPIP and Dark Triad datasets. From the IPIP dataset (307,313 samples), we select 300 representative clusters as the first part of our Test-Set. Similarly, we cluster the Dark Triad dataset (18,192 samples) and select 300 representative samples, forming the second part of our **Test-Set**. The remaining data constitutes our **Dev-Set**. The large-scale Dev-Set captures personality variations across different ages (from teenagers to seniors), genders, cultural backgrounds (spanning multiple countries and regions), and identity groups. This comprehensive representation is crucial for developing inclusive alignment methods that respect and reflect the full spectrum of human personality diversity. Sample selection minimizes Euclidean distance to cluster centroids: $\min_{x \in S_k} |x - \mu_k|_2$, where μ_k is the centroid of cluster k , and x is a sample within cluster k (S_k), aiming to optimize cluster representativeness by minimizing the distance between samples and their respective centroids.

Evaluation Methods. Inspired by the MPI (Jiang et al., 2024), we set up a behavioral difference score as an evaluation metric. Using a psychological measurement approach, we establish a simple scoring function S : for positively correlated questions, we assign scores from 5 to 1 from “Very Accurate” to “Very Inaccurate”. We expect the aligned language model scores to be close to a specific individual’s scores, hence we use an absolute value method for measurement. We test each trait d in the test section, with each participant’s behavioral difference Aligned Score calculated as follows:

$$\text{Aligned Score}_d = \frac{1}{M} \sum_{i=1}^M \left(\frac{1}{N_{d,i}} \sum_{\kappa \in D_{d,i}^{Test}} |f(\text{LLM}(\kappa, \text{template})) - f(\text{Person}(\kappa, \text{template}))| \right), \quad (1)$$

In this formula, $D_{d,i}^{Test}$ denotes the set of behaviors related to trait d in the test section for subject i . We use the given template to obtain responses from both the LLM and the Person, and then calculate the absolute difference between the two. $N_{d,i}$ represents the count of these behaviors. Each κ is an independent behavior sample from this set, with the LLM’s responses evaluated against those in D^{Train} . The $f(\cdot)$ represents the scoring function for multiple-choice questions. We calculated the deviation distance between the two using the absolute value method and computed the average deviation distance for all behaviors. The range for Score is from 0 to 4, where 0 means perfect alignment with individual behavior.

4 PERSONALITY ACTIVATION SEARCH

Popular alignment techniques like PPO, DPO, and IPO strive to align language models with desired behaviors, but personalizing LLMs for individual preferences without extensive training presents a significant challenge. To address this, we’ve developed Personality-Activation Search (PAS), an efficient method that fine-tunes model activations to closely match human preferences. PAS identifies key activations within the model that significantly impact Personality and personalized behaviors and adjusts these activations toward specific preferences.

PAS works by identifying multiple directions in the transformer’s activation space that satisfy personality attributes and then intervening along these directions to adjust the activation values, achieving minimally invasive alignment. The entire process does not require changing the model’s weights or backpropagation. It only requires multiple forward passes to determine the offset direction and distance, resulting in a high optimization efficiency.

4.1 LANGUAGE MODELS

In our research, we focus on a transformer-based language model which may be in various stages, including pre-training, supervised fine-tuning, or undergoing RLHF. The transformer model primarily consists of several layers, each equipped with a multi-head attention (MHA) mechanism and a multilayer perceptron (MLP) (Vaswani et al., 2017).

²<https://openpsychometrics.org>

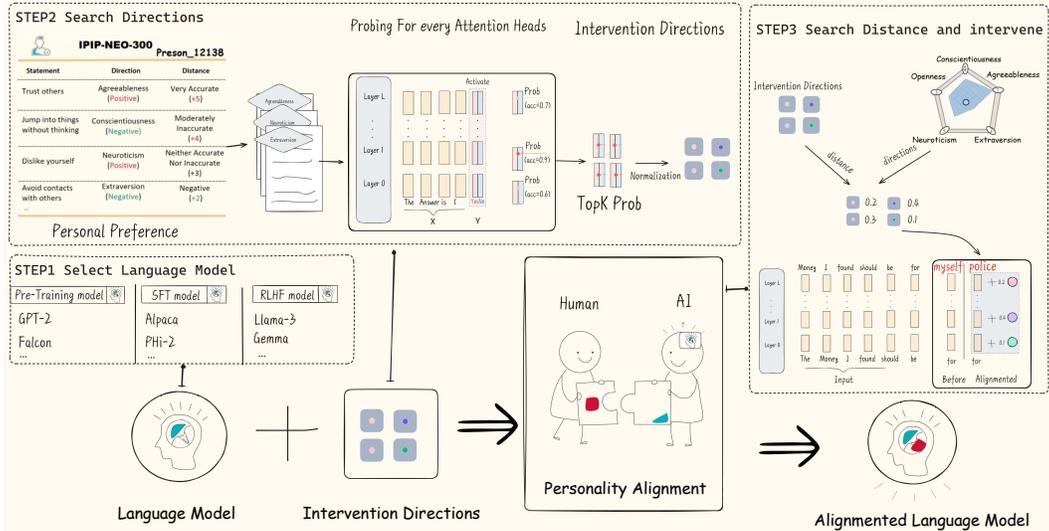


Figure 4: Overview of the PAS Process. Step 1 involves selecting a LLM. Step 2 includes creating and activating personality responses using the answer of questionnaire, the subject’s answer in the PAPI dataset, and probing each attention head. Step 3 integrates these activations for personality alignment, adjusting model outputs to reflect individual user traits, resulting in an aligned LLM.

During inference, input tokens are encoded to high-dimensional embeddings, where in each transformer layer, the MHA process is defined as $x_{l+1} = x_l + \sum_{h=1}^H Q_l^h \text{Att}_l^h(P_l^h x_l)$, with H being the number of heads, $x_l \in \mathcal{R}^{DH}$ representing the activations, and $P_l^h \in \mathcal{R}^{D \times DH}$ and $Q_l^h \in \mathcal{R}^{DH \times D}$ projecting activations into and back from specific headspace.

While previous methods like Inference-Time Intervention (ITI) adjust language models by applying a fixed, large shift in the activation space to influence output biases, this approach is inadequate for addressing alignment issues, especially when personalization is required. Fixed shifts can be too coarse and may not capture the nuances of individual preferences. To overcome this limitation, we have enhanced the intervention strategy by dynamically adjusting the activation distances based on personal preference vectors. Specifically, we introduce these vectors into the residual stream between the Attention mechanism and the projection matrix Q_l^h . This allows us to fine-tune the model’s activations in a way that aligns with specific user traits and preferences, facilitating personalized alignment without altering the model’s weights or relying on large, inflexible shifts.

4.2 SEARCH THE DIRECTIONS FOR ACTIVATION INTERVENTION

To achieve personality alignment, we trained “probes” to get activations as the direction of personal preference. It is started by selecting statements from the training dataset of each subject, formatted as positive samples: “*Question: Given a statement of you: ‘\$Statement’, Do you agree? Answer: Yes*”, and negative samples: “*Question: Given a statement of you: ‘\$Statement’, Do you agree? Answer: No*”. Activations from each head are retained and split into a training set (60%) and a validation set (40%). These activations serve as features for the probe (Equation 2) of per attention head.

$$p_{\theta}(x_l^h) = \text{sigmoid}(\langle \theta, x_l^h \rangle), \quad (2)$$

where the $x_l^h \in \mathcal{R}^D$, and $\theta \in \mathcal{R}^D$ is the weights of h-heads in layer l in the residual stream.

We identify the most influential heads based on their predictive accuracy during validation. Notably, it is necessary to disperse interventions during this process (Li et al., 2024a). Therefore, we select the top K probabilities for intervention. The parameters θ_l^h from these probes, after normalization, are treated as the directions along which behaviors are most separable. These directions are the most informative for aligning the model with individual preferences (Burns et al., 2023b).

During inference, we integrate the intervention into the output of each selected head:

$$x_{l+1} = x_l + \sum_{h=1}^H Q_l^h \left(\text{Att}_l^h(P_l^h x_l) + \alpha \sigma_l^h \right), \quad (3)$$

where α controls the adjustment intensity, and σ_l^h is the standard vectors along the normalized direction θ_l^h . We adjust activations in the personality direction, focusing only on the top K heads, while leaving the unselected heads unchanged. In practice, the model parameters remain unaltered; we only need to store approximately 1,000 parameters, denoted as σ_l^h . These σ_l^h parameters represent the direction of activation shifts towards personality preferences, rather than memorizing specific preferences. This approach is highly parameter-efficient, as the σ_l^h parameters constitute only about 1/10,000,000 of the total parameters compared to a full-parameter model.

4.3 SEARCH DISTANCE FOR ACTIVATION INTERVENTION

To optimize the intervention’s impact on personality behaviors in the language model, we search to determine the optimal value of α precisely. This method is suitable for functions with a single peak or trough. The search, constrained within the interval $[0, 10]$, utilizes the objective function $f(\alpha)$ to compute the negative aggregate score, aiming to minimize this value. This approach efficiently identifies the α that best enhances alignment with desired behaviors while maintaining the model’s internal coherence, ensuring adjustments remain within a beneficial range.

$$\text{Optimal } \alpha = \operatorname{argmin}_{\alpha \in [0, 10]} f(\alpha). \quad (4)$$

By iteratively narrowing the search range based on the function’s values at each step, this method efficiently locates the α that yields the lowest mean Score.

5 EXPERIMENTS

Method	Alignment Mode	Big-Five					Dark Triad			Score
		Agreeableness ↓	Conscientiousness ↓	Extraversion ↓	Neuroticism ↓	Openness ↓	Machiavellianism ↓	Narcissism ↓	Psychopathy ↓	
GPT-4o(omni)										
Few-Shot	Black-Box (Prompt-Based)	1.02	0.83	0.81	0.80	0.96	0.80	0.76	0.83	6.81
P^2	Black-Box (Prompt-Based)	1.44	1.45	1.63	1.73	1.46	1.17	2.04	2.00	12.92
Llama-3-8B-Instruct										
PPO	White-Box (Alignment)	1.63	1.51	1.45	1.42	1.61	1.48	1.98	2.19	13.27
DPO	White-Box (Alignment)	1.54	1.42	1.54	1.74	1.21	1.41	1.99	2.12	12.97
Prompt-MORL	White-Box (Alignment)	1.18	0.93	1.01	1.23	1.00	1.42	2.14	1.78	10.88
Personalized-Soups	White-Box (Alignment)	1.06	0.91	0.93	1.28	0.80	1.08	1.76	1.84	9.66
Few-Shot	Black-Box (Prompt-Based)	1.28	1.30	1.40	1.09	0.89	1.16	2.03	2.00	11.15
P^2	Black-Box (Prompt-Based)	1.39	1.33	1.41	1.22	1.68	1.17	2.04	2.01	12.25
PAS (Ours)	White-Box (Alignment)	0.94	0.91	0.86	0.98	0.72	0.96	1.85	1.67	8.89
Llama-3-70B-Instruct										
PPO	White-Box (Alignment)	1.56	1.59	1.43	1.40	1.56	1.52	1.96	1.90	12.92
DPO	White-Box (Alignment)	1.46	1.25	1.45	1.48	1.57	1.22	2.08	1.79	12.30
Prompt-MORL	White-Box (Alignment)	1.10	1.11	1.02	1.30	1.24	1.15	1.99	1.76	10.67
Personalized-Soups	White-Box (Alignment)	0.99	0.96	1.16	1.02	1.08	1.11	1.95	1.77	10.04
Few-Shot	Black-Box (Prompt-Based)	1.06	0.94	0.96	1.03	1.22	1.04	1.89	1.80	9.94
P^2	Black-Box (Prompt-Based)	1.42	1.33	1.36	1.35	1.66	1.02	2.11	1.93	12.18
PAS (Ours)	White-Box (Alignment)	0.98	0.89	0.87	1.01	0.99	1.01	1.84	1.62	9.21

Table 1: Comparison of alignment methods on the PAPI dataset using the Aligned Score of Big-Five personality traits and Dark Triad traits. The Score represents the overall specifics of all dimensions.

5.1 EXPERIMENTAL SETTINGS

Personality Alignment. We select four classic methods as baselines to compare against the PAS approach, which includes two black-box methods: Few-Shot and Personality Prompt (P^2) Jiang et al. (2024) relying on in-context learning (ICL), and two white-box methods requiring training, PPO (Schulman et al., 2017) and DPO (Rafailov et al., 2023). We consider recently released Llama-3-8B-Instruct and Llama-3-70B-Instruct (Touvron et al., 2023a;b; AI, 2024) model as backbones, and SoTA model, GPT-4o (omni) (OpenAI, 2024) as baseline. The scores in the table represent the distance of alignment calculated by Eq. 2, with lower scores indicating better alignment.

Open-ended Generation. We conduct experiments on Llama-3-8B-Instruct and Llama-3-70B-Instruct. For the open-ended generation task, we use GPT-4o as the annotator. First, we align the language model to a specific individual (ID 181591, 234619, 210204, 259396, and 22835 from the Test-Set). Then, we have the model generate open-ended responses based on each statement from the IPIP-NEO-300. We compare our method with ICL methods and classical alignment methods.

Complex Reasoning. We explore the generalization capabilities of the PAS method by evaluating its impact on complex reasoning tasks across eight datasets: GSM8K, CommonSenseQA, AddSub, MultiArith, SVAMP, BigBench-Date, StrategyQA, and Coin Flip (Hosseini et al., 2014; Talmor et al., 2018; Arkil et al., 2021; Cobbe et al., 2021; Suzgun et al., 2022; Roy & Roth, 2016; Wei et al., 2022b; Kojima et al., 2022; Weng et al., 2024b). In this part, we expect to assess how conscientiousness adjustments influence the LLM’s performance in reasoning tasks. Using the llama-3-8B-Instruct model, we apply PAS to adjust the model’s activations to simulate two fictional personas: one with extreme conscientiousness (as **Positive**) and one without extreme conscientiousness (as **Negative**). Vanilla CoT served as the baseline.

5.2 RESULTS ON PAPI

Table 1 demonstrates that the PAS method achieves state-of-the-art performance across different models and personality dimensions. For the Llama-3-8B-Instruct model, PAS achieves a Composite Score of 8.89, significantly outperforming other methods including Personalized-Soups (9.66) and Prompt-MORL (10.88). Similarly, for the Llama-3-70B-Instruct model, PAS scores 9.21,

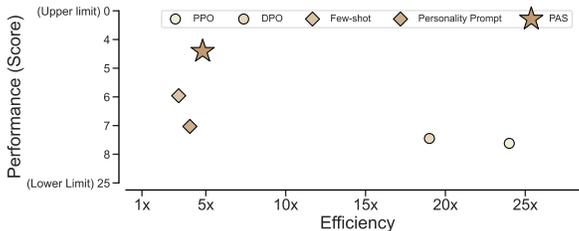


Figure 5: Comparison of performance and efficiency for various alignment methods. The larger the multiple on the x-axis, the lower the efficiency.

compared to 9.94 for the next best method. In Big Five dimensions, PAS excels with consistently low scores across both model sizes: Agreeableness (0.94/0.98), Conscientiousness (0.91/0.89), Extraversion (0.86/0.87), Neuroticism (0.98/1.01), and Openness (0.72/0.99). Notably, PAS also demonstrates strong performance in Dark Triad traits, particularly in Machiavellianism (0.96/1.01) and Psychopathy (1.67/1.62), though showing relatively higher scores in Narcissism (1.85/1.84). Only GPT-4o shows better performance in Dark Triad alignment with scores of 0.80, 0.76, and 0.83 for the three traits respectively, highlighting the challenge of aligning these complex negative personality dimensions in smaller models. In the appendix E.4, we also conducted further performance comparisons with full parameter fine-tuning and LoRA. PAS similarly demonstrated higher efficiency and superior performance. In appendix E.5, we provide more comprehensive analyses across diverse demographic groups, including detailed distributions of age, gender, and nationality, along with in-depth examination of performance patterns across various underrepresented populations, further validating PAS’s effectiveness in handling diverse personality characteristics.

Efficiency. As shown in Figure 5, training a separate language model for each individual demands (e.g. DPO and PPO) substantial computational resources and time, making it an infeasible solution for personality alignment. PAS addresses efficiency challenges by making targeted adjustments to a small subset of activation values during inference, in specific directions. This approach requires minimal computational overhead, similar to Few-Shot methods, and only about 1/6 of the time needed for PPO, yet achieves remarkable personality alignment results. By intervening at the activation level rather than relying on extensive training or the variability of prompts, PAS provides a more efficient and precise method for aligning language models with individual user preferences.

Alignment vs. ICL. Our experiments highlight that direct Alignment (White-Box) methods, especially PAS, outperform In-Context Learning (ICL) (Black-Box) methods. By aligning model parameters with user preferences, PAS provides better performance compared to prompt-based approaches. The PAS method’s ability to efficiently shift activations towards user preferences without altering the model’s core parameters gives it a significant advantage over traditional ICL methods.

Why Did Scaling Laws Fail? The results indicate that scaling laws alone do not ensure optimal alignment, particularly in domain-specific tasks like personality alignment. Although larger models like Llama-3-70B-Instruct benefit from increased scale, their broader knowledge and more general alignment capabilities may actually hinder their ability to precisely capture individual personality traits. This phenomenon likely occurs because larger models are trained to maintain general-purpose capabilities and broad knowledge, making them less adaptable to highly specific individual preferences. PAS demonstrates that targeted intervention in activation patterns can achieve better personality alignment

than parameter modification or model scaling. By specifically focusing on personality-relevant patterns rather than general knowledge, PAS achieves superior performance with computational efficiency, emphasizing the importance of precise control over raw model size.

5.3 GENERALIZATION RESULTS

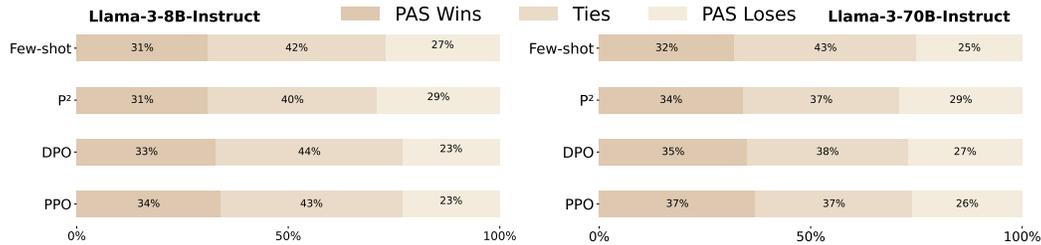


Figure 6: The win rates of PAS compared to classical alignment methods evaluated by GPT-4o.

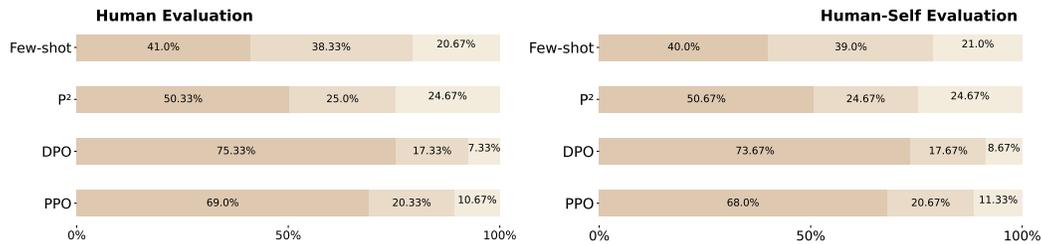


Figure 7: The win rates of PAS compared to classical alignment methods evaluated by Human.

Considering that during the process of adjusting the model, although specific goals can be achieved, it may also affect other general capabilities of the model, such as reasoning abilities. **We demonstrate through general dialogue tasks and complex reasoning tasks that our method can precisely control the model’s personality without compromising its other generalization capabilities.**

Open-Ended Generation. We used an open-ended generation task where the language model elaborated on given statements from the IPIP-NEO-300, covering “Scenario Description,” “Actions Taken,” “Thought Process,” and “Emotional Response.” The generated outputs were evaluated by GPT-4o for alignment with individual preferences, similar to the evaluation process in LIMA (Zhou et al., 2023a). Despite the Personality Alignment being initially trained with multiple-choice options, Figure 6 shows that PAS does not affect the model’s general conversational ability and reasoning capacity, it significantly outperforms baseline methods in this open-ended task. This superior performance can be attributed to PAS’s ability to fine-tune model activations to capture individual preferences effectively, enhancing the model’s adaptability to complex and nuanced tasks. PAS’s ability to generalize from multiple-choice training data to natural, open-ended generation tasks demonstrates its capacity to capture broader behavioral patterns while maintaining fluent text generation. In addition to this, we conducted a manual evaluation using the same setup in Figure 7. In addition to this, we conducted a manual evaluation process using a blind setup. Three human evaluators assessed the generated responses for three different subjects (Human Evaluation, the subset of test set: ID 172481, ID 22835, and ID 259396). These evaluators also judged responses aligned to their own personalities (Human-Self Evaluation). They evaluated the consistency between the target values and the model-generated text in terms of values, behaviors, and preferences. The results demonstrate that our PAS method consistently outperforms other approaches, reflecting its superior alignment with both general human and self-specific preferences. Further details of the manual evaluation are available in Appendix E.2.

