

AesBiasBench: Evaluating Bias and Alignment in Multimodal Language Models for Personalized Image Aesthetic Assessment

Kun Li¹, Lai-Man Po¹, Hongzheng Yang², Xuyuan Xu³, Kangcheng Liu⁴, Yuzhi Zhao^{1*}

¹City University of Hong Kong ²The Chinese University of Hong Kong

³Magilight (HK) Limited ⁴Hunan University

kunli25-c@my.cityu.edu.hk; yzzhao2-c@my.cityu.edu.hk

Abstract

Multimodal Large Language Models (MLLMs) are increasingly applied in Personalized Image Aesthetic Assessment (PIAA) as a scalable alternative to expert evaluations. However, their predictions may reflect subtle biases influenced by demographic factors such as gender, age, and education. In this work, we propose **AesBiasBench**, a benchmark designed to evaluate MLLMs along two complementary dimensions: (1) **stereotype bias**, quantified by measuring variations in aesthetic evaluations across demographic groups; and (2) **alignment** between model outputs and genuine human aesthetic preferences. Our benchmark covers three subtasks (Aesthetic Perception, Assessment, Empathy) and introduces structured metrics (IFD, NRD, AAS) to assess both bias and alignment. We evaluate 19 MLLMs, including proprietary models (e.g., GPT-4o, Claude-3.5-Sonnet) and open-source models (e.g., InternVL-2.5, Qwen2.5-VL). Results indicate that smaller models exhibit stronger stereotype biases, whereas larger models align more closely with human preferences. Incorporating identity information often exacerbates bias, particularly in emotional judgments. These findings underscore the importance of identity-aware evaluation frameworks in subjective vision-language tasks.

1 Introduction

Multimodal Large Language Models (MLLMs) have demonstrated impressive capabilities in vision-language tasks such as image recognition (Alayrac et al., 2022; Zhu et al., 2023), visual reasoning (Achiam et al., 2023; Wu et al., 2025), and visual question answering (Wu et al., 2023; Xu et al., 2025). Recently, these models have also been applied to Personalized Image Aesthetic Assessment (PIAA), which estimates the photographic or artistic quality of images based on individual




Figure 1: Examples illustrate bias in the image aesthetic empathy task. (a) and (b) show stereotypical bias in model outputs that arise from inherited cognitive priors. (c) presents human preferences for the image, which serve as a reference for evaluating the alignment of model predictions with human judgments.

preferences (Yang et al., 2022). PIAA applications include image retrieval, photo ranking, and creative recommendation (Ren et al., 2017).

Despite their promise, MLLMs may exhibit aesthetic bias, systematic differences in output driven by demographic attributes such as gender, age, geographic region, or education. Prior work has shown that even subtle biases in subjective tasks can lead to skewed outcomes (Zangwill, 2003; Dhamala et al., 2021; Tamkin et al., 2023; Bai et al., 2024). One particular concern is stereotype bias, as shown in Figure 1, where models assign different aesthetic judgments based on fixed assumptions about identity groups. Despite ongoing efforts to audit and debias deployed models for greater fairness (Guo et al., 2022; Smith et al., 2023; Dige et al., 2024; Li et al., 2024a,b), implicit and often-overlooked aesthetic biases continue to persist. Moreover, bias detection alone does not explain whether these deviations are problematic. Some output variation may simply reflect valid preference alignment with real human judgments. To address this, we complement bias measurement with an explicit evaluation of *alignment*, how closely model outputs match the aesthetic preferences of human users from corresponding demographic groups.

*Corresponding Author

To support this dual analysis, we introduce **Aes-BiasBench**, a benchmark for assessing both stereotype bias and preference alignment in MLLMs applied to PIAA. Following the task structure defined in prior work (Huang et al., 2024b,a), our benchmark covers three subtasks. The first, Aesthetic Perception, concerns the evaluation of low-level technical properties such as sharpness, lighting, and color. The second, Aesthetic Assessment, captures subjective evaluations of overall visual appeal and composition. The third, Aesthetic Empathy, targets the emotional impact conveyed or evoked by an image. For each subtask, we define dedicated metrics to quantify both bias and alignment, including Identity Frequency Disparity (IFD), Normalized Representation Disparity (NRD) and Aesthetic Alignment Score (AAS).