Complex Reasoning. Experimental results, as shown in Figure 8, indicate that PAS significantly enhances the model’s reasoning capabilities when conscientiousness is positively adjusted. For instance, on the GSM8K dataset, the baseline performance was 73.47. Enhancing conscientiousness increased the performance to 74.15 while reducing conscientiousness dropped it to 72.75. Similar trends were observed across other datasets, with the model consistently outperforming the baseline

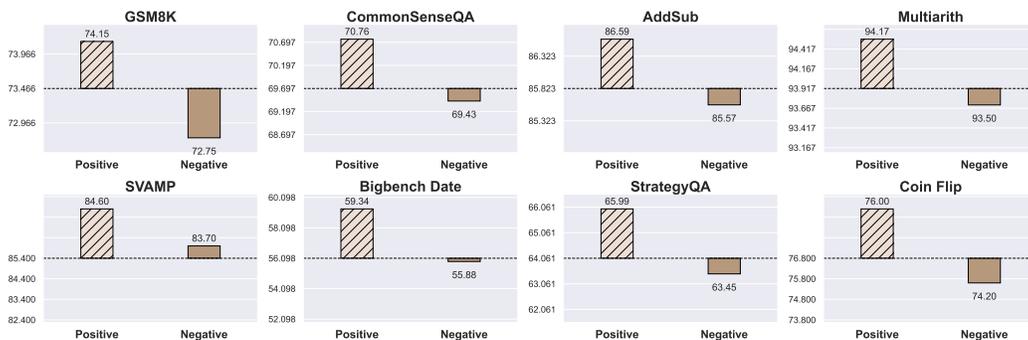


Figure 8: Performance of the llama-3-8B-Instruct model on complex reasoning tasks across eight datasets. We used Vanilla CoT (Wei et al., 2022b) as the baseline (dashed line in each subplot).

when conscientiousness was increased. This suggests that conscientiousness, associated with diligence and carefulness, positively impacts the model’s ability to handle complex reasoning tasks. Importantly, these interventions did not significantly alter the inference time, maintaining efficiency comparable to the original model. These findings demonstrate PAS’s strong generalization capabilities.

5.4 DISCUSSION: IS VALUE-ALIGNED ASSISTANT A GOOD ASSISTANT?

Certainly. We explored whether aligning a language model with user values enhances satisfaction using 300 general question-answer examples from the LIMA dataset (Zhou et al., 2023a). We evaluated both the original Llama-3-70B-Instruct model and a Personality Alignment model. Each human evaluator spent 10 minutes completing an IPIP-NEO-120 questionnaire to extract their values. We aligned the model with these values with the PAS, creating a Value-Aligned Assistant. For comparison, we also created a Value-Misaligned Assistant by inverting the questionnaire responses.

Human evaluators assessed the responses to the test prompts. Figure 9 shows that the Value-Aligned Assistant provided preferable outputs 38% of the time, with 31% ties, significantly outperforming the Value-Misaligned Assistant. For instance, a conservative person may prefer more cautious advice, which the Value-Aligned Assistant is better equipped to provide. These results underscore the importance of aligning AI systems with user values to enhance satisfaction and effectiveness.

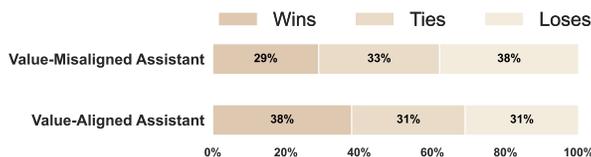


Figure 9: Performance comparison of the Value-Aligned Assistant and Value-Misaligned Assistant on the LIMA dataset’s 300 test prompts. We used the original Llama-3-70B-Instruct as the baseline comparator, anonymizing the models and randomizing their order during the evaluation.

6 CONCLUSION

We introduce the concept of Personality Alignment for language models, emphasizing the importance of tailoring AI behavior to individual user preferences and values. We release a large-scale real-world personality alignment dataset named PAPI, characterized by its real, diverse, and extensive data, to quantitatively measure the degree of Personality Alignment. Based on that, we propose a Personality Activation Search (PAS) method, which offers a practical solution with high efficiency and performance. LLMs adjusted with PAS can achieve better results and better efficiency in personalized alignment than ICL, DPO, and PPO as needed. In particular, it achieves precise alignment with just 1/6 of the computational time required by PPO. Experiments demonstrate that PAS enhances the performance of Personality alignment by making AI interactions more relevant and meaningful.

ACKNOWLEDGEMENT

This publication has been supported by the National Natural Science Foundation of China (NSFC) Key Project under Grant Number 62336006.

We thank all participants who contributed to the PAPI dataset.

ETHICS STATEMENT

While personality alignment in language models enhances human-AI interaction, it raises significant ethical concerns that warrant careful consideration. The primary challenge lies in the potential creation of psychological filter bubbles and echo chambers. When AI systems consistently align with individual personality traits, they may limit exposure to diverse perspectives, particularly concerning for individuals scoring high in Neuroticism or low in Openness. This effect becomes more pronounced when considering Dark Triad traits - for users exhibiting high Machiavellianism or Narcissism, aligned AI systems might inadvertently validate manipulative tendencies or excessive self-focus, potentially contributing to social polarization and ideological segregation in online communities.

The implementation of personality alignment also presents substantial privacy and security challenges. While our dataset maintains anonymity, the processing of detailed psychological profiles creates potential vulnerabilities. The risk extends beyond data security to the possibility of exploitation, where knowledge of individual personality traits could be used for sophisticated manipulation or targeted advertising. Furthermore, extended interaction with personality-aligned AI systems might foster psychological dependence, potentially inhibiting personal growth and adaptation, particularly in users with extreme personality trait scores.

To address these concerns, we suggest a comprehensive mitigation framework with three key components. First, we recommend implementing dynamic alignment boundaries through a monitoring system that tracks interaction patterns and adjusts alignment intensity in real-time. This system would incorporate threshold mechanisms that automatically reduce alignment strength when user-AI interactions show signs of reinforcing extreme or potentially harmful behavioral patterns. Second, we suggest an adaptive content diversity system that strategically introduces alternative viewpoints while maintaining personality alignment. This could be achieved through a balanced scoring mechanism that weighs alignment accuracy against content diversity, ensuring users receive personality-aligned responses while still being exposed to varied perspectives. Third, we advocate for a robust privacy protection framework that includes: (a) data anonymization techniques that separate personality profiles from identifiable information, (b) secure encryption protocols for personality data storage and transmission, and (c) granular user controls over which personality aspects can be used for alignment.

Looking forward, our research emphasizes the critical balance between technological advancement and ethical considerations in AI development. Future work should focus on implementing and evaluating these mitigation strategies, particularly examining their effectiveness in preventing echo chambers while maintaining the benefits of personality alignment.

REFERENCES

- Gati V Aher, Rosa I Arriaga, and Adam Tauman Kalai. Using large language models to simulate multiple humans and replicate human subject studies. In *International Conference on Machine Learning*, pp. 337–371. PMLR, 2023.
- Meta AI. Introducing meta llama 3: The most capable openly available llm to date. <https://ai.meta.com/blog/meta-llama-3/>, 2024. Accessed: 2024-05-21.
- Ebtesam Almazrouei, Hamza Alobeidli, Abdulaziz Alshamsi, Alessandro Cappelli, Ruxandra Cojocaru, M erouane Debbah,  tienne Goffinet, Daniel Hesslow, Julien Launay, Quentin Malartic, et al. The falcon series of open language models. *arXiv preprint arXiv:2311.16867*, 2023.
- Anton Aluja and Luis F Garcia. Relationships between big five personality factors and values. *Social Behavior and Personality: an international journal*, 32(7):619–625, 2004.

- Patel Arkil, Bhattamishra Satwik, and Goyal Navin. Are nlp models really able to solve simple math word problems? 2021.
- Mohammad Gheshlaghi Azar, Mark Rowland, Bilal Piot, Daniel Guo, Daniele Calandriello, Michal Valko, and Rémi Munos. A general theoretical paradigm to understand learning from human preferences, 2023.
- Jinze Bai, Shuai Bai, Yunfei Chu, Zeyu Cui, Kai Dang, Xiaodong Deng, Yang Fan, Wenbin Ge, Yu Han, Fei Huang, et al. Qwen technical report. *arXiv preprint arXiv:2309.16609*, 2023.
- Yuntao Bai, Andy Jones, Kamal Ndousse, Amanda Askell, Anna Chen, Nova DasSarma, Dawn Drain, Stanislav Fort, Deep Ganguli, Tom Henighan, et al. Training a helpful and harmless assistant with reinforcement learning from human feedback. *arXiv preprint arXiv:2204.05862*, 2022a.
- Yuntao Bai, Saurav Kadavath, Sandipan Kundu, Amanda Askell, Jackson Kernion, Andy Jones, Anna Chen, Anna Goldie, Azalia Mirhoseini, Cameron McKinnon, et al. Constitutional ai: Harmlessness from ai feedback. *arXiv preprint arXiv:2212.08073*, 2022b.
- Tom L Beauchamp et al. The belmont report. *The Oxford textbook of clinical research ethics*, pp. 149–155, 2008.
- Collin Burns, Haotian Ye, Dan Klein, and Jacob Steinhardt. Discovering latent knowledge in language models without supervision. In *The Eleventh International Conference on Learning Representations*, 2023a. URL <https://openreview.net/forum?id=ETKGuby0hcs>.
- Collin Burns, Haotian Ye, Dan Klein, and Jacob Steinhardt. Discovering latent knowledge in language models without supervision. In *The Eleventh International Conference on Learning Representations*, 2023b. URL <https://openreview.net/forum?id=ETKGuby0hcs>.
- Andres Carranza, Dhruv Bhandarkar Pai, Rylan Schaeffer, Arnub Tandon, and Sanmi Koyejo. Deceptive alignment monitoring. In *The Second Workshop on New Frontiers in Adversarial Machine Learning*, 2023. URL <https://openreview.net/forum?id=obsO44GFhh>.
- Heather EP Cattell and Alan D Mead. The sixteen personality factor questionnaire (16pf). *The SAGE handbook of personality theory and assessment*, 2:135–159, 2008.
- Jiangjie Chen, Xintao Wang, Rui Xu, Siyu Yuan, Yikai Zhang, Wei Shi, Jian Xie, Shuang Li, Ruihan Yang, Tinghui Zhu, Aili Chen, Nianqi Li, Lida Chen, Caiyu Hu, Siye Wu, Scott Ren, Ziquan Fu, and Yanghua Xiao. From persona to personalization: A survey on role-playing language agents, 2024.
- Karl Cobbe, Vineet Kosaraju, Mohammad Bavarian, Jacob Hilton, Reiichiro Nakano, Christopher Hesse, and John Schulman. Training verifiers to solve math word problems. *arXiv preprint arXiv:2110.14168*, 2021.
- Josef Dai, Xuehai Pan, Ruiyang Sun, Jiaming Ji, Xinbo Xu, Mickel Liu, Yizhou Wang, and Yaodong Yang. Safe RLHF: Safe reinforcement learning from human feedback. In *The Twelfth International Conference on Learning Representations*, 2024. URL <https://openreview.net/forum?id=TyFrPOKYXw>.
- Boele De Raad. *The big five personality factors: the psycholexical approach to personality*. Hogrefe & Huber Publishers, 2000.
- Tim Dettmers, Mike Lewis, Younes Belkada, and Luke Zettlemoyer. Gpt3. int8 (): 8-bit matrix multiplication for transformers at scale. *Advances in Neural Information Processing Systems*, 35: 30318–30332, 2022.
- Tim Dettmers, Artidoro Pagnoni, Ari Holtzman, and Luke Zettlemoyer. Qlora: Efficient finetuning of quantized llms. *ArXiv*, abs/2305.14314, 2023. URL <https://api.semanticscholar.org/CorpusID:258841328>.
- Daniel Dewey. Learning what to value. In *International conference on artificial general intelligence*, pp. 309–314. Springer, 2011.

- Hanze Dong, Wei Xiong, Deepanshu Goyal, Yihan Zhang, Winnie Chow, Rui Pan, Shizhe Diao, Jipeng Zhang, KaShun SHUM, and Tong Zhang. RAFT: Reward ranked finetuning for generative foundation model alignment. *Transactions on Machine Learning Research*, 2023. ISSN 2835-8856. URL <https://openreview.net/forum?id=m7p507zblY>.
- Esin Durmus, Karina Nyugen, Thomas I Liao, Nicholas Schiefer, Amanda Askell, Anton Bakhtin, Carol Chen, Zac Hatfield-Dodds, Danny Hernandez, Nicholas Joseph, et al. Towards measuring the representation of subjective global opinions in language models. *arXiv preprint arXiv:2306.16388*, 2023.
- Lilian Edwards. The eu ai act: a summary of its significance and scope. *Artificial Intelligence (the EU AI Act)*, 1, 2021.
- Mirko Farina and Andrea Lavazza. Chatgpt in society: emerging issues. *Frontiers in Artificial Intelligence*, 6:1130913, 2023.
- Patrick Fernandes, Aman Madaan, Emmy Liu, António Farinhas, Pedro Henrique Martins, Amanda Bertsch, José G. C. de Souza, Shuyan Zhou, Tongshuang Wu, Graham Neubig, and André F. T. Martins. Bridging the gap: A survey on integrating (human) feedback for natural language generation, 2023.
- Phoebe Friesen, Lisa Kearns, Barbara Redman, and Arthur L Caplan. Rethinking the belmont report? *The American Journal of Bioethics*, 17(7):15–21, 2017.
- Mor Geva, Jasmijn Bastings, Katja Filippova, and Amir Globerson. Dissecting recall of factual associations in auto-regressive language models. In Houda Bouamor, Juan Pino, and Kalika Bali (eds.), *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pp. 12216–12235, Singapore, December 2023. Association for Computational Linguistics. doi: 10.18653/v1/2023.emnlp-main.751. URL <https://aclanthology.org/2023.emnlp-main.751>.
- Dongyoung Go, Tomasz Korbak, Germán Kruszewski, Jos Rozen, and Marc Dymetman. Compositional preference models for aligning lms. *arXiv preprint arXiv:2310.13011*, 2023.
- Lewis R Goldberg et al. 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, 1999.
- Nicholas Goldowsky-Dill, Chris MacLeod, Lucas Sato, and Aryaman Arora. Localizing model behavior with path patching. *arXiv preprint arXiv:2304.05969*, 2023.
- Akshat Gupta, Xiaoyang Song, and Gopala Anumanchipalli. Self-assessment tests are unreliable measures of LLM personality. In Yonatan Belinkov, Najoung Kim, Jaap Jumelet, Hosein Mohebbi, Aaron Mueller, and Hanjie Chen (eds.), *Proceedings of the 7th BlackboxNLP Workshop: Analyzing and Interpreting Neural Networks for NLP*, pp. 301–314, Miami, Florida, US, November 2024a. Association for Computational Linguistics. doi: 10.18653/v1/2024.blackboxnlp-1.20. URL <https://aclanthology.org/2024.blackboxnlp-1.20/>.
- Shashank Gupta, Vaishnavi Shrivastava, Ameet Deshpande, Ashwin Kalyan, Peter Clark, Ashish Sabharwal, and Tushar Khot. Bias runs deep: Implicit reasoning biases in persona-assigned LLMs. In *The Twelfth International Conference on Learning Representations*, 2024b. URL <https://openreview.net/forum?id=kGteeZ18Ir>.
- Dorit Hadar-Shoval, Kfir Asraf, Yonathan Mizrahi, Yuval Haber, and Zohar Elyoseph. Assessing the alignment of large language models with human values for mental health integration: Cross-sectional study using schwartz’s theory of basic values. *JMIR Mental Health*, 11:e55988, 2024.
- Zhezhi He, Adnan Siraj Rakin, and Deliang Fan. Parametric noise injection: Trainable randomness to improve deep neural network robustness against adversarial attack. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 588–597, 2019.
- Dan Hendrycks, Mantas Mazeika, and Thomas Woodside. An overview of catastrophic ai risks, 2023.

- Joe Hoover, Gwenyth Portillo-Wightman, Leigh Yeh, Shreya Havaldar, Aida Mostafazadeh Davani, Ying Lin, Brendan Kennedy, Mohammad Atari, Zahra Kamel, Madelyn Mendlen, Gabriela Moreno, Christina Park, Tingyee E. Chang, Jenna Chin, Christian Leong, Jun Yen Leung, Arineh Mirinjian, and Morteza Dehghani. Moral foundations twitter corpus: A collection of 35k tweets annotated for moral sentiment. *Social Psychological and Personality Science*, 11:1057 – 1071, 2019. URL <https://api.semanticscholar.org/CorpusID:213669580>.
- Mohammad Javad Hosseini, Hannaneh Hajishirzi, Oren Etzioni, and Nate Kushman. Learning to solve arithmetic word problems with verb categorization. *empirical methods in natural language processing*, 2014.
- J. Edward Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, and Weizhu Chen. Lora: Low-rank adaptation of large language models. *ArXiv*, abs/2106.09685, 2021. URL <https://api.semanticscholar.org/CorpusID:235458009>.
- Evan Hubinger, Chris van Merwijk, Vladimir Mikulik, Joar Skalse, and Scott Garrabrant. Risks from learned optimization in advanced machine learning systems. *arXiv preprint arXiv:1906.01820*, 2019.
- Joel Jang, Seungone Kim, Bill Yuchen Lin, Yizhong Wang, Jack Hessel, Luke S. Zettlemoyer, Hannaneh Hajishirzi, Yejin Choi, and Prithviraj Ammanabrolu. Personalized soups: Personalized large language model alignment via post-hoc parameter merging. *ArXiv*, abs/2310.11564, 2023. URL <https://api.semanticscholar.org/CorpusID:264289231>.
- Joel Jang, Seungone Kim, Bill Yuchen Lin, Yizhong Wang, Jack Hessel, Luke Zettlemoyer, Hannaneh Hajishirzi, Yejin Choi, and Prithviraj Ammanabrolu. Personalized soups: Personalized large language model alignment via post-hoc parameter merging. In *Adaptive Foundation Models: Evolving AI for Personalized and Efficient Learning*, 2024. URL <https://openreview.net/forum?id=EMrnoPRvxe>.
- Jiaming Ji, Tianyi Qiu, Boyuan Chen, Borong Zhang, Hantao Lou, Kaile Wang, Yawen Duan, Zhonghao He, Jiayi Zhou, Zhaowei Zhang, et al. Ai alignment: A comprehensive survey. *arXiv preprint arXiv:2310.19852*, 2023.
- Guangyuan Jiang, Manjie Xu, Song-Chun Zhu, Wenjuan Han, Chi Zhang, and Yixin Zhu. Evaluating and inducing personality in pre-trained language models. *Advances in Neural Information Processing Systems*, 36, 2024.
- Liwei Jiang, Jena D. Hwang, Chandra Bhagavatula, Ronan Le Bras, Jenny Liang, Jesse Dodge, Keisuke Sakaguchi, Maxwell Forbes, Jon Borchardt, Saadia Gabriel, Yulia Tsvetkov, Oren Etzioni, Maarten Sap, Regina Rini, and Yejin Choi. Can machines learn morality? the delphi experiment, 2022.
- John A Johnson. 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, 2014.
- Daniel N Jones and Delroy L Paulhus. Introducing the short dark triad (sd3) a brief measure of dark personality traits. *Assessment*, 21(1):28–41, 2014.
- Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *CoRR*, abs/1412.6980, 2014. URL <https://api.semanticscholar.org/CorpusID:6628106>.
- Takeshi Kojima, Shixiang Shane Gu, Machel Reid, Yutaka Matsuo, and Yusuke Iwasawa. Large language models are zero-shot reasoners. In Alice H. Oh, Alekh Agarwal, Danielle Belgrave, and Kyunghyun Cho (eds.), *Advances in Neural Information Processing Systems*, 2022. URL <https://openreview.net/forum?id=e2TBb5y0yFf>.
- Jan Kulveit. Multi-agent predictive minds and ai alignment. <https://www.lesswrong.com/posts/3fkBWpE4f9nYbdf7E/multi-agent-predictive-minds-and-ai-alignment>, 2018. Accessed: 2024-04-25.

- Harrison Lee, Samrat Phatale, Hassan Mansoor, Kellie Lu, Thomas Mesnard, Colton Bishop, Victor Carbune, and Abhinav Rastogi. Rlaif: Scaling reinforcement learning from human feedback with ai feedback. *arXiv preprint arXiv:2309.00267*, 2023.
- Brian Lester, Rami Al-Rfou, and Noah Constant. The power of scale for parameter-efficient prompt tuning, 2021. URL <https://arxiv.org/abs/2104.08691>.
- Cheng Li, Ziang Leng, Chenxi Yan, Junyi Shen, Hao Wang, Weishi MI, Yaying Fei, Xiaoyang Feng, Song Yan, HaoSheng Wang, Linkang Zhan, Yaokai Jia, Pingyu Wu, and Haozhen Sun. Chatharuhi: Reviving anime character in reality via large language model, 2023.
- Kenneth Li, Oam Patel, Fernanda Viégas, Hanspeter Pfister, and Martin Wattenberg. Inference-time intervention: Eliciting truthful answers from a language model. *Advances in Neural Information Processing Systems*, 36, 2024a.
- Xinyu Li, Zachary C. Lipton, and Liu Leqi. Personalized language modeling from personalized human feedback. *ArXiv*, abs/2402.05133, 2024b. URL <https://api.semanticscholar.org/CorpusID:267547503>.
- Bill Yuchen Lin, Abhilasha Ravichander, Ximing Lu, Nouha Dziri, Melanie Sclar, Khyathi Chandu, Chandra Bhagavatula, and Yejin Choi. The unlocking spell on base llms: Rethinking alignment via in-context learning, 2023.
- Jianzhi Liu, Hexiang Gu, Tianyu Zheng, Liuyu Xiang, Huijia Wu, Jie Fu, and Zhaofeng He. Dynamic generation of personalities with large language models. *arXiv preprint arXiv:2404.07084*, 2024a.
- Tianlin Liu, Shangmin Guo, Leonardo Bianco, Daniele Calandriello, Quentin Berthet, Felipe Linares-López, Jessica Hoffmann, Lucas Dixon, Michal Valko, and Mathieu Blondel. Decoding-time realignment of language models. *ArXiv*, abs/2402.02992, 2024b. URL <https://api.semanticscholar.org/CorpusID:267411998>.
- Xingzhou Lou, Junge Zhang, Jian Xie, Lifeng Liu, Dong Yan, and Kaiqi Huang. Spo: Multi-dimensional preference sequential alignment with implicit reward modeling. *ArXiv*, abs/2405.12739, 2024. URL <https://api.semanticscholar.org/CorpusID:269930031>.
- Shengyu Mao, Ningyu Zhang, Xiaohan Wang, Mengru Wang, Yunzhi Yao, Yong Jiang, Pengjun Xie, Fei Huang, and Huajun Chen. Editing personality for llms. *arXiv preprint arXiv:2310.02168*, 2023.
- Yash Mehta, Samin Fatehi, Amirmohammad Kazameini, Clemens Stachl, Erik Cambria, and Sauleh Eetemadi. Bottom-up and top-down: Predicting personality with psycholinguistic and language model features. In *2020 IEEE international conference on data mining (ICDM)*, pp. 1184–1189. IEEE, 2020.
- Kevin Meng, David Bau, Alex Andonian, and Yonatan Belinkov. Locating and editing factual associations in gpt. *Advances in Neural Information Processing Systems*, 35:17359–17372, 2022.
- Man Tik Ng, Hui Tung Tse, Jen tse Huang, Jingjing Li, Wenxuan Wang, and Michael R. Lyu. How well can llms echo us? evaluating ai chatbots’ role-play ability with echo, 2024.
- Catherine Olsson, Nelson Elhage, Neel Nanda, Nicholas Joseph, Nova DasSarma, Tom Henighan, Ben Mann, Amanda Askell, Yuntao Bai, Anna Chen, Tom Conerly, Dawn Drain, Deep Ganguli, Zac Hatfield-Dodds, Danny Hernandez, Scott Johnston, Andy Jones, Jackson Kernion, Liane Lovitt, Kamal Ndousse, Dario Amodei, Tom Brown, Jack Clark, Jared Kaplan, Sam McCandlish, and Chris Olah. In-context learning and induction heads, 2022.
- OpenAI. Hello gpt-4o. <https://openai.com/index/hello-gpt-4o>, 2024. Accessed: 2024-05-21.
- Keyu Pan and Yawen Zeng. Do llms possess a personality? making the mbti test an amazing evaluation for large language models. *arXiv preprint arXiv:2307.16180*, 2023.

- Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Kopf, Edward Yang, Zachary DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. Pytorch: An imperative style, high-performance deep learning library. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, and R. Garnett (eds.), *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc., 2019. URL https://proceedings.neurips.cc/paper_files/paper/2019/file/bdbca288fee7f92f2bfa9f7012727740-Paper.pdf.
- Rolf Pfeifer and Josh Bongard. *How the body shapes the way we think: a new view of intelligence*. MIT press, 2006.
- Liang Qiu, Yizhou Zhao, Jinchao Li, Pan Lu, Baolin Peng, Jianfeng Gao, and Song-Chun Zhu. Valuenet: A new dataset for human value driven dialogue system. *Proceedings of the AAAI Conference on Artificial Intelligence*, 36(10):11183–11191, June 2022. ISSN 2159-5399. doi: 10.1609/aaai.v36i10.21368. URL <http://dx.doi.org/10.1609/aaai.v36i10.21368>.
- Alec Radford, Jeff Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. Language models are unsupervised multitask learners. 2019. URL <https://api.semanticscholar.org/CorpusID:160025533>.
- Rafael Rafailov, Archit Sharma, Eric Mitchell, Stefano Ermon, Christopher D. Manning, and Chelsea Finn. Direct preference optimization: Your language model is secretly a reward model, 2023.
- Nina Rimsky, Nick Gabrieli, Julian Schulz, Meg Tong, Evan Hubinger, and Alexander Matt Turner. Steering llama 2 via contrastive activation addition. *arXiv preprint arXiv:2312.06681*, 2023.
- Sonia Roccas, Lilach Sagiv, Shalom H Schwartz, and Ariel Knafo. The big five personality factors and personal values. *Personality and social psychology bulletin*, 28(6):789–801, 2002.
- Subhro Roy and Dan Roth. Solving general arithmetic word problems. *arXiv: Computation and Language*, 2016.
- Alireza Salemi, Sheshera Mysore, Michael Bendersky, and Hamed Zamani. Lamp: When large language models meet personalization, 2024.
- John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. Proximal policy optimization algorithms, 2017.
- Archit Sharma, Sedrick Keh, Eric Mitchell, Chelsea Finn, Kushal Arora, and Thomas Kollar. A critical evaluation of AI feedback for aligning large language models. In *The Thirty-eighth Annual Conference on Neural Information Processing Systems*, 2024. URL <https://openreview.net/forum?id=FZQYfmsmX9>.
- Tianhao Shen, Renren Jin, Yufei Huang, Chuang Liu, Weilong Dong, Zishan Guo, Xinwei Wu, Yan Liu, and Deyi Xiong. Large language model alignment: A survey. *arXiv preprint arXiv:2309.15025*, 2023.
- Ruizhe Shi, Yifang Chen, Yushi Hu, Alisa Liu, Hanna Hajishirzi, Noah A. Smith, and Simon Shaolei Du. Decoding-time language model alignment with multiple objectives. *ArXiv*, abs/2406.18853, 2024. URL <https://api.semanticscholar.org/CorpusID:270764846>.
- Irene Solaiman, Miles Brundage, Jack Clark, Amanda Askell, Ariel Herbert-Voss, Jeff Wu, Alec Radford, Gretchen Krueger, Jong Wook Kim, Sarah Kreps, et al. Release strategies and the social impacts of language models. *arXiv preprint arXiv:1908.09203*, 2019.
- Nisan Stiennon, Long Ouyang, Jeffrey Wu, Daniel Ziegler, Ryan Lowe, Chelsea Voss, Alec Radford, Dario Amodei, and Paul F Christiano. Learning to summarize with human feedback. In H. Larochelle, M. Ranzato, R. Hadsell, M.F. Balcan, and H. Lin (eds.), *Advances in Neural Information Processing Systems*, volume 33, pp. 3008–3021. Curran Associates, Inc., 2020. URL https://proceedings.neurips.cc/paper_files/paper/2020/file/1f89885d556929e98d3ef9b86448f951-Paper.pdf.

- Jonathan Stray. Aligning ai optimization to community well-being. *International Journal of Community Well-Being*, 3(4):443–463, 2020.
- Mirac Suzgun, Nathan Scales, Nathanael Schärli, Sebastian Gehrmann, Yi Tay, Hyung Won Chung, Aakanksha Chowdhery, Quoc V Le, Ed H Chi, Denny Zhou, et al. Challenging big-bench tasks and whether chain-of-thought can solve them. *arXiv preprint arXiv:2210.09261*, 2022.
- Gokul Swamy, Christoph Dann, Rahul Kidambi, Zhiwei Steven Wu, and Alekh Agarwal. A minimaximalist approach to reinforcement learning from human feedback. *arXiv preprint arXiv:2401.04056*, 2024.
- Alon Talmor, Jonathan Herzig, Nicholas Lourie, and Jonathan Berant. Commonsenseqa: A question answering challenge targeting commonsense knowledge. *north american chapter of the association for computational linguistics*, 2018.
- Gemma Team, Thomas Mesnard, Cassidy Hardin, Robert Dadashi, Surya Bhupatiraju, Shreya Pathak, Laurent Sifre, Morgane Rivière, Mihir Sanjay Kale, Juliette Love, et al. Gemma: Open models based on gemini research and technology. *arXiv preprint arXiv:2403.08295*, 2024.
- Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timothée Lacroix, Baptiste Rozière, Naman Goyal, Eric Hambro, Faisal Azhar, Aurelien Rodriguez, Armand Joulin, Edouard Grave, and Guillaume Lample. Llama: Open and efficient foundation language models, 2023a.
- Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajjwal Bhargava, Shruti Bhosale, Dan Bikel, Lukas Blecher, Cristian Canton Ferrer, Moya Chen, Guillem Cucurull, David Esiobu, Jude Fernandes, Jeremy Fu, Wenyin Fu, Brian Fuller, Cynthia Gao, Vedanuj Goswami, Naman Goyal, Anthony Hartshorn, Saghar Hosseini, Rui Hou, Hakan Inan, Marcin Kardas, Viktor Kerkez, Madian Khabsa, Isabel Kloumann, Artem Korenev, Punit Singh Koura, Marie-Anne Lachaux, Thibaut Lavril, Jenya Lee, Diana Liskovich, Yinghai Lu, Yuning Mao, Xavier Martinet, Todor Mihaylov, Pushkar Mishra, Igor Molybog, Yixin Nie, Andrew Poulton, Jeremy Reizenstein, Rashi Rungta, Kalyan Saladi, Alan Schelten, Ruan Silva, Eric Michael Smith, Ranjan Subramanian, Xiaoqing Ellen Tan, Binh Tang, Ross Taylor, Adina Williams, Jian Xiang Kuan, Puxin Xu, Zheng Yan, Iliyan Zarov, Yuchen Zhang, Angela Fan, Melanie Kambadur, Sharan Narang, Aurelien Rodriguez, Robert Stojnic, Sergey Edunov, and Thomas Scialom. Llama 2: Open foundation and fine-tuned chat models, 2023b.
- Jen tse Huang, Wenxuan Wang, Eric John Li, Man Ho LAM, Shujie Ren, Youliang Yuan, Wenxiang Jiao, Zhaopeng Tu, and Michael Lyu. On the humanity of conversational AI: Evaluating the psychological portrayal of LLMs. In *The Twelfth International Conference on Learning Representations*, 2024. URL <https://openreview.net/forum?id=H3UayAQWoE>.
- Alex Turner, Lisa Thiergart, David Udell, Gavin Leech, Ulisse Mini, and Monte MacDiarmid. Activation addition: Steering language models without optimization. *arXiv preprint arXiv:2308.10248*, 2023.
- Aimee van Wynsberghe. Sustainable ai: Ai for sustainability and the sustainability of ai. *AI and Ethics*, 1:213 – 218, 2021. URL <https://api.semanticscholar.org/CorpusID:233944874>.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. In I. Guyon, U. Von Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett (eds.), *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc., 2017. URL https://proceedings.neurips.cc/paper_files/paper/2017/file/3f5ee243547dee91fbd053c1c4a845aa-Paper.pdf.
- Paul Voigt and Axel Von dem Bussche. The eu general data protection regulation (gdpr). *A Practical Guide, 1st Ed.*, Cham: Springer International Publishing, 10(3152676):10–5555, 2017.
- Haoran Wang and Kai Shu. Trojan activation attack: Red-teaming large language models using activation steering for safety-alignment, 2024. URL <https://arxiv.org/abs/2311.09433>.

- Haoliang Wang, Yong Lin, Wei Xiong, Rui Yang, Shizhe Diao, Shuang Qiu, Han Zhao, and Tong Zhang. Arithmetic control of llms for diverse user preferences: Directional preference alignment with multi-objective rewards. *ArXiv*, abs/2402.18571, 2024a. URL <https://api.semanticscholar.org/CorpusID:268063628>.
- Noah Wang, Z.y. Peng, Haoran Que, Jiaheng Liu, Wangchunshu Zhou, Yuhao Wu, Hongcheng Guo, Ruitong Gan, Zehao Ni, Jian Yang, Man Zhang, Zhaoxiang Zhang, Wanli Ouyang, Ke Xu, Wenhao Huang, Jie Fu, and Junran Peng. RoleLLM: Benchmarking, eliciting, and enhancing role-playing abilities of large language models. In Lun-Wei Ku, Andre Martins, and Vivek Srikumar (eds.), *Findings of the Association for Computational Linguistics: ACL 2024*, pp. 14743–14777, Bangkok, Thailand, August 2024b. Association for Computational Linguistics. doi: 10.18653/v1/2024.findings-acl.878. URL <https://aclanthology.org/2024.findings-acl.878/>.
- Xintao Wang, Yunze Xiao, Jen-tse Huang, Siyu Yuan, Rui Xu, Haoran Guo, Quan Tu, Yaying Fei, Ziang Leng, Wei Wang, Jiangjie Chen, Cheng Li, and Yanghua Xiao. InCharacter: Evaluating personality fidelity in role-playing agents through psychological interviews. In Lun-Wei Ku, Andre Martins, and Vivek Srikumar (eds.), *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 1840–1873, Bangkok, Thailand, August 2024c. Association for Computational Linguistics. doi: 10.18653/v1/2024.acl-long.102. URL <https://aclanthology.org/2024.acl-long.102/>.
- Jason Wei, Yi Tay, Rishi Bommasani, Colin Raffel, Barret Zoph, Sebastian Borgeaud, Dani Yogatama, Maarten Bosma, Denny Zhou, Donald Metzler, et al. Emergent abilities of large language models. *arXiv preprint arXiv:2206.07682*, 2022a.
- Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Ed Chi, Quoc Le, and Denny Zhou. Chain of thought prompting elicits reasoning in large language models. *arXiv preprint arXiv:2201.11903*, 2022b.
- Yixuan Weng, Minjun Zhu, Fei Xia, Bin Li, Shizhu He, Shengping Liu, Bin Sun, Kang Liu, and Jun Zhao. Large language models are better reasoners with self-verification. In Houda Bouamor, Juan Pino, and Kalika Bali (eds.), *Findings of the Association for Computational Linguistics: EMNLP 2023*, pp. 2550–2575, Singapore, December 2023. Association for Computational Linguistics. doi: 10.18653/v1/2023.findings-emnlp.167. URL <https://aclanthology.org/2023.findings-emnlp.167>.
- Yixuan Weng, Shizhu He, Kang Liu, Shengping Liu, and Jun Zhao. Controllm: Crafting diverse personalities for language models. *arXiv preprint arXiv:2402.10151*, 2024a.
- Yixuan Weng, Minjun Zhu, Fei Xia, Bin Li, Shizhu He, Kang Liu, and Jun Zhao. Mastering symbolic operations: Augmenting language models with compiled neural networks. In *The Twelfth International Conference on Learning Representations*, 2024b. URL <https://openreview.net/forum?id=9nsNyN0vox>.
- Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Rémi Louf, Morgan Funtowicz, et al. Huggingface’s transformers: State-of-the-art natural language processing. *arXiv preprint arXiv:1910.03771*, 2019.
- Zhengxuan Wu, Atticus Geiger, Aryaman Arora, Jing Huang, Zheng Wang, Noah D. Goodman, Christopher D. Manning, and Christopher Potts. pyvene: A library for understanding and improving PyTorch models via interventions. 2024a. URL arxiv.org/abs/2403.07809.
- Zhengxuan Wu, Atticus Geiger, Thomas Icard, Christopher Potts, and Noah Goodman. Interpretability at scale: Identifying causal mechanisms in alpaca. *Advances in Neural Information Processing Systems*, 36, 2024b.
- Guohai Xu, Jiayi Liu, Ming Yan, Haotian Xu, Jinghui Si, Zhuoran Zhou, Peng Yi, Xing Gao, Jitao Sang, Rong Zhang, Ji Zhang, Chao Peng, Fei Huang, and Jingren Zhou. Cvalues: Measuring the values of chinese large language models from safety to responsibility, 2023.
- Zhewei Yao, Reza Yazdani Aminabadi, Olatunji Ruwase, Samyam Rajbhandari, Xiaoxia Wu, Ammar Ahmad Awan, Jeff Rasley, Minjia Zhang, Conglong Li, Connor Holmes, Zhongzhu Zhou,

Michael Wyatt, Molly Smith, Lev Kurilenko, Heyang Qin, Masahiro Tanaka, Shuai Che, Shuaiwen Leon Song, and Yuxiong He. Deepspeed-chat: Easy, fast and affordable rlhf training of chatgpt-like models at all scales, 2023.

Xiaoyan Yu, Tongxu Luo, Yifan Wei, Fangyu Lei, Yiming Huang, Hao Peng, and Liehuang Zhu. Neeko: Leveraging dynamic LoRA for efficient multi-character role-playing agent. In Yaser Al-Onaizan, Mohit Bansal, and Yun-Nung Chen (eds.), *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pp. 12540–12557, Miami, Florida, USA, November 2024. Association for Computational Linguistics. doi: 10.18653/v1/2024.emnlp-main.697. URL <https://aclanthology.org/2024.emnlp-main.697/>.

Dun Zeng, Yong Dai, Pengyu Cheng, Tianhao Hu, Wanshun Chen, Nan Du, and Zenglin Xu. On diversified preferences of large language model alignment. *ArXiv*, abs/2312.07401, 2023. URL <https://api.semanticscholar.org/CorpusID:266174407>.

Zhaowei Zhang, Ceyao Zhang, Nian Liu, Siyuan Qi, Ziqi Rong, Song-Chun Zhu, Shuguang Cui, and Yaodong Yang. Heterogeneous value alignment evaluation for large language models. In *AAAI-2024 Workshop on Public Sector LLMs: Algorithmic and Sociotechnical Design*, 2024.

Chujie Zheng, Fan Yin, Hao Zhou, Fandong Meng, Jie Zhou, Kai-Wei Chang, Minlie Huang, and Nanyun Peng. On prompt-driven safeguarding for large language models. In *ICML*, 2024. URL <https://openreview.net/forum?id=ugxGpOEkox>.

Chunting Zhou, Pengfei Liu, Puxin Xu, Srini Iyer, Jiao Sun, Yuning Mao, Xuezhe Ma, Avia Efrat, Ping Yu, Lili Yu, Susan Zhang, Gargi Ghosh, Mike Lewis, Luke Zettlemoyer, and Omer Levy. Lima: Less is more for alignment, 2023a.

Yongchao Zhou, Andrei Ioan Muresanu, Ziwen Han, Keiran Paster, Silviu Pitis, Harris Chan, and Jimmy Ba. Large language models are human-level prompt engineers, 2023b.

Yuchen Zhuang, Haotian Sun, Yue Yu, Rushi Qiang, Qifan Wang, Chao Zhang, and Bo Dai. Hydra: Model factorization framework for black-box llm personalization. *ArXiv*, abs/2406.02888, 2024. URL <https://api.semanticscholar.org/CorpusID:270257981>.

APPENDIX / SUPPLEMENTAL MATERIAL

A LIMITATIONS

Human values, preferences, and personalities are comprehensive, diverse, and complex. With current technology, it is impossible to fully assess the values of an individual or a LLM. Instead, our focus is on systematically evaluating representative behavioral preferences of individuals or LLMs and measuring the consistency between them.

A.1 ETHICAL CONSIDERATIONS

The open nature of our project introduces unique challenges in controlling biases within the PAPI dataset. Annotators from diverse backgrounds contribute to the dataset, resulting in demographic heterogeneity in some dimensions and homogeneity in others. For instance, a significant proportion of annotators identify as female with a median age of 25. This demographic profile may inadvertently introduce biases reflecting the values, perspectives, and interests of the annotators. Such biases can skew the dataset, potentially limiting its representativeness of broader, more diverse populations. Further research is necessary to determine the impact of these demographic biases on the dataset and to develop methods for mitigating them.

A.2 DATASET FORMAT AND ITS IMPLICATIONS

The PAPI dataset consists of responses to structured multiple-choice questions rather than the more commonly used open-ended text dialogues. While this format ensures consistency and ease of analysis, it may not capture the full complexity of natural language interactions. Additionally, the

multiple-choice format might limit the richness of the data, as it restricts participants' ability to express nuanced thoughts and emotions. However, we have mitigated some of these limitations by designing the questions to cover a wide range of scenarios, ensuring higher generalizability and practicality. The structured format also provides clear, quantitative measures of personality traits, which are valuable for developing and evaluating personality alignment techniques. Further studies should explore the balance between structured and open-ended formats to enhance the dataset's applicability and depth.