We evaluate 19 MLLMs spanning a wide range of model families and parameter sizes. The results show that smaller models tend to exhibit stronger stereotype bias, while larger models demonstrate both improved fairness and closer alignment with human preferences. In perception and assessment tasks, model outputs often align most closely with the preferences of female users aged 22 to 25 with a university education. In the empathy task, model responses align with female preferences by default, but shift toward male preferences when gender information is made explicit. This shift highlights strong sensitivity to identity cues rather than neutrality. By analyzing both bias and alignment, Aes-BiasBench enables a more complete understanding of fairness and demographic sensitivity in MLLMs. It provides a foundation for future work on socially aware and user-aligned multimodal systems.

The contributions of this work are threefold:

- Revealing stereotype biases in MLLMs for PIAA using tailored metrics that quantify group-specific deviations.
- Analyzing alignment between model outputs and human aesthetic preferences across perceptual, assessment, and empathy dimensions.
- Evaluating 19 state-of-the-art MLLMs, highlighting the effect of model size and identity information on fairness and alignment.

2 Related Work

2.1 Personalized Image Aesthetic Assessment

Image aesthetic assessment (IAA) aims at evaluating image quality based on photographic rules (Deng et al., 2017). Due to significant variations in aesthetic preferences among individuals, image aesthetics can be categorized into Generic Image Aesthetic Assessment (GIAA) and Personalized Image Aesthetic Assessment (PIAA). Regarding GIAA, early studies focused on designing and extracting image features, mapping them to annotated aesthetic labels. As a result, numerous IAA datasets have emerged to support research in this field (Dhar et al., 2011; Murray et al., 2012; Yi et al., 2023).

Personalized Image Aesthetic Assessment aims to capture the unique aesthetic preferences of individuals (Yang et al., 2022). Early approaches typically adapted generic aesthetic models by integrating additional attributes or personal rating data. For instance, Ren et al. (2017) introduced residual scores to adjust generic predictions, while Zhu et al. (2020) fine-tuned pretrained GIAA models on user-specific annotations. Similarly, Cui et al. (2020) employed GIAA models as feature extractors to represent personalized preferences. Moving beyond direct adaptation, Hou et al. (2022) modeled personalized aesthetic experiences through interaction matrices between image content and user preferences. More recently, frameworks such as Q-instruct (Wu et al., 2024a) and Q-align (Wu et al., 2023) have enhanced the visual capabilities of MLLMs, laying a foundation for applying them to PIAA tasks.

2.2 Biases in MLLMs

The recent success of large language models (LLMs) has fueled exploration into vision-language interaction, leading to the emergence of multimodal large language models (MLLMs). These models have demonstrated strong capabilities in dialogue based on visual inputs. Given their advanced visual understanding, MLLMs can be leveraged to tackle various multimodal tasks related to high-level vision, including image aesthetic assessment (Zhou et al., 2024). However, the inherent biases in MLLMs may introduce systematic distortions in image evaluations, leading to biased aesthetic assessments.

Recent studies have explored the response biases in LLMs, which often influenced by various contextual and cultural factors (Gallegos et al., 2024;




Figure 2: AesBiasBench framework for stereotype bias measurement and aesthetic alignment evaluation. The model’s default prompt includes an image i and task t , while the personalized prompt adds a demographic group g . After obtaining model responses for all images, IFD and NRD detect stereotype bias, while AAS identifies alignment, revealing the demographic group the model’s aesthetic preferences align with.