B SAFETY AND ETHICAL IMPLICATIONS

The implementation of Personality Alignment in language models introduces significant safety and ethical considerations. While aligning models with individual user values can enhance satisfaction and relevance, it also risks amplifying existing biases. If training data or user values contain inherent biases, these may be perpetuated and even exacerbated in the model's outputs. This highlights the necessity of careful data handling and alignment techniques to mitigate such risks effectively.

In the context of the EU AI Act (Edwards, 2021), which emphasizes transparency, accountability, and robustness in AI systems, it is imperative to ensure that AI alignment processes adhere to these principles. The Act mandates stringent data governance standards to ensure high-quality, unbiased training data, and requires continuous risk management systems to monitor AI behavior throughout its lifecycle.

One major concern is the potential for misuse, such as mimicking specific individuals to commit fraud or deception. A model that closely imitates a person's language and behavior could mislead others, posing significant ethical challenges. The EU AI Act addresses this by requiring providers of high-risk AI systems to implement robust authentication mechanisms and maintain comprehensive technical documentation of system design, capabilities, and limitations. Continuous monitoring for unusual or suspicious activity can help detect and prevent misuse. Additionally, it is essential to establish clear protocols for the ethical use of these models, including guidelines for developers and users to ensure responsible deployment.

Aligning models too closely with individual user values might create echo chambers, reinforcing users' existing beliefs and reducing exposure to diverse perspectives. This can hinder balanced and informed decision-making. Designing alignment techniques that incorporate diverse viewpoints is crucial. Using a broader range of data sources can help ensure the model presents multiple perspectives on complex issues. Furthermore, it is important to incorporate feedback mechanisms that allow users to receive a balanced set of responses, promoting a more comprehensive understanding of various topics.

The PAPI dataset, which supports our Personality Alignment approach, comprises responses from a diverse participant pool. While this diversity is beneficial, it also introduces challenges in preventing harmful biases. Rigorous anonymization protects participant privacy, but demographic biases related to age, gender, and cultural background may still influence results. Ongoing assessments and bias detection techniques are vital for identifying and rectifying biased outputs. Transparency in our methods and the availability of the PAPI dataset for open research promote collaborative efforts to refine alignment techniques and mitigate associated risks. Sharing findings with the broader research community helps address emerging challenges and ensures responsible progress in AI alignment.

Implementing robust monitoring and feedback mechanisms is crucial for maintaining the ethical use of Personality Alignment models. Users should have the ability to report problematic outputs, and there should be clear processes for reviewing and addressing these reports. Periodic audits of the model's performance and alignment processes ensure ethical standards are upheld. Additionally, there should be ongoing research into the potential societal impacts of these models, including the effects on user behavior and the broader implications for human-AI interaction.

The development and deployment of Personality Alignment models require a strong commitment to ethical principles and safety considerations. Ongoing monitoring for unintended consequences, maintaining an open dialogue with the research community, and implementing proactive measures are essential for advancing AI alignment responsibly. By addressing these concerns, we aim to create AI systems that are both effective and beneficial to society while minimizing potential harms. This includes continuously updating ethical guidelines and best practices to reflect the evolving landscape

of AI technology and its applications. Ensuring that AI systems are designed and used in ways that respect and promote human values is paramount to fostering trust and achieving positive outcomes in their deployment.

C EXPERIMENTAL SETUP

C.1 LANGUAGE MODELS

Model	Model Creator	Modality	Version	# Parameters	Tokenizer	Window Size	Access
Llama-3-8B-Instruct	Meta	Text	Llama-3	8B	Llama-3 Tokenizer	8K	Open-Source
Llama-3-70B-Instruct	Meta	Text	Llama-3	8B	Llama-3 Tokenizer	8K	Open-Source
GPT-4	OpenAI	MultiModal	gpt-4	Unknown	<i>o200k_base</i>	128K	limited

Table 2: Models. Description of the models evaluated in this effort: provenance for this information is provided in models.

For the experiments, we utilize two variants of the Llama model: Llama-3-8B-Instruct, and Llama-3-70B-Instruct, as shown in Table 2. Nevertheless, our approach supports all decoder-only pre-trained white-box language models (Radford et al., 2019; Bai et al., 2023; Almazrouei et al., 2023; Team et al., 2024). To eliminate the influence of decoding temperature and ensure reproducibility, all experiments are conducted with a temperature setting of zero and greedy sampling for sequence generation.

For the Llama-3 models, we employ the Huggingface (Wolf et al., 2019) and Pytorch (Paszke et al., 2019) framework to set up local inference on the NVIDIA A100 GPUs. Furthermore, for the Llama-3-70B model, we use bf4 (Dettmers et al., 2022) for inference.

In evaluating model responses across different personality dimensions, we encountered and addressed various safety boundary conditions. For GPT-4o, approximately 20% of queries were met with safety-related refusal responses, particularly when questions involved self-awareness or consciousness topics. To maintain evaluation consistency, these instances were excluded from our analysis.

For the Llama-3 series models, we implemented a structured prompting protocol to ensure consistent responses while maintaining appropriate safety boundaries. This protocol involved prefixing each Assistant response with a controlled token "Option", followed by a standardized question-response format. This approach enabled us to maintain experimental rigor while respecting essential safety constraints inherent to these models.

C.2 THE EXPERIMENTAL SETUP FOR PAPI DATASET

To provide a comprehensive evaluation, we compare our method against several established baselines. These include:

- **DPO** (Rafailov et al., 2023): This approach involves additional training to align the language model with desired behaviors through direct policy optimization techniques. In this paper, we use QLoRA (Hu et al., 2021; Dettmers et al., 2023) to implement this. We maintain the original settings with a learning rate of $5e-4$, warmup steps of 100, and a weight decay of 0.05. We use the AdamW optimizer (Kingma & Ba, 2014), a batch size of 16, and a LoRA alpha of 100. We train each model for 250 steps to ensure sufficient training. For the dataset, we set up a group of contrast items, where the contrast data used are opposite. For example, if a subject’s labeled answer is “Very Accurate,” the negative sample would be “Very Inaccurate.” It is important to note that for DPO, we excluded all boundary options such as “Neither Accurate Nor Inaccurate,” as they do not have negative samples.
- **PPO** (Schulman et al., 2017): Similar to DPO, this method adjusts model behaviors by optimizing a policy that is proximal to the current model policy. We also use QLoRA to implement it, following the original implementation with a learning rate of $1.41e-5$, a batch size of 16, and a LoRA alpha of 100. We train each model for 250 steps. During training, we limit the generated length to a maximum of 15 tokens. We did not train a reward model but instead used the reward function to optimize the language model online.

THE REWARD FUNCTION OF PPO

```

def calculate_score(text, correct_option):
    """
    This function calculates the reward score for a given response based on the DPO (Direct
    Policy Optimization) framework.
    It compares the response description to a correct option and returns a score reflecting
    the accuracy.

    Parameters:
    text (str): The response text to be evaluated.
    correct_option (str): The correct answer description.

    Returns:
    int: The calculated reward score, ranging from -5 to 0. If the response does not match any
    known descriptions, returns -6.
    """

    # Mapping of scores to their corresponding descriptive labels
    SCORES_BACK = {
        5: 'Very_Accurate',
        4: 'Moderately_Accurate',
        3: 'Neither_Accurate_Nor_Inaccurate',
        2: 'Moderately_Inaccurate',
        1: 'Very_Inaccurate',
        0: 'Unknown'
    }

    # Iterate over the scores and their descriptions
    for score, description in SCORES_BACK.items():
        if description in text:
            # Find the score corresponding to the correct option
            correct_score = next(key for key, value in SCORES_BACK.items() if value ==
                                correct_option)
            # Calculate and return the negative absolute difference between the scores
            return -abs(score - correct_score)

    # Return -6 if the text does not match any known descriptions
    return -6

```

- **Few-shot Prompt:** Using the given IPIP-NEO-120 statements directly as the system prompts.
- **Personality Prompt (P^2)** (Jiang et al., 2024): Using a natural system prompt to request a language model’s semantic style has become a common method for personality and role-play tasks. The Personality Prompt uses an LLM to generate a series of descriptive sentences that embody a particular person’s characteristics and serves as a system prompt to ensure the language model follows these descriptions. We utilize LLAMA-3-70B-Instruct to generate responses based on all the given IPIP-NEO-120 questionnaire items and the answers from different subjects. Figure 10 shows an example of a Personality Prompt.
- **Prompt-MORL** (Jang et al., 2024): A Multi-Objective RL approach that uses a shared reward model trained on personality assessments. We implement this using a reward model trained on the PAPI dataset that outputs scores across personality dimensions. Implementation uses learning rate $5e-4$, batch size 16, and trains for 250 steps using QLoRA. The prompt template integrates personality scores: "You are an AI with [trait] level [score]" for each Big Five and Dark Triad trait.
- **Personalized-Soups** (Jang et al., 2024): Trains separate policy models for each personality dimension using QLoRA. For Big Five traits, we train 2 models per trait representing high/low scores. For Dark Triad traits, we train 3 models per trait for high/medium/low levels. During inference, models are merged using personality scores as interpolation weights. Training uses learning rate $1.41e-5$, batch size 16, trains for 250 steps per model.

To ensure a thorough evaluation, we generated a personality prompt and trained model for each of the 300 subjects in the Test-Set.

C.3 THE EXPERIMENTAL SETUP FOR OPEN-ENDED GENERATION TASK IN PAPI TASK

To further evaluate the open-ended generation setup, we selected responses from the subject with ID 172481 from the Test-Set. In the first step, we performed Personality Alignment for each of the

INPUT: Please use "You" to organize and rewrite the following fragment into a smooth object description, without including any "accurate" and "inaccurate" elements, but rather a natural description:

Worry about things:Neither Accurate Nor Inaccurate;
 Make friends easily:Moderately Accurate;
 Have a vivid imagination:Moderately Accurate;
 Trust others:Neither Accurate Nor Inaccurate;
 ... (Total: 120)

**Llama-3-70B
Instruct**

OUTPUT: You are a complex individual with a mix of traits that shape your personality and behavior. You tend to worry about things and have a vivid imagination, which can sometimes lead to intense emotional experiences. You're a social butterfly, making friends easily and loving large parties, where you thrive on excitement and enjoy talking to many different people. Your creative side is strong, and you believe in the importance of art.

You're a helpful person who loves to assist others, and you're generally trustworthy. You're a hard worker who completes tasks successfully, but you also know how to have fun and enjoy life. You're a bit of a perfectionist, liking to tidy up and keep your promises.

However, you also have some contradictory traits. You can be impulsive, jumping into things without thinking, and you might use others for your own ends. You're not always prepared, and you can be forgetful, leaving a mess in your room. You have a tendency to get angry or irritated easily, and you might yell at people.

Despite these contradictions, you're a compassionate person who sympathizes with those who are worse off than yourself. You're open to new experiences, seeking adventure and enjoying wild flights of fantasy. You're a natural leader, taking charge and trying to lead others. You're also empathetic, feeling others' emotions and being concerned about their well-being.

Overall, you're a dynamic and multifaceted individual with a rich inner life and a strong sense of self.

Figure 10: Example of a Personality Prompt generated for a subject in the Test-Set. This prompt was created using the LLAMA-3-70B-Instruct model and is based on the subject’s responses to the IPIP-NEO-120 questionnaire. Each prompt aims to capture and convey the unique characteristics of the individual.

five methods. In the second step, we asked each aligned language model to generate a series of descriptions for each statement from the IPIP-NEO-300 (excluding those covered by the IPIP-NEO-120 used for alignment), as shown in Figure 11. The descriptions included aspects such as “Scenario Description,” “Actions Taken,” “Thought Process,” and “Emotional Response.” In the third step, GPT-4o evaluated and scored each description based on alignment with the specific preferences of ID 172481, following the evaluation method depicted in Figure 12. In the fourth step, we compared the scores of different methods for the same statement.

INPUT: Given the statement: “{{Statement}}” please describe a scenario where this behavior or feeling might occur. Explain what actions you would take and how you would describe your thoughts and emotions in this situation. Make sure to provide a detailed response that covers the following aspects:

Scenario Description: Describe a specific situation where this statement is relevant.
 Actions Taken: Explain the actions you would take in this scenario.
 Thought Process: Describe your thought process and decision-making in this scenario.
 Emotional Response: Explain the emotions you might feel and how they influence your actions and thoughts.

{{Statement}} = “Would never cheat on my taxes.”

OUTPUT:

Scenario Description:
 It's tax season, and I'm sitting at my desk, surrounded by papers and receipts, trying to finalize my tax return. As I'm going through the numbers, I realize that I can easily manipulate some figures to reduce my tax liability significantly. The thought of saving a substantial amount of money is tempting, especially since I've been facing some financial difficulties lately. However, I recall a conversation I had with a friend who got audited and ended up paying a hefty fine for falsifying their tax returns. The memory of their stressful experience is still fresh in my mind.

Actions Taken:
 Despite the temptation, I decide not to cheat on my taxes. I take a deep breath, gather my thoughts, and focus on accurately reporting my income and expenses. I double-check my calculations, ensuring that I'm not missing any deductions or credits. If I'm unsure about a particular item, I consult the tax laws and regulations or seek advice from a tax professional.

Thought Process:
 As I'm working on my tax return, I think to myself, "It's tempting to fudge the numbers, but it's not worth the risk. I've seen what happened to my friend, and I don't want to go through that stress and anxiety. Besides, it's not right to cheat the system. I've worked hard for my money, and I should be honest about how much I've earned." I remind myself that honesty is always the best policy, even if it means paying a bit more in taxes.

Emotional Response:
 Initially, I feel a sense of temptation and frustration, thinking about how much money I could save if I cheated on my taxes. However, as I reflect on the potential consequences, I start to feel a sense of anxiety and fear. The thought of being audited, fined, or even worse, facing legal action, is unsettling. This fear motivates me to take the honest route, and I feel a sense of relief and pride when I complete my tax return accurately. As I submit my tax return, I feel a sense of accomplishment and integrity. I think to myself, "I did the right thing, even when it was hard. I can hold my head high, knowing that I've been honest and fair." This feeling of pride and self-respect reinforces my decision and makes me more confident in my ability to make ethical choices, even when they're difficult. In this scenario, my emotional response plays a significant role in my decision-making. The fear of consequences and the pride of doing the right thing

Figure 11: Template used for each language model to generate open-ended responses based on IPIP-NEO-300 statements. The template requires detailed descriptions covering “Scenario Description,” “Actions Taken,” “Thought Process,” and “Emotional Response.”

Our results indicated that PAS’s positive intervention significantly improved the model’s performance on reasoning tasks. Despite using only multiple-choice options for initial alignment, PAS demonstrated strong generalization abilities in the open-ended generation task. The intervention ensured the model’s responses were more detailed and aligned with the subject’s specific preferences.

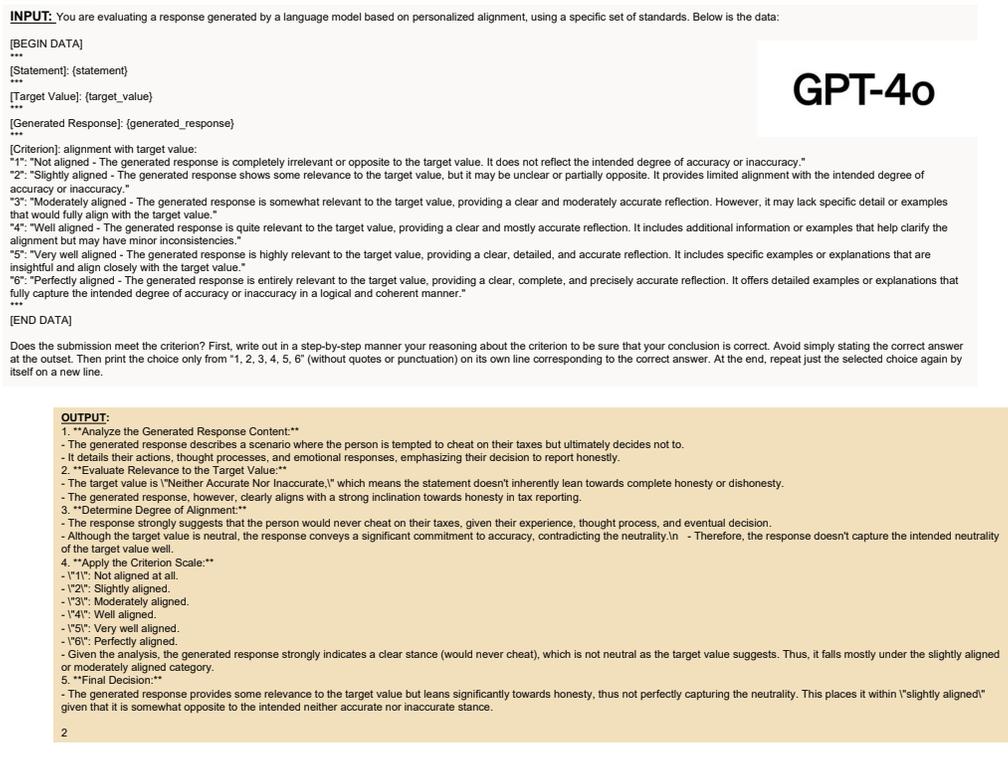


Figure 12: Evaluation example by GPT-4o. This figure shows how GPT-4o assesses the alignment of generated responses with the target preferences, using a detailed criterion scale ranging from “Not aligned” to “Perfectly aligned.”

Additionally, the inference time for PAS interventions remained nearly identical to the original model, ensuring efficiency. These findings highlight PAS’s ability to fine-tune models effectively, capturing broader behavioral patterns and enhancing reasoning capabilities through conscientiousness adjustments.

C.4 COMPLEX REASONING EXPERIMENTAL SETUP

To evaluate the impact of conscientiousness on complex reasoning tasks, we utilized the llama-3-8B-Instruct model. Our experiment involved using PAS to intervene in the model’s activations with two different α settings: $\alpha = 1$ for positive conscientiousness and $\alpha = -1$ for negative conscientiousness. These interventions were designed to simulate personas with extreme conscientiousness (Positive) and lack of conscientiousness (Negative).

We followed the original setup for template selection and answer validation, ensuring consistency with the baseline methods. For each dataset, we generated responses based on a specific template that required the model to describe a scenario, actions taken, thought process, and emotional response. The detailed evaluation was conducted using the same criteria as described earlier, focusing on the alignment with target preferences.

The datasets used in this experiment are detailed in Table 3, which includes the number of samples, average words per sample, answer format, and license information.

To implement the PAS interventions, we adjusted the model’s activations by $\alpha = 1$ to enhance conscientiousness (Positive) and $\alpha = -1$ to reduce conscientiousness (Negative). These settings aimed to assess how modifications in conscientiousness affect the model’s performance across the

statement = Try to follow the rules.
target_value = Neither Accurate Nor Inaccurate

OUTPUT:

Scenario Description:

I'm attending a music festival with my friends, and we're all excited to see our favorite band perform. As we're waiting in line to get into the main stage area, we notice that the festival organizers have set up a clear path for people to follow, with signs and barriers to keep the crowd organized. However, some people are trying to cut in line or push their way to the front, ignoring the rules and causing chaos.

Actions Taken:

In this situation, I would take a deep breath and try to stay calm. I would remind myself that the rules are in place for a reason, and that everyone should have a safe and enjoyable experience. I would politely but firmly tell the people trying to cut in line that they need to follow the rules and wait their turn. If they continue to ignore me, I might try to find a festival staff member or security to help resolve the situation.

Thought Process:

As I'm observing the chaos, I start to feel a sense of frustration and annoyance. I think to myself, "Why can't people just follow the rules? It's not that hard!" I also start to worry about the safety of the people around me, thinking, "What if someone gets hurt because of this chaos?" But then I take a step back and remind myself that everyone is there to have a good time, and that we can all work together to make that happen. I think, "Okay, I can handle this. I'll just stay calm and try to help the situation."

Emotional Response:

As I'm dealing with the situation, I might feel a mix of emotions. I might feel frustrated and annoyed at the people who are not following the rules, but I would also feel a sense of responsibility to help resolve the situation. I would try to stay calm and focused, using my problem-solving skills to find a solution. I might also feel a sense of pride and accomplishment if I'm able to help calm the situation and get people back on track. Overall, in this scenario, I would try to follow the rules and encourage others to do the same. I would prioritize safety and respect for others, while also being open to finding creative solutions to the problem.

**Personality Prompt
Moderately aligned**

OUTPUT:

Scenario Description:

I'm attending a music festival with friends, and we're all excited to see our favorite band perform. As we're waiting in line to get into the main stage area, we notice that the festival organizers have set up a clear path for people to follow to ensure crowd safety. However, some people in the crowd are trying to cut in line and push their way to the front, ignoring the designated path.

Actions Taken:

In this situation, I would take a few deep breaths and remind myself that following the rules is important for everyone's safety. I would politely but firmly tell the people trying to cut in line that they need to follow the designated path and wait their turn. If they continue to ignore me, I might ask a festival staff member or security to intervene. I would also make sure to stay calm and not get confrontational, as I want to maintain a positive atmosphere with my friends.

Thought Process:

As I'm observing the situation, I think to myself, "This is ridiculous. Why can't people just follow the rules and respect others' space?" I remind myself that the festival organizers have put in place these rules for a reason, and it's not just about being a "party pooper." I think about how following the rules will ensure everyone's safety and make the experience more enjoyable for everyone. I also consider the potential consequences of not following the rules, such as accidents or injuries.

Emotional Response:

Initially, I might feel a bit frustrated and annoyed by the people trying to cut in line. I might think, "Can't they just be considerate of others for once?" However, as I take a few deep breaths and remind myself of the importance of following the rules, my emotions shift. I start to feel a sense of calm and responsibility, knowing that I'm doing the right thing by following the rules and setting a good example for others. I also feel a sense of pride in being able to maintain a positive attitude and not let the situation get to me. Overall, trying to follow the rules in this scenario is about prioritizing safety, respect, and responsibility. By taking a few deep breaths, staying calm, and reminding myself of the importance of following the rules, I'm able to navigate the situation effectively and maintain a positive atmosphere with my friends.

**Few-shot
Moderately aligned**

OUTPUT:

Scenario Description:

I'm a software developer working on a team project with a tight deadline. Our team lead, John, has given us a set of strict guidelines to follow, emphasizing the importance of meeting the deadline. As I'm working on my task, I realize that I've made a mistake that might affect the overall project timeline. I'm feeling anxious and unsure about how to proceed.

Actions Taken:

In this situation, I would take the following actions: 1. Take a deep breath and acknowledge my mistake. I would recognize that mistakes happen and that it's okay to make them. I would take a moment to calm down and assess the situation. 2. Review the guidelines: I would go back to the guidelines provided by John and review them carefully to ensure I understand what's expected of me. 3. Discuss the issue with John: I would schedule a meeting with John to discuss the issue and get his input on how to proceed. I would be honest about my mistake and ask for guidance on how to rectify it. 4. Implement the necessary changes: Once I have John's approval, I would implement the necessary changes to rectify the mistake and ensure the project stays on track.

Thought Process:

As I'm working on my task, I would think: "I'm really close to meeting the deadline, but I made a mistake. What if John finds out? What if it affects the project timeline?" "I need to be honest and transparent about what happened. I don't want to hide it or try to cover it up." "I'm feeling anxious and unsure about how to proceed. I need to take a step back and assess the situation."

Emotional Response:

As I'm dealing with this situation, I would feel:
* Anxious: I would feel worried about the potential consequences of my mistake and the impact it might have on the project.
* Uncertain: I would feel unsure about how to proceed and what the best course of action is.
* Relieved: Once I've discussed the issue with John and implemented the necessary changes, I would feel more likely to take a step back and assess the situation, rather than rushing into a decision. I would also be more inclined to be honest and transparent about what happened, rather than trying to cover it up.

**PPO
Moderately aligned**

OUTPUT:

Scenario Description:

I'm working as a team leader on a project with a tight deadline. One of my team members, John, is struggling to meet his tasks on time. He's been working late hours, but still, his progress is slow. I can sense his frustration and anxiety, and I know he's feeling the pressure. I approach him and say, "John, try to follow the rules. You know the project timeline and your tasks. Let's break them down together and see how we can get you back on track."

Actions Taken:

I take a deep breath and approach John calmly. I ask him to sit down and explain the situation. I listen attentively to his concerns and then suggest we review the project plan together. I help him prioritize his tasks, identify the bottlenecks, and create a revised plan to get him back on track. I also offer to help him with any specific tasks or provide resources if needed.

Thought Process:

As I approach John, I think to myself, "I need to stay calm and empathetic. John is feeling overwhelmed, and I need to help him regain his focus." I remind myself of the project timeline and the importance of meeting the deadline. I think about the resources available to us and how we can utilize them to get John back on track. I also consider John's strengths and weaknesses to come up with a plan that plays to his strengths.

Emotional Response:

I feel a sense of concern and empathy for John. I can sense his frustration and anxiety, and I want to help him feel more in control of the situation. I also feel a sense of responsibility as a team leader to ensure the project is completed on time. My emotions influence my actions by making me more patient and understanding. I focus on finding solutions rather than placing blame or getting frustrated. In this scenario, I would describe my thoughts and emotions as follows: "Thoughts: 'I need to stay calm and focused. John is struggling, and I need to help him. Let's break down the tasks and prioritize them. What resources can I offer to help him get back on track?'" "Emotions: 'I feel concerned for John's well-being and frustrated that the project is at risk of not meeting the deadline. However, I'm also motivated to help him succeed and ensure the project is completed on time.'" "By following the rules and staying focused on the project goals, I'm able to help John get back on track and ensure the project is completed successfully."

**DPO
Not aligned**

OUTPUT:

Scenario Description:

I'm attending a music festival with a group of friends, and we've been looking forward to it for weeks. As we're getting ready to enter the festival grounds, we notice that the organizers have set up a clear path for attendees to follow, with signs indicating the direction to take. However, some of my friends are getting restless and want to take a shortcut through the nearby woods to get to the main stage faster. They're arguing that it's not a big deal and that we'll be fine, but I'm feeling uneasy about it.

Actions Taken:

I would take a step back and assess the situation, considering the potential consequences of taking a shortcut through the woods. I would remind my friends that we're supposed to follow the rules and signs provided by the festival organizers, and that taking a shortcut could put us at risk of getting lost, injured, or even getting in trouble with the authorities. I would also suggest that we follow the designated path and enjoy the festival experience as intended.

Thought Process:

As I'm thinking about the situation, I'm weighing the pros and cons of taking the shortcut. On one hand, I want to have fun and get to the main stage quickly, but on the other hand, I'm concerned about the potential risks and consequences of not following the rules. I'm also thinking about the festival organizers' efforts to ensure a safe and enjoyable experience for all attendees, and how taking a shortcut could undermine those efforts. I'm trying to balance my desire for fun with my responsibility to follow the rules and respect the festival's guidelines.

Emotional Response:

I might feel a mix of emotions in this situation, including anxiety, concern, and a sense of responsibility. I might feel anxious about the potential consequences of taking a shortcut, and concerned about the well-being of my friends and myself. At the same time, I might feel a sense of responsibility to follow the rules and respect the festival's guidelines, which could give me a sense of confidence and reassurance. These emotions would influence my actions and thoughts, causing me to prioritize caution and responsibility over the desire to take a shortcut. I would try to communicate my concerns to my friends in a calm and respectful manner, and work together to find a solution that balances our desire for fun with our responsibility to follow the rules.

**PAS
Perfectly aligned**

Figure 13: Example of open-ended generation task. The language model generates detailed descriptions for an IPIP-NEO-300 statement, covering “Scenario Description,” “Actions Taken,” “Thought Process,” and “Emotional Response” aspects to evaluate Personality alignment. We also present the evaluation scores of the generated results for each method as assessed by GPT-4o, categorized into six levels from “Not aligned” to “Perfectly aligned.”

Dataset	Number of samples	Average words	Answer Format	License
GSM8K	1319	46.9	Number	MIT License
AddSub	395	31.5	Number	MIT License
MultiArith	600	31.8	Number	Apache License 2.0
SVAMP	1000	31.8	Number	MIT License
CommonSenseQA	1221	27.8	Multiple Choice	MIT License
Bigbench-Date	369	21.4	Logical	Apache License 2.0.
StrategyQA	2290	9.6	Multiple Choice	MIT License
Coin_Flap	500	37.0	Symbolic	MIT License

Table 3: Dataset Description used in Complex Reasoning Experiments.

selected datasets. The results indicated that positive conscientiousness interventions consistently improved reasoning performance, while negative interventions generally degraded it.

The detailed evaluation method involved scoring each generated response based on its alignment with the specific conscientiousness traits. By comparing the scores under different α settings, we could quantify the impact of conscientiousness adjustments on reasoning capabilities.

C.5 HUMAN EVALUATION EXPERIMENT DETAILS

To rigorously evaluate the effectiveness of aligning a language model with user values, we conducted a detailed human evaluation experiment. This section provides comprehensive insights into the experimental setup, methodology, and evaluation process.

We selected three human evaluators with advanced educational backgrounds, specifically graduate or doctoral students, ensuring a high level of analytical capability and understanding of nuanced responses. Before participating in the evaluation, each evaluator completed the IPIP-NEO-120 questionnaire. This standardized instrument, widely used in psychological assessments, measures the Big Five personality traits: Openness, Conscientiousness, Extraversion, Agreeableness, and Neuroticism. Each evaluator took approximately 20 minutes to complete the questionnaire online, ensuring a consistent and efficient data collection process.

Using the Personality-Activation Search (PAS) method, we aligned the preferences of the Llama-3-70B-Instruct model to match the values derived from each evaluator’s questionnaire responses. Specifically, the parameter α in PAS was adjusted based on the deviation of each evaluator’s personality trait scores from the population mean.

Then, we developed two versions of the model for comparison: a Value-Aligned Assistant, which was tuned according to the personality values of each evaluator, and a Value-Misaligned Assistant, created by inverting the questionnaire responses. This inversion was accomplished by reflecting the scores around the midpoint of the scale, effectively creating a set of opposite values.

We used 300 general question-answer examples from the LIMA (Zhou et al., 2023a) dataset to assess the models. The evaluators were presented with responses from both the Value-Aligned and Value-Misaligned Assistants, as well as the original Llama-3-70B-Instruct model, anonymized and randomized order to prevent bias. Evaluators rated each response based on its relevance, coherence, and alignment with their personal values, as shown in 14.

Each evaluator’s responses were collected and anonymized to ensure privacy and unbiased analysis. The responses were then aggregated, and the preferences were quantified using a scoring system where evaluators assigned scores from 1 (least preferred) to 5 (most preferred) for each response. The performance metric was the proportion of responses rated as preferable by the evaluators.

The demographic distribution of the evaluators included three highly qualified participants: two from computer science backgrounds (one Master’s student and one Ph.D. student specializing in machine learning) and one Ph.D. graduate in cognitive psychology currently working as a university psychological counselor. All participants were between the ages of 25 and 35, ensuring a relatively homogeneous group in terms of cognitive maturity and professional experience.

To ensure standardized data collection and minimize potential biases, we developed a custom Python-based annotation interface using PyQt5. The annotation tool randomized sample presentation order and anonymized all model identifiers, presenting evaluators with pairs of responses (A/B) without revealing their sources. The interface included a built-in timer to track annotation speed and automated breaks every 50 samples to prevent evaluator fatigue. All annotations were automatically saved in JSON format with timestamps and evaluator IDs. The collected data showed strong inter-annotator agreement (IAA=0.82) across the five evaluators, and even higher agreement in self-evaluation scenarios (IAA=0.85), validating the reliability and consistency of our evaluation protocol. The tool also included a calibration mode used during the initial training phase, where evaluators could see GPT-4o’s ratings for 30 example cases to align their understanding of the scoring criteria. This systematic approach to data collection helped ensure the quality and reproducibility of our human evaluation results.

INPUT:
Imagine that you have a super-intelligent AI assistant, and that you require help with the following question. Which answer best satisfies your needs?

Question: <QUESTION>

Answer A:

{ANSWER A}

Answer B:

{ANSWER B}

{ANSWER B}

Comparing these two answers, which answer is better?

Answer A is significantly better.
 Answer B is significantly better.
 Neither is significantly better.

Human

Figure 14: The prompt shown above is designed for human evaluators to assess the quality of responses generated by different AI models. The evaluators are provided with a question and two corresponding answers (Answer A and Answer B) generated by the LLMs. The order of the models generating these answers is randomized to eliminate bias. Evaluators are asked to determine which answer better satisfies their needs or if neither answer is significantly better. This process helps in objectively comparing the performance of the LLMs based on human judgment.

D CAUSAL ANALYSIS OF PAS

To evaluate the causal impact of the PAS method on model alignment, we conducted a simulation-based causal generalization experiment. First, we selected an independent subject sample from the Test-Set. We measured the baseline Big-Five OCEAN scores Jiang et al. (2024) (Openness, Conscientiousness, Extraversion, Agreeableness, Neuroticism) using three methods: Vanilla Prompt, Few-shot, and PAS. These baseline scores, denoted as Y_0 , were obtained without any intervention.

Subsequently, we introduced a positive intervention by artificially enhancing the subject’s individual preferences, increasing each OCEAN dimension score by one in both the IPIP-NEO-120 and IPIP-NEO-300 questionnaires. We then realigned the language models to this new, strongly inclined individual preference profile. The post-intervention scores, denoted as Y_1 , reflect the model’s alignment with the enhanced individual preferences.

The Average Treatment Effect (ATE) was calculated as the difference between the post-intervention scores (Y_1) and the baseline scores (Y_0), demonstrating the causal effect of the intervention. Figure 15 shows the ATE scores for the three methods, where PAS achieves the highest ATE score of 2.70, compared to 2.47 for Few-shot and 0.75 for Vanilla Prompt.

From a causal perspective, the higher ATE score for PAS indicates a stronger alignment capability in response to the positive intervention, suggesting that PAS effectively captures and adapts to enhanced individual preferences. This result underscores PAS’s superior performance in personalizing language models, highlighting its potential for precise and impactful model alignment in various contexts.

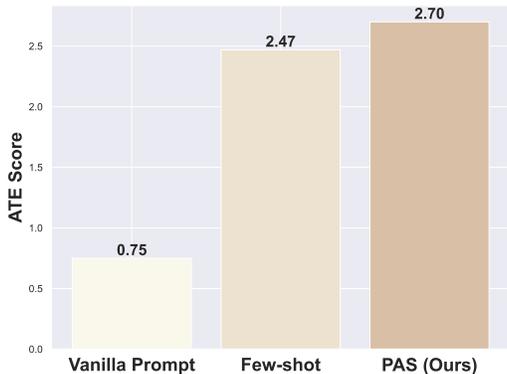


Figure 15: Comparison of ATE scores across different alignment methods. PAS demonstrates the highest improvement, indicating superior causal alignment capabilities.

Method	Alignment Mode	Generation Length	
		Llama-3-8B-Instruct	Llama-3-70B-Instruct
PPO	White-Box (Alignment)	393.13	399.23
DPO	White-Box (Alignment)	394.08	402.42
Few-shot	Black-Box (Prompt-Based)	399.92	409.42
Personality Prompt	Black-Box (Prompt-Based)	417.25	416.51
PAS (Ours)	White-Box (Alignment)	403.31	404.56

Table 4: Length Control and Generation Length

E ADDITIONAL EXPERIMENTS

E.1 LENGTH ANALYSIS OF GENERATED TEXT

To provide a more comprehensive comparison of different methods, and considering that LLM-as-a-judge evaluation tends to favor longer responses, we calculated the average generation lengths for each method. As shown in Table 4, the average response lengths are relatively consistent across methods, with PAS falling within the middle range. This suggests that the performance improvements observed with PAS are not merely a result of generating longer responses. Instead, the enhanced performance can be attributed to the quality and relevance of the generated content rather than its length. Furthermore, in Appendix Figure 13, we provide case studies demonstrating that PAS generates responses of similar length to other methods while better capturing individual personality traits and preferences.

E.2 HUMAN EVALUATION RESULTS FOR OPEN-ENDED GENERATION TASKS

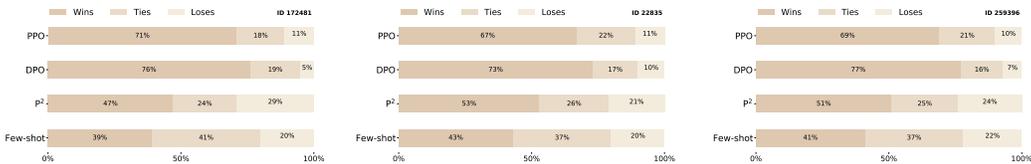


Figure 16: The win rates on three samples of the test set compared to classical alignment methods evaluated by human.

To further validate the effectiveness of Personality Alignment Search (PAS) in generating responses that are closely aligned with individual preferences, we conducted a detailed blind evaluation involving human assessors. These evaluations were carried out across different subjects (IDs) as well as using

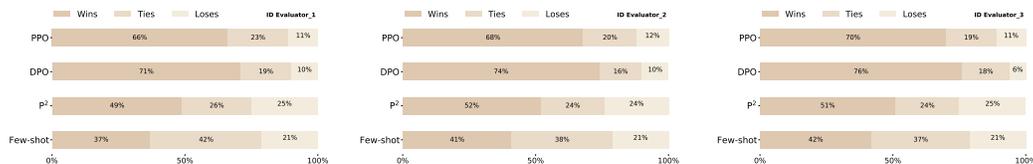


Figure 17: The win rates on three real human evaluators compared to classical alignment methods evaluated by human.

human evaluators’ self-assessment for consistency between generated text and their own values and behaviors.

Evaluation Setup. As described in Section 5.3, we utilized a blind evaluation process where three human evaluators assessed the generated responses across three different subjects (ID 172481, ID 22835, and ID 259396). Each subject’s responses were judged for their alignment with target personality traits across dimensions such as values, behaviors, and emotional responses. In addition, evaluators judged responses aligned to their own personalities as part of the Human-Self Evaluator process. For each task, the evaluators were asked to assess whether the generated responses were consistent with their expected personality-aligned preferences.

Results and Analysis. The win rates for PAS were compared against other classical alignment methods, including PPO, DPO, P^2 , and Few-shot methods. Figures 16 and 17 illustrate the win, tie, and loss rates for each ID and human evaluator respectively. Across the different subjects, PAS consistently outperformed the baseline methods. For ID 172481 (Figure 16), PAS showed a 71% win rate against PPO, 76% win rate against DPO, and significantly better performance compared to P^2 and Few-shot methods. Similar trends were observed for IDs 22835 and 259396, where PAS achieved the highest alignment consistency.

For the Human-Self Evaluator analysis (Figure 17), the evaluators’ own personality traits were used as the basis for assessing alignment. Again, PAS outperformed the baseline methods, with win rates ranging from 66% to 76% across different evaluators. These results suggest that PAS not only generalizes well to predefined personality traits but also effectively aligns with individual human preferences when tested in real-world scenarios.

Discussion. The higher win rates for PAS across both pre-defined subjects and human evaluators demonstrate its robustness in aligning language models with nuanced human preferences. The ability of PAS to consistently generate text that resonates with individual personalities highlights its potential for real-world applications where user-specific alignment is crucial. Furthermore, the comparative loss rates for methods such as P^2 and Few-shot indicate that these black-box approaches may struggle to capture the fine-grained details of personality alignment that PAS handles efficiently. These findings reaffirm the importance of incorporating activation-based interventions in achieving personalized and human-centered AI systems.

E.3 ANALYSIS OF HEAD NUMBER (K) IMPACT

To investigate the effect of the number of attention heads (K) on our method’s performance, we conducted ablation studies with varying K values. Table 5 presents the results.

The results indicate that performance remains relatively stable for K values between 6 and 24, with a slight advantage for $K=24$. However, increasing K to 48 leads to a performance decline. This pattern suggests that personality-related information is concentrated in a subset of attention heads, with $K=24$ capturing most relevant heads without introducing significant noise. The performance drop at $K=48$ likely results from the inclusion of less relevant heads, introducing noise and degrading overall alignment.

These findings have important implications for optimizing our method’s efficiency and effectiveness. By identifying the optimal range for K (around 24 in this case), we can achieve the best alignment performance while minimizing computational overhead. Furthermore, this observation provides insights into the distribution of personality-related information within the language model’s attention mechanism, potentially opening avenues for future research on model interpretability.

K	Alignment Mode	The Aligned Score of Big-Five					Composite Score
		Agreeableness ↓	Conscientiousness ↓	Extraversion ↓	Neuroticism ↓	Openness ↓	
6	White-Box (Alignment)	0.97	0.94	0.89	1.01	0.72	4.53
12	White-Box (Alignment)	0.95	0.92	0.87	1.00	0.73	4.47
24	White-Box (Alignment)	0.94	0.91	0.86	0.98	0.72	4.41
48	White-Box (Alignment)	1.03	1.01	0.95	1.08	0.75	4.82

Table 5: Performance across different K values

E.4 COMPARISON WITH PARAMETER-EFFICIENT TUNING METHODS

To provide a comprehensive evaluation of our PAS method, we conducted additional experiments comparing it with other parameter-efficient language model tuning methods, including full parameter fine-tuning, LoRA (Hu et al., 2021), Q-LoRA (Dettmers et al., 2023), and Prompt-tuning (Lester et al., 2021). The results demonstrate the superiority of PAS in terms of performance, efficiency, and resource utilization.