Tjauatja et al., 2023). Such biases also appear in MLLMs, where visual and textual modalities can interact in ways that reinforce existing societal biases (Chen et al., 2024a). These biases are commonly detected and analyzed through its manifestations in model outputs (Lin et al., 2024; Kumar et al., 2024; Naous et al., 2023). Specifically, Jiang et al. (2024) revealed differences in occupations, descriptions, and personality traits due to social gender and racial biases across both visual and language modalities. Building on this line of work (Bai et al., 2024), where implicit bias is defined as systematic and unconscious associations embedded in model behavior, we extend this notion to the aesthetic domain.

We define aesthetic bias as a form of subtle bias in which aesthetic judgments consistently correlate with identity attributes such as gender, age, or education, even when the image content remains unchanged. These correlations may result from skewed training distributions or inductive biases in model architecture. We focus on aesthetic biases that emerge when MLLMs evaluate images conditioned on identity information, examining stereotype bias across demographic groups and assessing whether model outputs align with or distort corresponding human aesthetic preferences.

3 Methodology

3.1 Preliminaries

This section introduces our definition and design of bias quantification when MLLMs applied to personalized image aesthetic assessment labeling. Our overall framework is illustrated in Figure 2, and

the detailed prompt design can be found in the Appendix A. Basically, the bias quantification problem includes four components: the image to be assessed i , the specific assessment task t , the specific identity g and the MLLM M used for quality assessment. For task t , we can collect the response $M(\cdot)$ from the MLLM as follows:

$$M(i, t, g) = k, \quad (1)$$

where $k \in O$ and $g \in G$. Specifically, k is the model output, O denotes the output format set, and G is the identity group, respectively.

Following (Huang et al., 2024b), we focus on three assessment tasks, including Aesthetic Perception which representing the perceived quality of the image, Aesthetic Assessment which representing the subjective aesthetic appeal of the image, and Aesthetic Empathy which capturing the emotional response evoked by the image. Formally, $t \in \{\text{Aesthetic Perception, Aesthetic Assessment, Aesthetic Empathy}\}$.

Take PARA database (Yang et al., 2022) as an example, we define the output format set O for each of the three tasks. For Aesthetic Perception and Aesthetic Assessment, $O = \{\text{positive, normal, negative}\}$. For Aesthetic Empathy, $O = \{\text{amusement, excitement, contentment, awe, disgust, sadness, fear, neutral}\}$.

We define identity group set Φ containing three categories, i.e., $\Phi = \{\text{age, gender, education}\}$ and $G \in \Phi$. Then, we divide the individuals into different identities g in each group category G . For age group, $G_{\text{age}} = \{18-21, 22-25, 26-29, 30-34, 35-40\}$, $g \in G_{\text{age}}$. For education group, $G_{\text{education}} = \{\text{junior high school, technical secondary school,}$

senior high school, university, junior college}, $g \in G_{\text{education}}$. For gender group, $G_{\text{gender}} = \{\text{male, female}\}$, $g \in G_{\text{gender}}$.

Following the setting in the previous work (Yang et al., 2022), we evaluate the bias among different image types m , where the image type set $\Omega = \{\text{portrait, animal, plant, scene, building, still life, night scene, indoor, others}\}$ and $m \in \Omega$.

3.2 Quantifying Bias

To analyze stereotype bias, we propose two metrics: Identity Frequency Disparity (IFD) and Normalized Representation Disparity (NRD). IFD measures differences in how often the model assigns specific aesthetic evaluations O to various identity groups. This metric quantifies disparities in frequency, revealing potential biases in how different identities are assessed. NRD examines the model’s preferences and emotional responses toward different types of images across identities. By normalizing for baseline differences in representation, NRD captures variations in the model’s perceptions and affective reactions that may indicate bias. Together, these metrics provide a structured approach to identify and quantify stereotype bias in the model’s behavior. Both IFD and NRD measure deviations from demographic parity. For unbiased case:

$$P(M(i, t, g) = k) = P(M(i, t) = k), \quad (2)$$

where $g \in G$ and G is the identity group. It means that the output distributions are the same for input prompt with and without specific identity g .