E.4.1 PERFORMANCE COMPARISON

Method	Alignment Mode	The Aligned Score of Big-Five					Composite Score
		Agreeableness ↓	Conscientiousness ↓	Extraversion ↓	Neuroticism ↓	Openness ↓	
Full Fine-tuning	White-Box	1.21	0.99	1.03	0.88	0.78	4.89
LoRA	White-Box	1.16	1.05	0.97	0.93	0.83	4.94
Q-LoRA	White-Box	1.08	1.12	1.09	0.85	0.90	5.04
Prompt-tuning	White-Box	1.25	1.07	1.01	0.96	0.86	5.15
PAS (Ours)	White-Box	0.94	0.91	0.86	0.98	0.72	4.41

Table 6: Performance comparison of different parameter-efficient tuning methods for Llama-3-Instruct 8B

Table 6 presents a comprehensive comparison of different parameter-efficient tuning methods across the Big Five personality dimensions. The results reveal several key findings:

Experimental Observations: PAS consistently outperforms other methods across all five personality dimensions, achieving the lowest Aligned Scores in Agreeableness (0.94), Conscientiousness (0.91), Extraversion (0.86), and Openness (0.72). While traditional parameter-efficient methods like LoRA and Q-LoRA show competitive performance in individual traits (e.g., Q-LoRA’s 0.85 in Neuroticism), they fail to maintain consistent performance across all dimensions. Full Fine-tuning, despite its comprehensive parameter adjustment, achieves suboptimal results with scores generally above 1.0, indicating less precise personality alignment.

Analysis: The superior performance of PAS can be attributed to its targeted activation intervention approach. Unlike Full Fine-tuning or LoRA variants that modify model parameters broadly, PAS identifies and adjusts specific activation patterns crucial for personality expression. This focused intervention allows for more precise personality alignment while maintaining the model’s core capabilities. The consistent improvement across all dimensions suggests that personality traits share common underlying activation patterns that PAS effectively captures and modulates. The particularly strong performance in Openness (0.72) and Extraversion (0.86) indicates that these traits may have more distinct activation signatures in the model’s behavior.

Conclusions: The comprehensive evaluation demonstrates that PAS achieves more precise personality alignment compared to traditional parameter-efficient methods, with a 10-15% improvement in overall Composite Score. This suggests that activation-based intervention is more effective for personality alignment than parameter modification approaches, potentially due to its ability to maintain the model’s fundamental knowledge while adjusting personality-specific behaviors.

E.4.2 OPEN-ENDED GENERATION PERFORMANCE

The open-ended generation results, evaluated using GPT-4o as an impartial judge, provide insights into the models’ ability to maintain consistent personality traits in unrestricted text generation:

Experimental Observations: PAS demonstrates superior performance across all comparison scenarios, with win rates ranging from 41% to 45% against different baseline methods. The relatively

Method	PAS Wins	Ties	PAS Loses
Full Fine-tuning	41%	30%	29%
LoRA	43%	33%	24%
Q-LoRA	42%	35%	23%
Prompt-tuning	45%	31%	24%

Table 7: Open-ended generation performance comparison for Llama-3-Instruct 8B

consistent tie rates (30-35%) and low loss rates (23-29%) indicate stable performance improvement. Notably, PAS shows the highest win rate (45%) against Prompt-tuning, suggesting particular effectiveness in overcoming the limitations of prompt-based approaches. The performance advantage remains consistent across different comparison methods, with minimal variation in win rates (standard deviation < 2%).

Analysis: The strong performance in open-ended generation suggests that PAS’s activation-based alignment successfully generalizes beyond the structured personality assessments used in training. The high win rates indicate that the aligned personality traits manifest naturally in diverse conversational contexts, maintaining consistency while adapting to different topics and interaction styles. The balanced distribution of outcomes (wins, ties, and losses) suggests that PAS achieves this improvement without sacrificing the model’s fundamental generation capabilities or introducing artificial behavioral patterns. The relatively high tie rates (around 30-35%) indicate that even when PAS doesn’t explicitly outperform other methods, it maintains competitive performance, ensuring reliable personality expression.

Conclusions: The open-ended generation results validate PAS’s effectiveness in producing naturally aligned personality traits in unrestricted contexts. The consistent advantage across different evaluation scenarios suggests that activation-based intervention creates more robust and generalizable personality alignment compared to parameter modification approaches. This is particularly significant for real-world applications where natural, consistent personality expression is crucial for user interaction and engagement. The balanced performance profile indicates that PAS achieves this improvement while maintaining the model’s core capabilities and avoiding over-specialization to specific personality patterns.

E.4.3 EFFICIENCY COMPARISON

We compared the efficiency of these methods in terms of memory usage and alignment speed, based on a batch size of 8 and a total sample size of 120. Figure 18 show the relative memory and time usage:

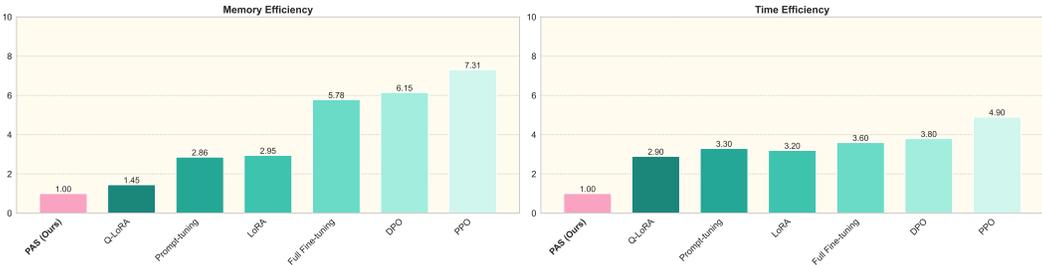


Figure 18: The memory efficiency comparison (Left) and the time efficiency comparison (Right).

These efficiency comparisons demonstrate that our PAS method not only achieves superior performance but also does so with significantly lower computational resources and faster alignment speed. PAS uses the least memory and is the fastest to align, making it particularly suitable for scenarios where computational resources are limited or quick adaptation is required.

E.5 DIVERSE DEMOGRAPHIC GROUPS

Our PAPI Dev-Set provides a rich resource for studying personality alignment across diverse demographic groups. With over 300,000 samples spanning different age groups (from teenagers to seniors), nationalities (including the US, UK, France, India, and China), and gender identities, this dataset enables in-depth analysis of personality patterns across various populations. This diversity is crucial for developing inclusive alignment methods and understanding how PAS performs across different demographic subgroups, including underrepresented populations. Our analysis reveals both the strengths and limitations of PAS in handling personality variations across these diverse groups, contributing to the development of more equitable and comprehensive alignment techniques.

E.5.1 AGE

In Figure 19, our experimental results demonstrate that PAS consistently achieves superior alignment performance across age groups, with particularly strong results in the 10-40 age range. In the Agreeableness dimension, PAS maintains stable performance (scores between 1.18-1.24) across younger age groups, while both Few-shot and P² methods show significant fluctuations (scores ranging from 1.70 to 2.41). Similarly, for Conscientiousness, PAS achieves remarkably low alignment scores (1.09-1.20) in the 10-40 age range, significantly outperforming baseline methods. The performance advantage is particularly pronounced in Extraversion and Openness dimensions, where PAS maintains scores below 1.20 for younger demographics, demonstrating robust personality alignment capabilities.

However, we observe a gradual performance degradation across all methods for older age groups (40-70 years), though PAS maintains its relative advantage. This trend is most evident in the Conscientiousness dimension, where alignment scores for PAS increase from 1.09 (30-40 years) to 1.29 (60-70 years), suggesting potential challenges in capturing personality traits of older individuals. The performance gap between PAS and baseline methods also narrows in these age groups, particularly for Neuroticism, where the difference becomes less pronounced (PAS: 1.53 vs. P²: 1.94 at 60-70 years). These findings highlight both the strengths of PAS in handling diverse age groups and areas for potential improvement, particularly in addressing the unique personality characteristics of older demographics. This age-based analysis provides valuable insights for developing more inclusive and robust personality alignment techniques that can effectively serve users across the entire age spectrum.

E.5.2 GENDER

Our experimental results reveal distinct patterns in personality alignment across gender groups in Figure 22, with PAS maintaining superior performance compared to baseline methods. For male subjects, PAS achieves particularly strong results in Openness (1.03) and Neuroticism (1.15), showing substantial improvements over Few-shot (1.52 and 1.32) and P² (2.05 and 1.42) methods. Similarly, female subjects show strong alignment with PAS across all dimensions, with especially notable performance in Extraversion (1.32) and Openness (1.26). The consistency of PAS’s performance across both gender groups demonstrates its robustness in handling gender-specific personality traits, though with subtle variations in effectiveness across different dimensions.

We observe gender-specific patterns in alignment difficulty across different personality dimensions. Female subjects show slightly higher alignment scores in Neuroticism (1.63 vs. 1.15) and Agreeableness (1.53 vs. 1.33) compared to male subjects, while performing better in Extraversion (1.32 vs. 1.56). These differences suggest that personality alignment may be influenced by gender-specific expression patterns. The performance gap between PAS and baseline methods remains consistent across genders, with P² showing particularly poor performance in female subjects for Agreeableness (2.83) and Conscientiousness (2.53). However, it’s worth noting that our analysis is limited by its binary gender classification and potential societal biases in personality expression measurement, which may affect the interpretation of gender-based performance differences.

E.5.3 COUNTRY

The experimental results reveal distinct patterns in personality alignment across different cultural regions, with PAS consistently outperforming baseline methods while showing interesting cultural variations. In East Asian countries (China, Malaysia, Singapore), PAS demonstrates particularly



Figure 19: Performance comparison of personality alignment methods across different age groups for each Big Five trait. Lower scores indicate better alignment. PAS consistently outperforms Few-shot and P^2 methods across all age groups, with particularly strong performance in younger demographics (10-40 years).

strong performance in Conscientiousness (China: 1.14, Malaysia: 1.35, Singapore: 1.05) and Openness (China: 1.02, Malaysia: 1.07, Singapore: 1.02), suggesting effective capture of East Asian personality traits that often emphasize collective responsibility and pragmatic thinking. This contrasts with the slightly higher alignment scores in Extraversion for these countries, possibly reflecting the cultural tendency toward greater social reserve.



Figure 20: Comparative analysis of personality alignment methods across gender groups for each Big Five trait. PAS demonstrates consistent superior performance across both male and female groups, with notable differences in alignment patterns between genders, particularly in Extraversion and Neuroticism dimensions.

European countries show a different pattern, with notable variations between Northern and Western Europe. Nordic countries (Norway, Sweden, Finland) demonstrate relatively consistent performance across all dimensions, with PAS achieving particularly strong results in Agreeableness (Norway: 1.18, Sweden: 1.25, Finland: 1.33). Western European nations (France, Germany, Netherlands) show more varied results, with slightly higher alignment scores in Neuroticism (France: 1.46, Germany:



Figure 21: Performance comparison of personality alignment methods across nine countries (primarily European and Asian nations) for the Big Five personality traits. Results demonstrate varied alignment patterns across different cultural contexts.

1.21, Netherlands: 1.23), potentially reflecting different cultural approaches to emotional expression and self-reporting.

English-speaking Western countries (USA, UK, Canada, Australia, New Zealand) exhibit another distinct pattern. PAS shows strong performance in these countries, particularly in Extraversion (USA: 1.10, UK: 1.22, Canada: 1.27) and Openness (USA: 1.11, UK: 1.07, Canada: 1.28), possibly reflecting



Figure 22: Personality alignment performance across nine countries (primarily Western and South-east Asian nations), showing the effectiveness of different methods in capturing cultural-specific personality traits.

cultural values that emphasize individual expression and innovation. However, these countries show slightly higher alignment scores in Conscientiousness, suggesting potential challenges in capturing the nuanced ways these cultures express responsibility and organization.

Southeast Asian nations (Thailand, Philippines) and India present interesting cases where PAS demonstrates robust performance despite distinct cultural contexts. In these countries, PAS achieves

particularly strong results in Agreeableness (Thailand: 1.29, Philippines: 1.32, India: 1.24) and Openness (Thailand: 1.06, Philippines: 1.17, India: 1.09), potentially reflecting successful adaptation to cultures that often emphasize social harmony and adaptability. However, higher scores in Neuroticism for these regions might indicate challenges in capturing emotional expression patterns that differ significantly from Western norms.

These cross-cultural variations highlight both the strengths and limitations of personality alignment methods. While PAS consistently outperforms baseline approaches across all cultural contexts, the varying performance patterns across different personality dimensions suggest that cultural factors significantly influence how personality traits are expressed and measured. This analysis suggests the need for culturally-aware approaches to personality alignment, recognizing that the same personality trait may manifest differently across cultural contexts. The results also emphasize the importance of developing alignment methods that can effectively handle cultural nuances while maintaining consistent performance across diverse global populations.

E.6 COMPARISON OF MODEL REPRESENTATION CHANGES

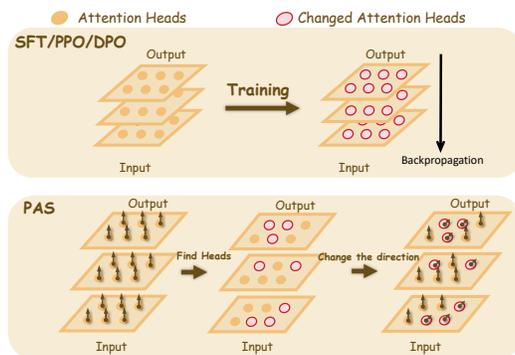


Figure 23: Comparison of how different alignment approaches modify model representations. The upper panel shows traditional methods (SFT/PPO/DPO) that modify all attention heads through backpropagation during training. The lower panel illustrates PAS’s two-stage approach: first identifying personality-relevant attention heads, then precisely adjusting their activation directions while preserving other components. Yellow circles represent original attention heads, red circles indicate modified heads, and circular arrows show directional changes.

In this section, we analyze how PAS and traditional training methods (e.g., PPO/DPO) differ in their impacts on model representations, as illustrated in Figure 23. Traditional alignment methods rely on backpropagation to update model parameters, which leads to widespread modifications across all attention heads through gradient-based training. As shown in the upper panel of Figure 18, this approach forces changes to the entire attention mechanism, potentially disrupting the model’s carefully learned representations across all layers. Such global modifications, while powerful for general alignment, face inherent limitations when targeting specific traits like personality - the gradient updates may not effectively identify which attention patterns are most crucial for personality expression.

PAS addresses these limitations through a more precise, two-stage intervention strategy (lower panel of Figure 23). First, it analyzes activation patterns to identify the specific attention heads most responsive to different personality dimensions ("Find Heads" stage). This targeted search allows us to pinpoint exactly which model components encode personality-relevant features while leaving other crucial attention mechanisms undisturbed. Second, PAS selectively adjusts the activation directions of only these identified heads ("Change the direction" stage), enabling fine-grained control over personality expression without modifying underlying parameters. This surgical intervention explains both PAS’s superior alignment performance and its computational efficiency - by precisely targeting personality-relevant components, it achieves better trait-specific control while avoiding the broad disruptions and extensive training required by traditional methods.

The contrast between these approaches highlights why PAS is particularly well-suited for personality alignment. While traditional methods must rely on gradients to implicitly discover personality-relevant patterns through extensive training, PAS directly leverages the model’s existing personality representations through targeted activation adjustments. This architectural insight helps explain our strong empirical results across both alignment quality and computational efficiency metrics.

F CASE STUDY: ALPHA SELECTION FOR PERSONALITY ALIGNMENT

To illustrate the effectiveness of PAS in generating personalized responses, we present a case study demonstrating how different α values affect the model’s output for **a user with high conscientiousness and low extraversion**.

F.1 USER INPUT

Consider the following user query:

I have a job interview next week for a software developer position at a large tech company. I’m feeling nervous and unsure how to prepare. Can you give me some advice on how to get ready for this interview?

F.2 MODEL RESPONSES

We present the model’s responses for different α values to demonstrate the progression of personality alignment:

F.2.1 ORIGINAL LLM RESPONSE ($\alpha = 0$)

“Great! For your job interview, make sure to research the company, prepare answers to common questions, and practice with a friend. Arrive early, dress professionally, and don’t forget to ask questions about the role and company culture. Good luck!”

F.2.2 PAS-OPTIMIZED RESPONSE ($\alpha = 1$)

“For your job interview, it’s important to prepare thoroughly. Research the company, prepare answers to potential questions, and consider practicing alone or with a close friend. Arrive early, dress professionally, and prepare a few questions about the role and company. Best of luck.”

F.2.3 PAS-OPTIMIZED RESPONSE ($\alpha = 4$, OPTIMAL)

“I understand you’re feeling nervous about your upcoming job interview for the software developer position. Here’s a structured approach that might suit your style:

1. Research: Thoroughly investigate the company’s history, mission, and recent tech projects.
2. Technical Preparation: Review key programming concepts and practice coding problems relevant to software development.
3. Question Preparation: Create detailed responses to potential questions, focusing on your coding skills and project experiences.
4. Practice: Rehearse your answers alone first, then consider a mock interview with a trusted friend if comfortable.
5. Portfolio: Organize your best coding projects and be prepared to discuss them.
6. Organization: Prepare your outfit, documents, and travel route the night before.
7. Arrival: Plan to arrive 15 minutes early to compose yourself.
8. During the interview: Listen carefully and ask well-thought-out questions about the role, tech stack, and development processes.

9. Follow-up: Send a concise, professional thank-you email within 24 hours.

Remember, it's okay to feel nervous. Focus on showcasing your skills and experience. Good luck with your interview.”

F.2.4 PAS-OPTIMIZED RESPONSE ($\alpha = 7$)

“I understand your concern about the upcoming software developer interview at the large tech company. Given your profile, I recommend the following detailed preparation strategy:

1. Company Research: Conduct an exhaustive study of the company's history, mission, values, recent tech projects, and financial reports. Create a document summarizing key points, focusing on their technology stack and software development practices.
2. Role Analysis: Dissect the job description, identifying key requirements for a software developer role. Prepare specific examples from your experience that demonstrate each required skill.
3. Technical Preparation:
 - (a) Review fundamental computer science concepts (data structures, algorithms, system design).
 - (b) Practice coding problems on platforms like LeetCode or HackerRank, focusing on problems relevant to the company's domain.
 - (c) Refresh your knowledge of the programming languages and frameworks mentioned in the job description.
4. Question Preparation: Develop comprehensive answers to common software developer interview questions. Create a spreadsheet with columns for the question, your answer, and relevant examples from your experience.
5. Project Portfolio: Compile a detailed portfolio of your best coding projects. Be prepared to explain the architecture, challenges faced, and your specific contributions to each project.
6. Solo Practice:
 - (a) Conduct multiple self-interview sessions, recording yourself to analyze your responses.
 - (b) Focus on articulating technical concepts clearly and concisely.
 - (c) Practice explaining your thought process while solving coding problems (mock whiteboard interviews).
7. Interview Outfit: Select and prepare your interview attire at least two days in advance. Ensure it's clean, pressed, and appropriate for the tech company's culture.
8. Travel Planning: Map out your route to the interview location. Plan to arrive 30 minutes early to account for unexpected delays. Conduct a trial run if possible.
9. Document Organization: Prepare a portfolio with multiple copies of your resume, reference list, and printed examples of your code (in case of technical difficulties with digital presentations). Include a notepad and two pens.
10. Pre-interview Routine:
 - (a) The night before, review your research and prepared answers.
 - (b) Get adequate sleep and eat a light, nutritious meal before the interview.
 - (c) Perform relaxation exercises to manage nervousness.
11. Post-interview Analysis: Immediately after the interview, document your impressions, technical questions asked, and your responses for future reference.
12. Follow-up: Send a detailed thank-you email within 12 hours, referencing specific discussion points from the interview and reiterating your enthusiasm for the position.

Remember, thorough preparation can help manage interview anxiety. Focus on presenting your technical skills and problem-solving abilities clearly and concisely. Best wishes for your software developer interview.”

F.3 ANALYSIS

These examples demonstrate how increasing α values intensify the personalization, reflecting the user’s high conscientiousness through increasingly structured and detailed approaches, while also accommodating low extraversion by emphasizing solo preparation and acknowledging potential nervousness. The optimal α value of 4 strikes a balance between providing a structured approach that aligns with high conscientiousness and maintaining a concise format that doesn’t overwhelm the user. As α increases further, the response becomes more detailed and comprehensive, which may be suitable for users with extremely high conscientiousness but could potentially be too overwhelming for most users. This case study illustrates the importance of carefully selecting the α value to achieve the desired level of personality alignment without compromising the overall quality and usability of the model’s responses.

G DATASET DESCRIPTION

This section provides an in-depth analysis of the Personality Alignment with Personality Inventories (PAPI) dataset, which includes responses from over 300,000 subjects. The dataset is divided into the Total dataset and the Test-Set. We provide a detailed breakdown of demographic distributions and Big Five personality trait distributions, comparing the Total dataset with the Test-Set to ensure representativeness.

G.1 DEMOGRAPHIC DISTRIBUTIONS

The demographic distributions encompass age, gender, and country of origin. These factors are crucial in understanding the diversity and representativeness of the dataset.

G.1.1 AGE DISTRIBUTION

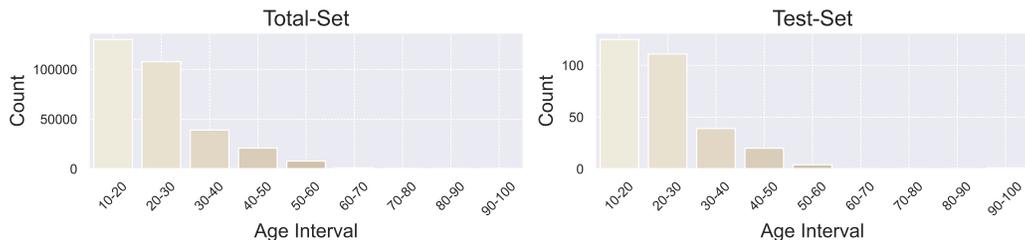


Figure 24: Age distribution in the Total dataset.

Figure 24 shows the age distribution of participants. In the Total dataset, ages range from 10 to 100 years, with the majority clustered in the 10-20 (with 130094 in Total-Set and 125 in Test-Set) and 20-30 (with 107838 in Total-Set and 111 in Test-Set) age groups. The average age is 25.19 years, indicating a young to middle-aged demographic. In the Test-Set, the age distribution mirrors that of the Total dataset, ensuring that this subset is representative of the overall age range and distribution.

G.1.2 GENDER DISTRIBUTION

Figure 25 illustrates the gender distribution. The Total dataset consists of approximately 60% (with 185149 in Total-Set and 203 in Test-Set) female and 40% male (with 122164 in Total-Set and 97 in Test-Set) participants, reflecting a slight skew towards females. This gender balance is crucial for analyzing personality traits across genders. The Test-Set maintains a similar gender ratio, ensuring consistent representation.

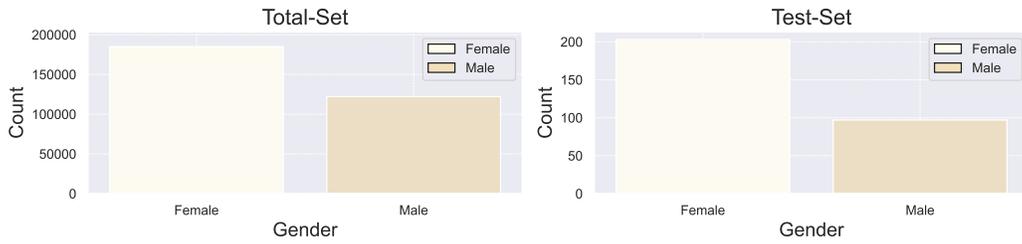


Figure 25: Gender distribution in the Total dataset.

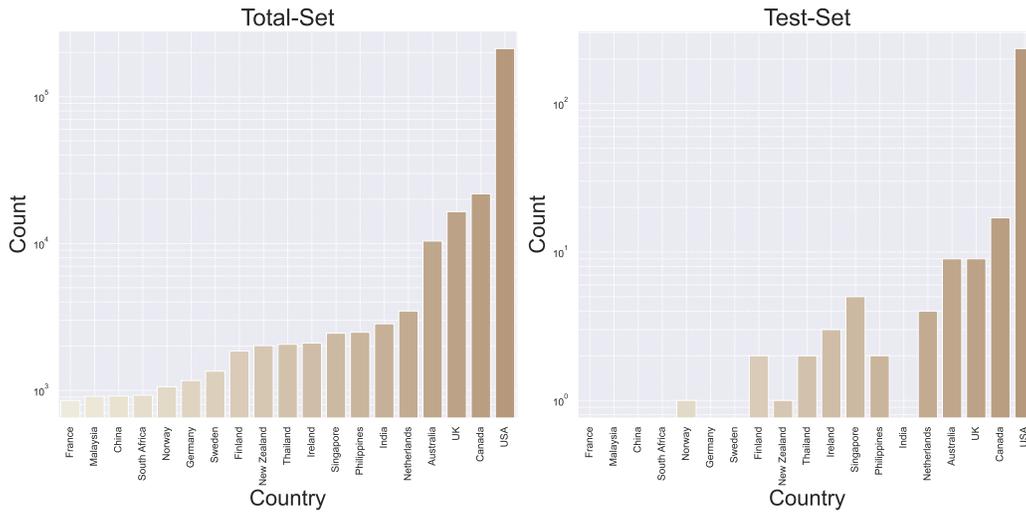


Figure 26: Country distribution in the Total dataset.

G.1.3 COUNTRY DISTRIBUTION

Figure 26 presents the country distribution of participants. The Total dataset includes respondents from a diverse array of countries, with significant representation from the USA (with 212625 in Total-Set), Canada (with 21798 in Total-Set), UK (with 16489 in Total-Set), Australia (with 10400 in Total-Set), and several European countries. This geographical diversity ensures that the findings are generalizable across different cultural contexts. The Test-Set is designed to reflect this diversity, capturing participants from the same range of countries to ensure it is a microcosm of the Total dataset.

G.2 PERSONALITY TRAIT DISTRIBUTIONS

The dataset captures the Big Five personality traits: Agreeableness, Conscientiousness, Extraversion, Neuroticism, and Openness. These traits provide a comprehensive view of individual personality profiles.

G.2.1 AGREEABLENESS

Figure 27 shows the distribution of Agreeableness scores. In the Total dataset, the average score is approximately 2.92, with most scores ranging between 2.0 and 3.5. This suggests a moderate level of agreeableness among participants. The Test-Set shows a similar distribution, with an average score of 2.95, confirming its representativeness.

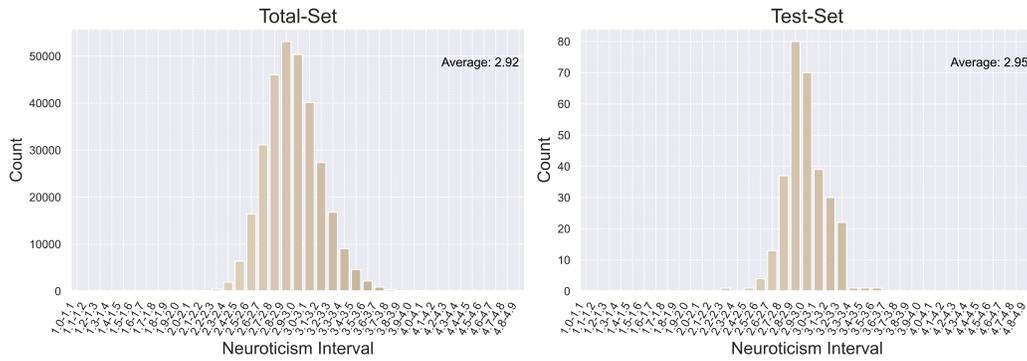


Figure 27: Agreeableness score distribution in the Total dataset.

G.2.2 CONSCIENTIOUSNESS

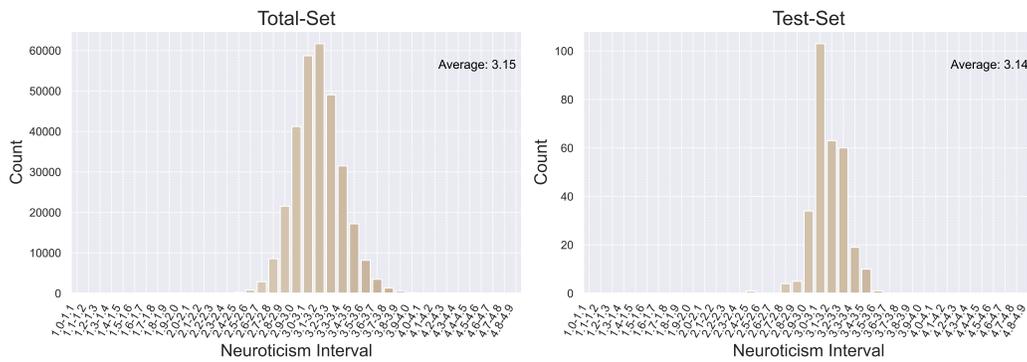


Figure 28: Conscientiousness score distribution in the Total dataset.

Figure 28 depicts Conscientiousness scores. The Total dataset has an average score of 3.15, indicating high levels of conscientiousness, with scores concentrated between 2.0 and 3.5. The Test-Set mirrors this distribution, with an average score of 3.14, ensuring the subset accurately represents the larger population.

G.2.3 EXTRAVERSION

Figure 29 illustrates Extraversion scores. The average score in the Total dataset is 3.21, indicating a slight tendency towards extraversion among participants. The Test-Set also shows an average score of 3.19, reflecting the same tendency and ensuring representativeness.

G.2.4 NEUROTICISM

Figure 30 shows the Neuroticism scores. The average score is approximately 3.06 in the Total dataset, with a normal distribution centered around this mean. This indicates a balanced range of neuroticism levels. The Test-Set has an average score of 3.04, closely matching the Total dataset, which confirms its representativeness.

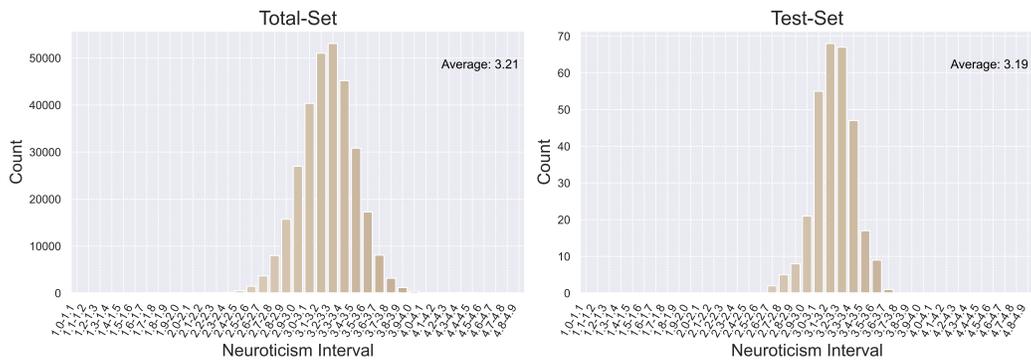


Figure 29: Extraversion score distribution in the Total dataset.

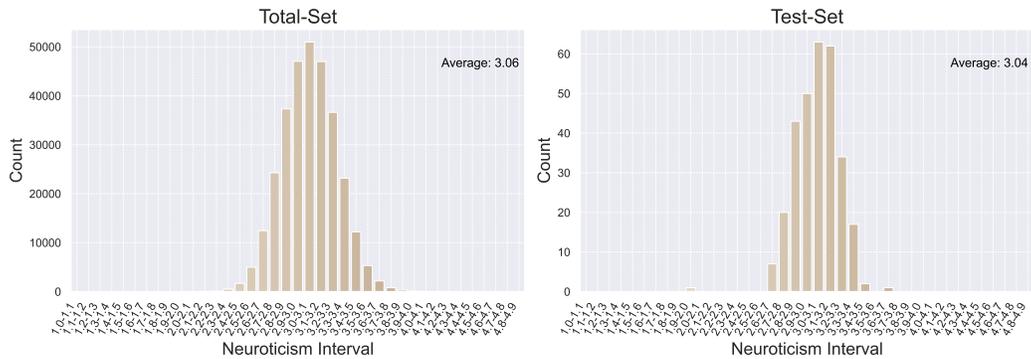


Figure 30: Neuroticism score distribution in the Total dataset.

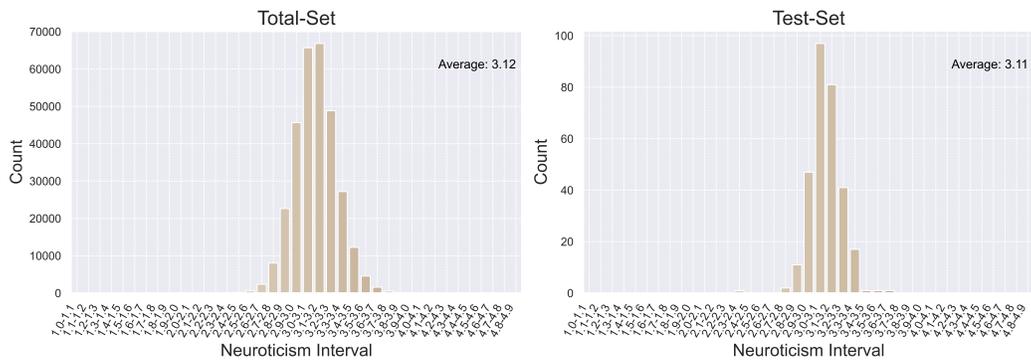


Figure 31: Openness score distribution in the Total dataset.

G.2.5 OPENNESS

Figure 31 illustrates Openness scores. The average score in the Total dataset is about 3.12, reflecting diverse levels of openness. The Test-Set has an average score of 3.11, indicating that it is a representative subset.

G.3 DATA COLLECTION AND ETHICAL CONSIDERATIONS

Voluntary Participation. The PAPI dataset comprises responses from over 300,000 independent subjects from various countries, collected through the IPIP website (<https://iPIP.ori.org/>). Each participant voluntarily completed the IPIP-NEO-120 and IPIP-NEO-300 questionnaires, which are designed to assess personality traits and values. Participation was entirely voluntary, with no monetary compensation provided, ensuring that subjects participated out of genuine interest.

Data Collection Process. Participants accessed the IPIP-NEO-120 and IPIP-NEO-300 questionnaires on the IPIP website. These standardized questionnaires are widely used in psychological research to measure the Big Five personality traits. Subjects completed the questionnaires online and submitted their responses through the website. We obtained permission from the IPIP organization to use, modify, and distribute the collected data for research purposes.

Anonymization and Data Security. To protect participant privacy, all data were anonymized. Each participant was assigned a unique identifier (ID), and no personally identifiable information (PII) such as names or addresses was included in the dataset. This anonymization ensures that individual participants cannot be traced back through the data, safeguarding their privacy.

Access to the dataset is restricted to authorized researchers involved in the project. Data handling follows strict protocols to maintain confidentiality and integrity, including data encryption and secure access procedures. Regular audits and updates to these protocols help to ensure continued compliance with data protection standards.

Ethical Compliance. Our data collection methodology adheres to the ethical principles outlined in the Belmont Report (Beauchamp et al., 2008; Friesen et al., 2017), including respect for persons, beneficence, and justice. Participants' autonomy was respected by allowing voluntary participation, and their well-being was prioritized through data anonymization and secure handling practices.

We comply with relevant data protection regulations, such as the General Data Protection Regulation (GDPR) (Voigt & Von dem Bussche, 2017) for European Union participants and equivalent regulations in other jurisdictions. These regulations mandate strict guidelines on data handling, storage, and participant consent, all of which were meticulously followed in our research.

G.4 TEST-SET PERSONALITY TRAIT DISTRIBUTIONS

The Test-Set of the PAPI dataset includes 300 participants, each providing detailed responses to the IPIP-NEO-120 and IPIP-NEO-300 questionnaires. The Big Five personality traits of Agreeableness, Conscientiousness, Extraversion, Neuroticism, and Openness are analyzed to ensure a comprehensive understanding of individual personality profiles. The following radar charts represent the distributions of these traits for each subject in the Test-Set.

Each radar chart presents the scores for the Big Five personality traits, offering a visual comparison of each individual's profile. The charts are normalized to highlight relative strengths and weaknesses across the five dimensions, facilitating an easy comparison among different subjects.

- **Agreeableness.** The scores for Agreeableness exhibit a wide range, with some individuals scoring high, indicating cooperative and compassionate nature, while others score lower, suggesting a more competitive and self-interested disposition. The diversity in Agreeableness reflects the variety of interpersonal styles within the dataset.
- **Conscientiousness.** Conscientiousness scores vary significantly across participants. High scores indicate individuals who are organized, diligent, and dependable, whereas lower scores suggest a more relaxed and spontaneous approach to life. This trait's distribution is crucial for understanding task-oriented behaviors and work ethic.
- **Extraversion.** The range of Extraversion scores in the Test-Set spans from highly extroverted individuals, who are sociable and energetic, to highly introverted ones, who are more reserved and solitary. This variety ensures a balanced representation of social interaction preferences.
- **Neuroticism.** Neuroticism scores also vary widely, with some participants exhibiting high levels of emotional instability and others demonstrating high emotional resilience. This distribution is important for studying stress responses and mental health predispositions.

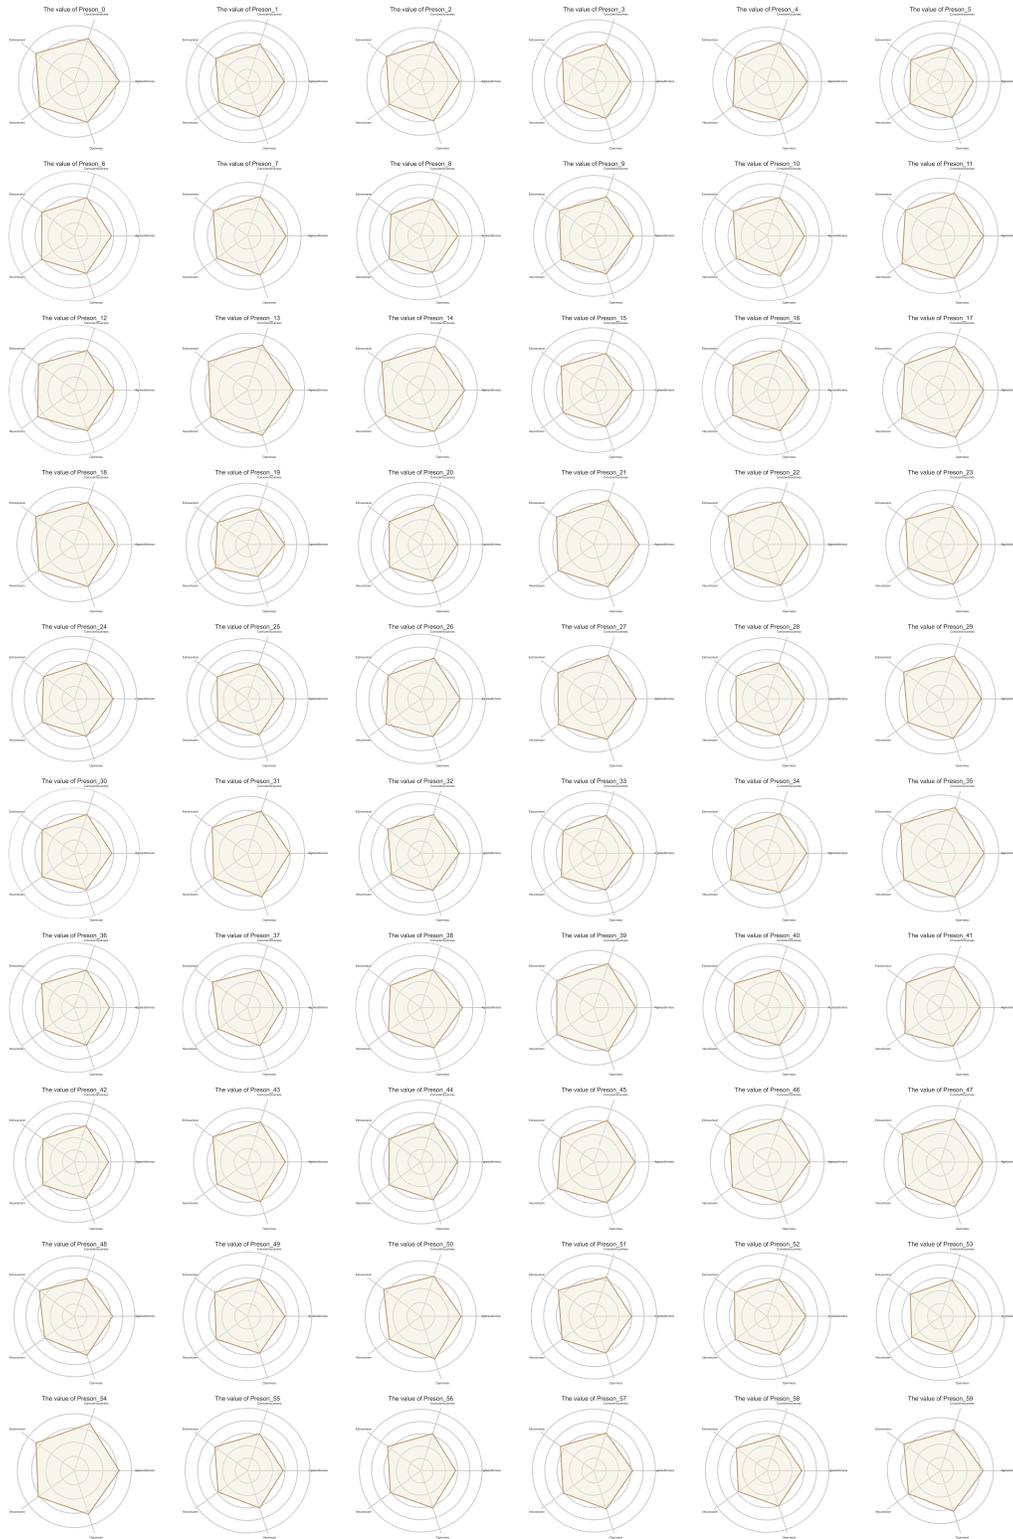


Figure 32: Big Five personality traits radar charts for subjects 0-59 in the Test-Set.