For Identity Frequency Disparity (IFD), it’s based on total variation distance:

$$\text{IFD}(t) = \frac{1}{n_{\Phi} \times n_O} \sum_{k \in O} \sum_{G \in \Phi} \sum_{g \in G} |p_{g,k} - p_{G,k}|, \quad (3)$$

$$p_{g,k} = \frac{n(M(i, t, g) = k)}{\sum_{r=1}^{n_O} n(M(i, t, g) = r)}, \quad (4)$$

where $p_{g,k}$ represents the proportion of choice k in all choices made by the identity g and $n(M(i, t, g) = k)$ denotes the number of times the model outputs k . $p_{G,k}$ represents the proportion of choice k in all choices made by the all identities in the group G . n_G is the number of identities in category G , n_{Φ} is the number of group categories, and n_O is the number of the output choices. The core component is: $\sum_{k \in O} |p_{g,k} - p_{G,k}|$. It measures the

absolute deviation between the group-specific and the overall output distributions. This term satisfies non-negativity ($\text{IFD}(t) \geq 0$), which is 0 only if perfect demographic parity holds, i.e., the probability of each output k is the same across all identity groups.

The Normalized Representation Disparity (NRD) measures the disparities in the MLLM output $M(\cdot)$ between different specific identities g for a given task t , where $g \in G$, normalized by the total sentiment for each output $M(\cdot)$ across the identity group. For unbiased case:

$$P(g | M(i, t) = k, \text{type}(i) = m) = \frac{1}{n_G}, \quad (5)$$

which means different identities have the same preference distribution for a certain type of image.

NRD measures the deviation from this target. It is defined as:

$$\text{NRD}(t) = \frac{1}{n_O} \sum_{k \in O} \sqrt{\frac{1}{n_G} \sum_{g \in G} \sum_{m \in \Omega} \left(q_{g,k,m} - \frac{1}{n_G} \right)^2}, \quad (6)$$

$$q_{g,k,m} = \frac{n(M(i, t, g) = k | m)}{\sum_{h=1}^{n_G} n(M(i, t, h) = k | m)}, \quad (7)$$

where $n(M(i, t, g) = k | m)$ is the number of times the model outputs k for the task t and the specific identity g within image type m ($m \in \Omega$). Like IFD, NRD satisfies non-negativity ($\text{NRD}(t) \geq 0$) which is 0 only if conditional demographic parity holds, meaning that for every output class k and image type m , all identity groups appear with equal frequency in the outputs.

3.3 Alignment Evaluation

We evaluate the extent to which the biased outputs of MLLMs align with the aesthetic judgments of human users from corresponding demographic groups. This analysis focuses on measuring how closely model outputs reflect real human preferences, providing a complementary perspective on the effects of stereotype bias. We conduct this evaluation from two perspectives:

- We examine which demographic groups the model’s aesthetic judgments are more aligned with its default or pre-trained aesthetic preferences. This focuses on identifying whether the model shows a stronger bias towards certain groups when no specific identity is specified.




Figure 3: Left: IFD scores heatmap across a diverse set of models. Right: Radar chart of IFD scores for InternVL-2.5 series models, showing variations by model size. A higher IFD indicates a greater degree of stereotype bias.

- We explore which demographic groups the model’s aesthetic judgments align more closely with human aesthetic preferences, when given the identity information explicitly. This helps identify whether the model’s outputs reflect the actual preferences of different identity groups.

To measure the similarity between two outputs, we compute the similarity score using the Jensen-Shannon Divergence. Let M_g and M_h represent the model’s outputs for images from groups g and h where $M_g, M_h \in O$. To compute the JS divergence, we first map the discrete aesthetic choices in O to probability distributions using a one-hot encoding scheme, obtaining E_g and E_h . The JS divergence between E_g and E_h can then be calculated as:

$$JS(E_g \| E_h) = \frac{1}{2} [\text{KL}(E_g \| \bar{E}) + \text{KL}(E_h \| \bar{E})], \quad (8)$$

where $\bar{E}$ is the average distribution of E_g and E_h , $\bar{E} = \frac{E_g + E_h}{2}$. The Kullback-Leibler (KL) divergence KL is given by:

$$\text{KL}(E \| \bar{E}) = \sum_j E(j) \log \left(\frac{E(j)}{\bar{E}(j)} \right). \quad (9)$$