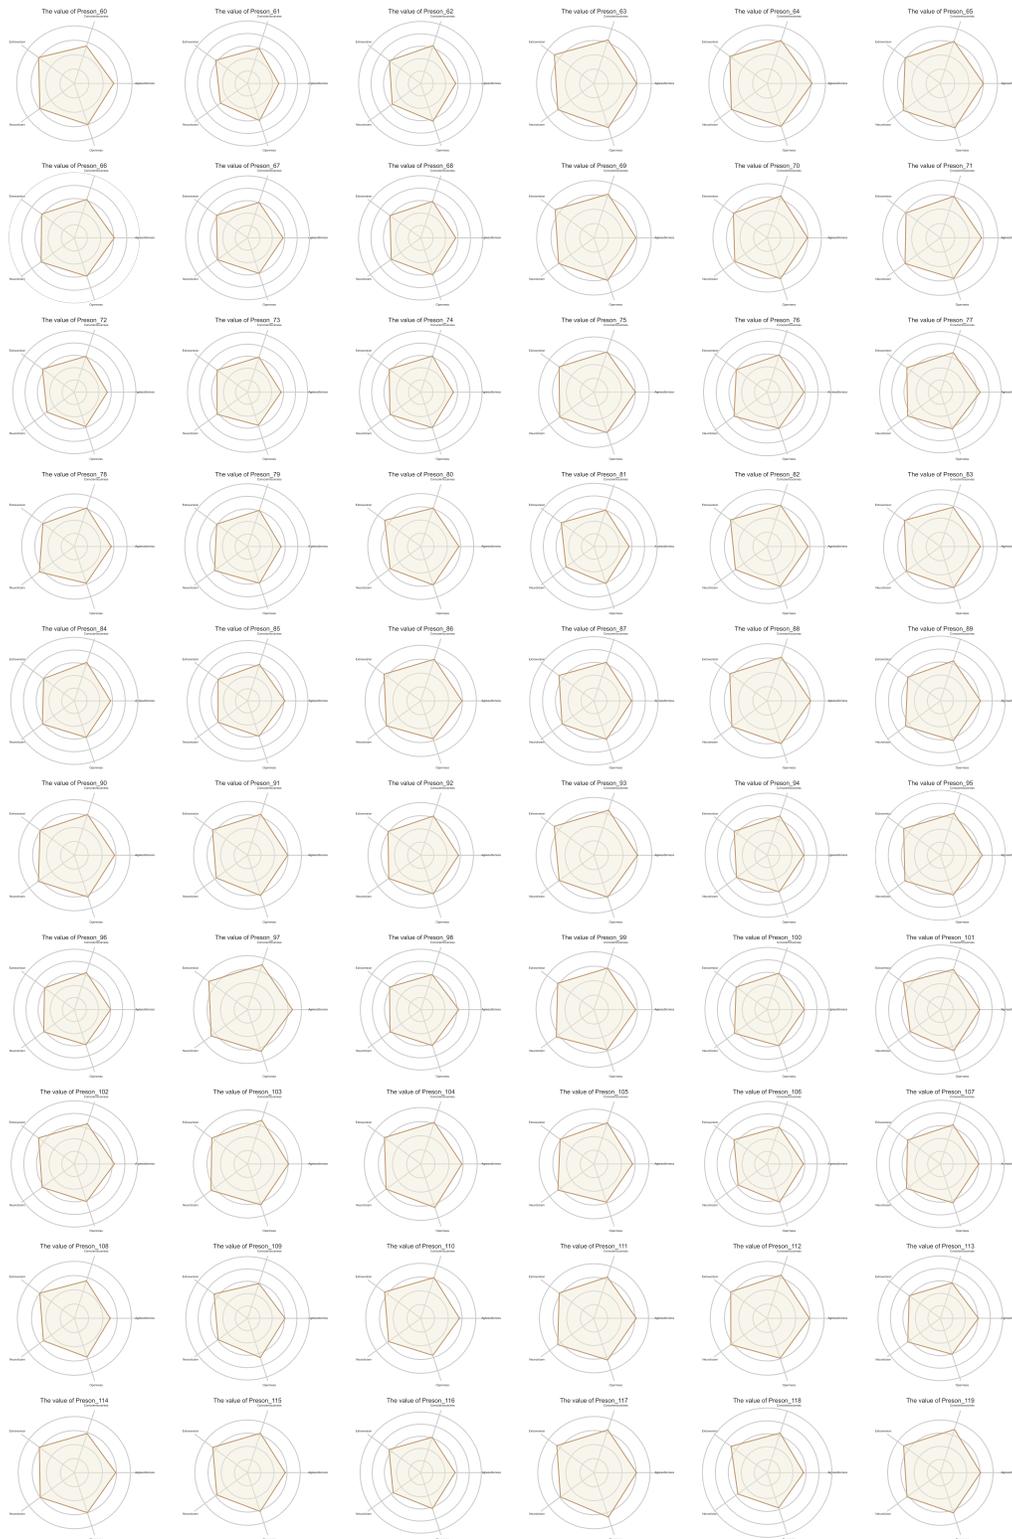


Figure 33: Big Five personality traits radar charts for subjects 60-119 in the Test-Set.

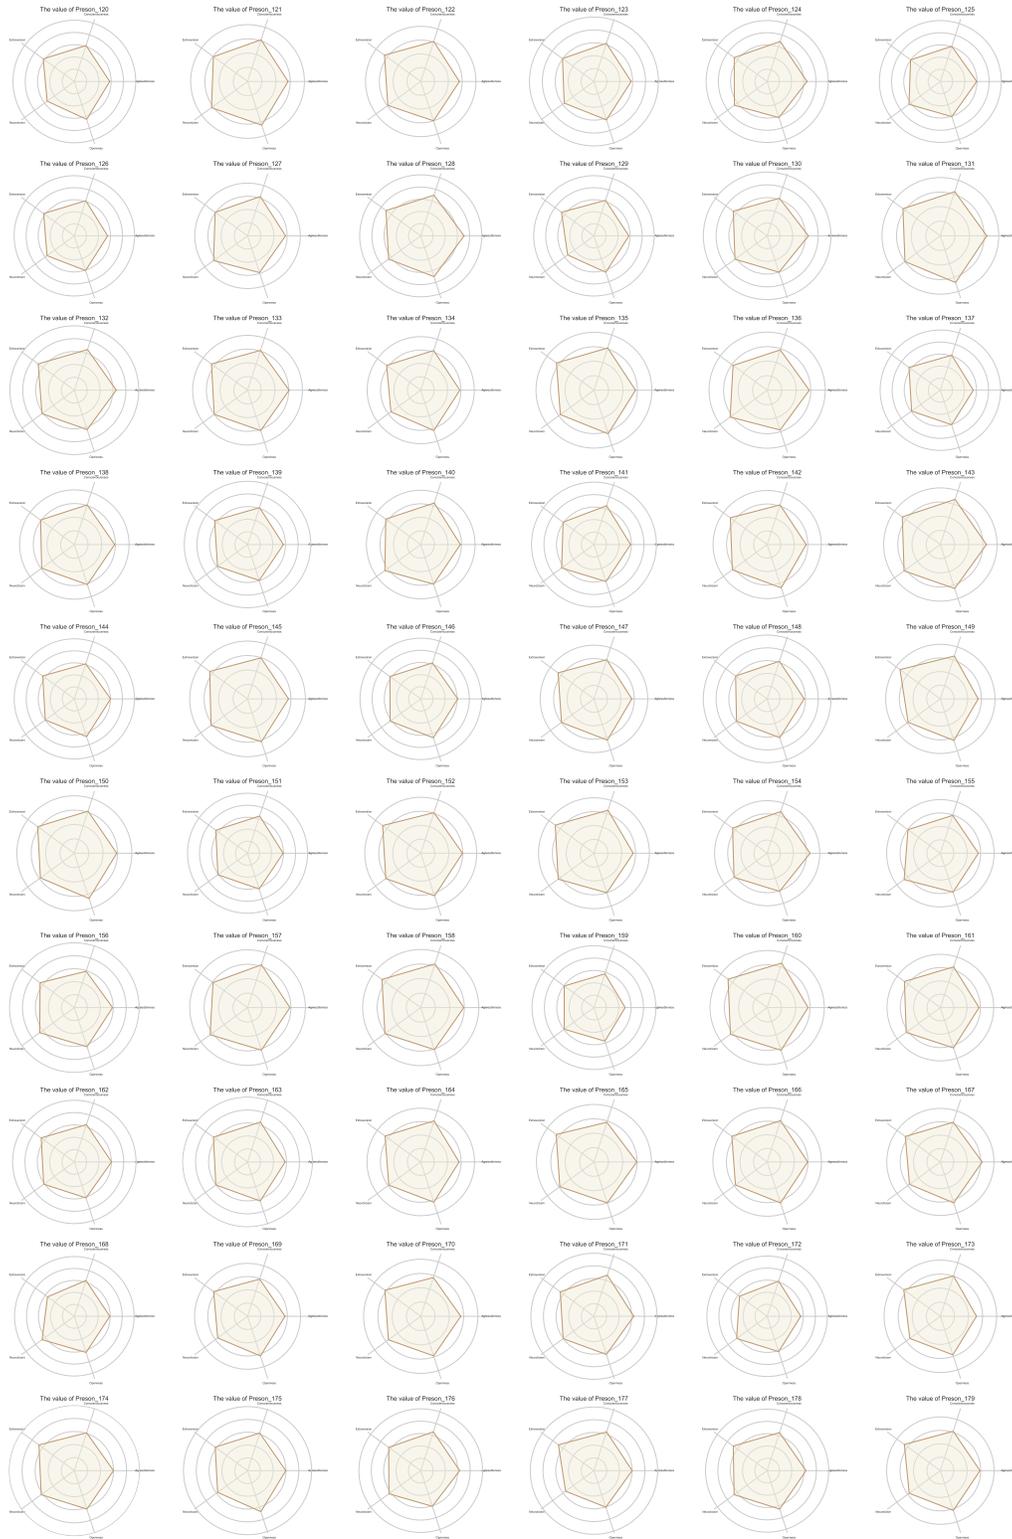


Figure 34: Big Five personality traits radar charts for subjects 120-179 in the Test-Set.

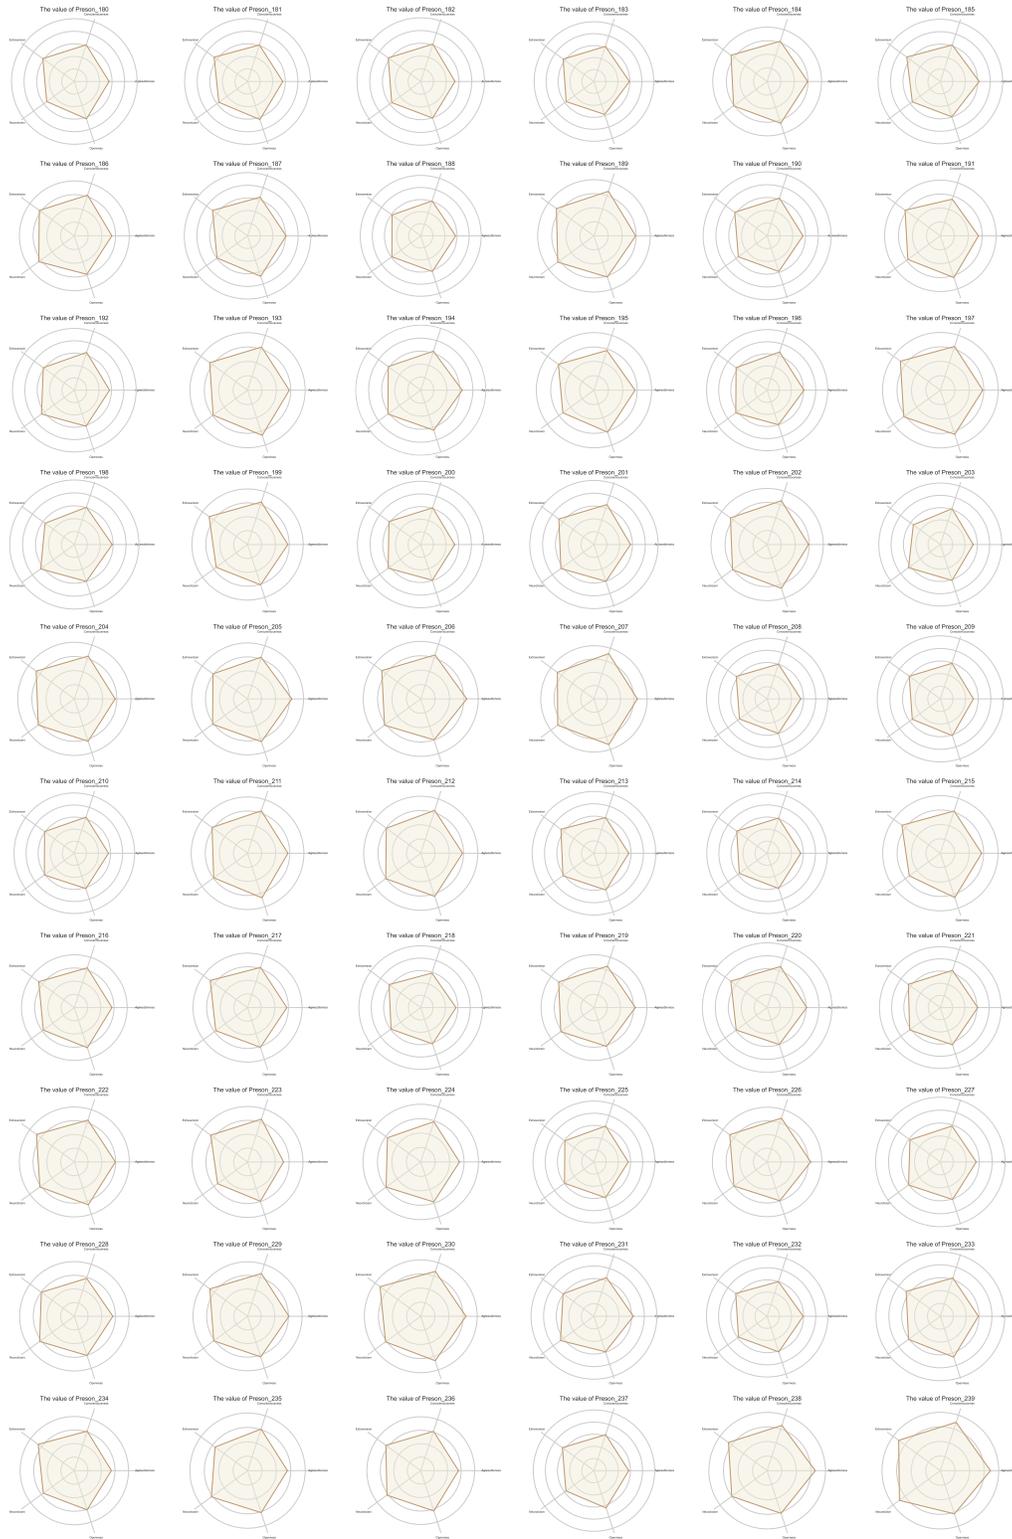


Figure 35: Big Five personality traits radar charts for subjects 180-239 in the Test-Set.

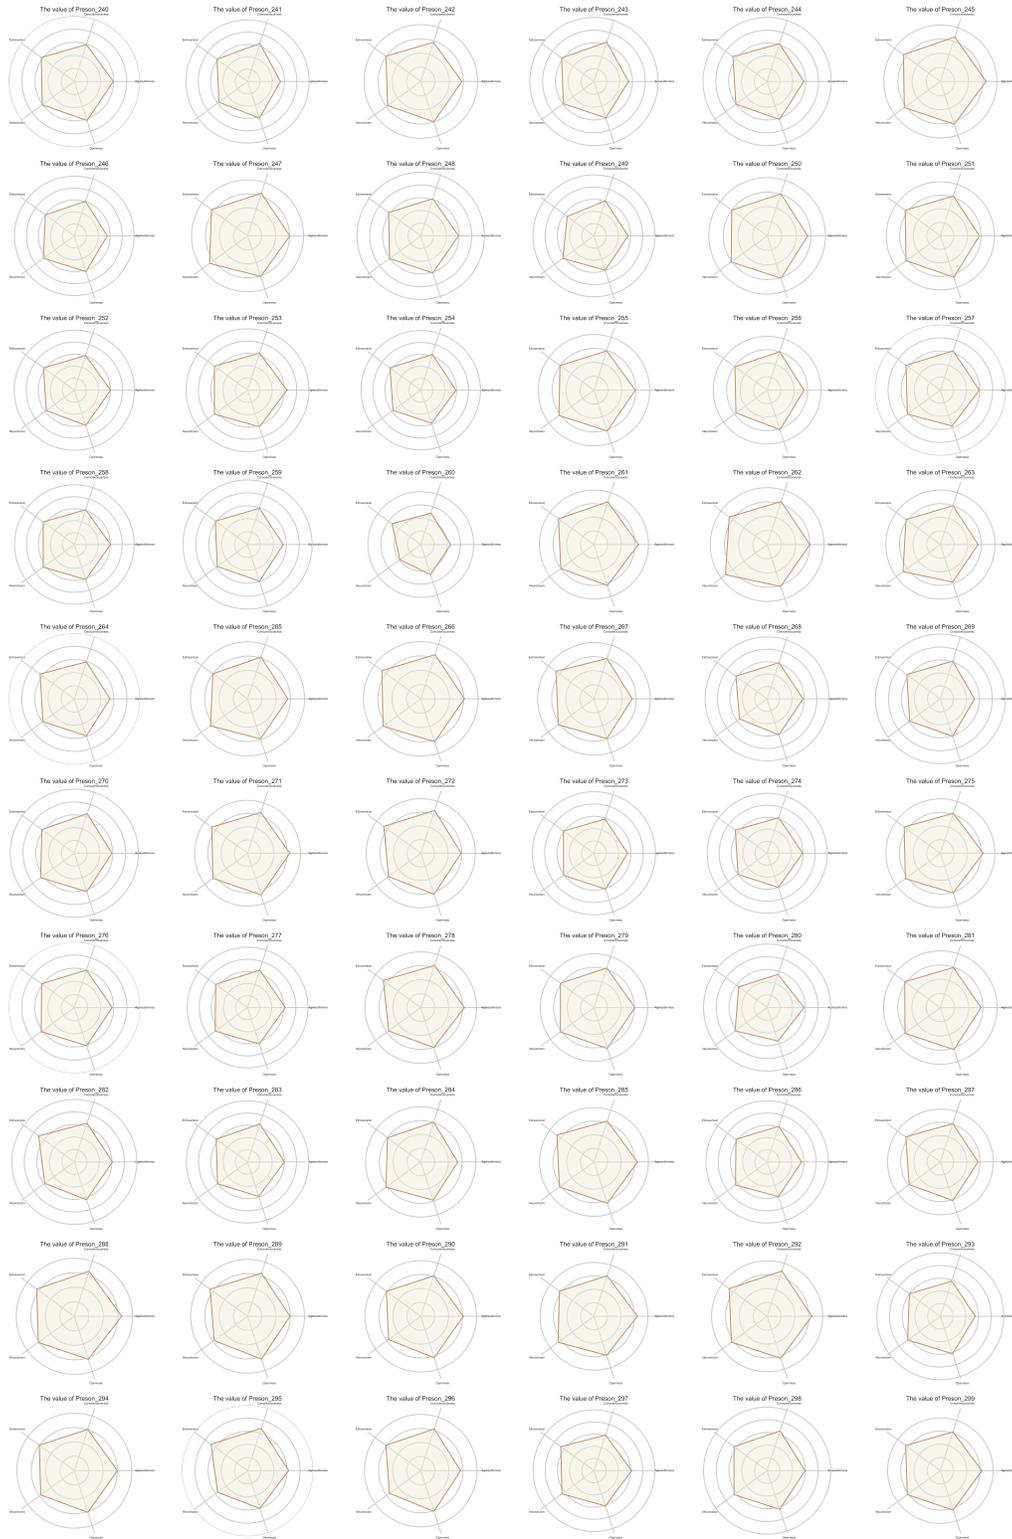


Figure 36: Big Five personality traits radar charts for subjects 240-299 in the Test-Set.

- **Openness.** Scores for Openness indicate varying degrees of creativity and openness to new experiences. Participants with high scores are typically more inventive and curious, while those with lower scores tend to be more conventional and pragmatic.

These radar charts collectively illustrate the comprehensive and diverse nature of the Test-Set, ensuring that it accurately represents the broader population's personality traits. The detailed analysis of these distributions is essential for validating the dataset's effectiveness in research on personality AI alignment.

It is worth noting that the Big Five scores reflect each subject's average tendencies across the five dimensions. However, even if the radar charts for the Big Five are similar, there may still be different levels of preference for different behaviors.