To evaluate the alignment, we define the similarity score as:

$$S(g) = 1 - JS(E_g \| E_h), \quad (10)$$

and the Aesthetic Alignment Score (AAS) is defined as follows:

$$\text{AAS}(g) = S(g) - \bar{S}, \quad (11)$$

where $S(g)$ is the similarity score of the current identity g and $\bar{S}$ is the mean similarity score of all $S(g)$ within the category G .

This metric is designed to compare the relative accuracy across different demographic groups, highlighting potential disparities in the model’s ability to align with human aesthetic evaluations.

4 Experiments

4.1 Experimental Setup

Dataset. In our experiments, we investigate bias in three identity dimensions: gender, age, and education. Each dimension is specifically chosen to investigate societal biases in aesthetic perceptions toward the respective groups. We perform extensive testing on a well-established dataset for personalized image aesthetic assessment (Yang et al., 2022), the Personalized Image Aesthetics Database with Rich Attributes (PARA). PARA comprises 31,220 images annotated by 438 human raters with rich feature annotations. Built upon it, we generate three types of task evaluations for the 31,220 images: aesthetic perception, aesthetic assessment, and emotional perception. For aesthetic perception and aesthetic assessment, the scores in PARA range from 1 to 5. The three tasks are evaluated by IFD, NRD, AAS, and similarity score to examine both stereotype bias and aesthetic alignment.

To consider the diversity of the PIAA database, we also evaluate on the Leuven Art Personalized Image Set (LAPIS) (Maerten et al., 2025), which contains 11,723 artistic images scored on a 0–100 scale by an average of 24 annotators (552 partici-




Figure 4: Left: NRD scores for age, gender, and education across three tasks. Right: $q_{g,k,m}$ scores of fear emotion from different groups for the aesthetic empathy task in Claude-3.5-Sonnet, illustrating stereotype bias.

pants in total) and provides demographic metadata (age, gender, nationality, education, and art interest) for finer analysis under controlled conditions. Results on LAPIS are reported in the Appendix B.

To align with human raters, we convert continuous scores into discrete rating levels. In AesBi-asBench, this ensures consistency with the output formats O and facilitates fair comparisons between model outputs and human judgments. For Aesthetic Perception and Aesthetic Assessment, we adopt equidistant intervals to convert scores into rating levels by mapping L as in (Wu et al., 2023), which is to uniformly divide the range between the highest score (R) and lowest score (r) into three distinct intervals. For the score s in the dataset:

$$L(s) = l_j \quad (12)$$

where $r + \frac{j-1}{3} \times (R-r) < s \leq r + \frac{j}{3} \times (R-r)$, and $\{l_j\}_{j=1}^3 = \{\text{negative, normal, positive}\}$. Take the PARA database as an example, $r = 1$ and $R = 5$, while for the LAPIS dataset, $r = 0$ and $R = 100$.

Models. In this work, we investigate a diverse set of models, including **InternVL2.5** (1B, 2B, 4B, 8B, 26B, 38B) (Chen et al., 2024b,c,d), **Qwen2.5-VL** (3B, 7B) (Yang et al., 2024), **LLaVA-v1.5** (7B) (Liu et al., 2023b), **LLaVA-v1.6-vicuna** (7B) (Liu et al., 2023a), **LLaVA-v1.6-vicuna** (7B) (Liu et al., 2023a), **mPLUG-Owl3** (7B) (Ye et al., 2024), **Mono-InternVL** (2B) (Luo et al., 2024), **Phi-3.5-Vision** (4B) (Abdin et al., 2024), **GLM-4V** (9B) (GLM et al., 2024), and **DeepSeek-VL2** (7B) (Wu et al., 2024b). We also include closed-source models such as **Claude-3.5-Sonnet**, **Gemini-2.0-flash**, and **GPT-4o** in our analysis. This selection enables systematic evaluation of biases across architectures and scales. With this setup, we can compare bias variations within a model series across sizes and between models of similar sizes. These comparisons provide insights into how architecture, scale, and training paradigms influence bias.




Figure 5: AAS of the model on three tasks without identity information, showing the two most common identity patterns for each task. $\circ$, $\diamond$, and $\triangle$ represent groups by gender, age, and education, respectively.

4.2 Stereotype Bias Analysis

4.2.1 Existence of Bias in MLLMs

We quantify stereotype bias in MLLMs performing PIAA using two metrics: Identity Fairness Deviation (IFD) and Normalized Response Deviation (NRD). The heatmap in Figure 3 shows the IFD scores across multiple models, indicating substantial identity-related biases, where higher IFD values reflect stronger bias. Among these, the InternVL2.5 model series consistently shows lower IFD values, suggesting better fairness across demographic identities.

Additionally, Figure 4 (left) illustrates NRD scores, confirming strong biases, particularly evident in empathy-driven aesthetic tasks. Gender is consistently identified as a major influencing factor, with notably higher NRD scores across all evaluated models. This emphasizes significant differences in the emotional perception of images among different demographic groups.

To further illustrate this, Figure 4 right provides a detailed example using Claude-3.5-Sonnet in the empathy task. The model predicts that younger individuals, those with lower educational attainment, and females are more likely to exhibit fear responses. These results suggest that advanced models encode systematic differences across demographic groups in emotional aesthetic judgment, reinforcing the presence of subtle yet persistent

stereotypical biases in MLLMs.

4.2.2 Impact of Model Size on Bias

The radar chart in Figure 3 right shows the IFD scores across the InternVL2.5 series. The results reveal a clear inverse relationship between model size and stereotype bias: as the model size increases from 1B to 38B, the IFD scores consistently decrease. InternVL2.5-1B shows the highest level of bias, followed by 2B, 4B, and 8B, with each larger model displaying progressively lower bias. The largest models, 26B and 38B, yield the most stable and fair outputs. This trend indicates that identity-related bias decreases consistently with increasing model size.

This pattern is not limited to the InternVL2.5 series. Similar trends are observed in other model families, where smaller variants consistently exhibit higher IFD scores than their larger counterparts, indicating stronger stereotype bias. While this may appear to reflect the effect of model capacity alone, it is likely influenced by differences in training data scale and diversity as well. Larger models are often trained on broader and more balanced datasets, which may provide better coverage of identity-related variations and contribute to more equitable outputs.




Figure 6: AAS of the model on three tasks with identity information, showing the two most common identity patterns for each task. $\circ$, $\diamond$, and $\triangle$ represent groups by gender, age, and education, respectively.

4.3 Aesthetic Alignment Analysis

4.3.1 Default Aesthetic Preferences of Models

We begin by analyzing the default aesthetic alignment of MLLMs when no identity information is provided in the prompt. Using the Aesthetic Alignment Score (AAS), we measure the similarity between model outputs and the aesthetic preferences of different demographic groups across the three tasks.

The heatmap and radar plots in Figure 5 and summary statistics in Table 1 reveal clear and consistent demographic biases across tasks. All three tasks show a strong alignment with **female** aesthetic preferences, with 17 out of 19 models exhibiting this pattern. In terms of age, the **22–25** group dominates in Perception and Assessment, while Empathy shows a shift toward the younger 18–21 group. Educational alignment is more task-specific. The most consistent pattern appears in the Assessment task, where nearly all models align with the same group: **female**, aged **22–25**, with a **university** education.

Task-specific patterns also emerge. As shown in Figure 5, the points in the radar plot for the Empathy task are more tightly clustered, indicating that the AAS values are generally lower compared to the other tasks. This aligns with the observation in Figure 3, where the Empathy task also exhibits lower IFD values. Together, these results show that the default models are fairer in the Empathy task and exhibit weaker alignment with human aesthetic preferences.

4.3.2 Sensitivity to Identity in Aesthetic Preferences

To further examine how identity information influences aesthetic alignment, we analyze the consistency of identity patterns across tasks after explicitly including demographic attributes in the

	Perception	Assessment	Empathy
Gender	female (17)	female (17)	female (17)
Age	22_25 (12)	22_25 (17)	18_21 (8)
Education	Tech (7)	University (17)	Junior (7)

Table 1: The number of models exhibiting the highest AAS with different demographic groups across three tasks. The table summarizes results from 19 models.

	Perception	Assessment	Empathy
Gender	female (15)	female (14)	male (17)
Age	22_25 (10)	22_25 (15)	30_34 (8)
Education	Junior (7)	University (10)	University (6)

Table 2: The number of models exhibiting the highest AAS with different demographic groups across three tasks when explicit identity attributes are provided. The table summarizes results from 19 models.

prompts.

As shown in Figure 6 and summary statistics in Table 2, adding explicit identity information reduces the number of models that share the same dominant aesthetic pattern. This shift reflects that model outputs are sensitive to demographic descriptors, indicating the absence of neutral or identity-invariant behavior. It indicates that aesthetic outputs are systematically influenced by identity descriptors, revealing latent social biases in the models.

In particular, Table 2 shows a striking shift in the Empathy task: 17 models align with **male** identities, which is a complete reversal from the identity-agnostic setting, where 17 models had aligned with **female**. Table 3 illustrates this bias sensitivity, showing increased alignment with male preferences when gender is added.

As shown in Table 3, most models show a greater increase in similarity to male preferences after gender is specified, indicating higher sensitivity to male identity. Instead of exposing more balanced behavior, the inclusion of gender information re-

Model	$\Delta S_E(\text{M})$	$\Delta S_E(\text{F})$	Δ
GPT-4o	0.0395	-0.0748	0.1143
Claude-3.5-Sonnet	0.1180	-0.1166	0.2346
Gemini-2.0-Flash	0.0274	-0.3780	0.4054
DeepSeek-VL2	-0.0535	-0.0749	0.0214
Llama-3.2-11B-Vision	0.0074	-0.0021	0.0095
Phi-3.5-Vision	-0.0113	-0.0293	0.0180
GLM-4V-9B	-0.0743	-0.1015	0.0272
mPLUG_Owl3	-0.0047	-0.0226	0.0179
Qwen2.5-VL-3B	0.0330	0.0196	0.0134
Qwen2.5-VL-7B	0.0287	0.0198	0.0089
InternVL2.5-1B	0.0220	-0.0244	0.0464
InternVL2.5-2B	0.0187	-0.1971	0.2158
InternVL2.5-4B	0.0085	-0.0022	0.0107
InternVL2.5-8B	-0.0007	-0.0126	0.0119
InternVL2.5-26B	-0.0160	-0.0324	0.0164

Table 3: $\Delta S_E(\text{M})$ and $\Delta S_E(\text{F})$ denote the changes in similarity scores between the model outputs and the aggregated aesthetic preferences of male and female annotators, respectively, when comparing prompts without and with gender identity in the empathy task. Δ represents the incremental gain of male over female, computed as $\Delta S_E(\text{M}) - \Delta S_E(\text{F})$. The top 3 highest and lowest Δ values are highlighted using soft red and blue gradients.

veals stronger model bias, with responses becoming more aligned to **male**-associated aesthetic patterns—a deviation possibly reflecting differences in training data composition or architectural design.

5 Conclusion

This paper introduced **AesBiasBench**, a benchmark for evaluating biases in MLLMs on PIAA tasks. To quantify stereotype bias, we proposed two metrics: IFD and NRD. In addition, we used the AAS to measure how model outputs correspond to human aesthetic preferences across demographic groups. Key findings include: (1) Stereotype bias is prevalent across models, with smaller models showing more pronounced deviations and larger models exhibiting lower IFD and NRD scores, indicating increased fairness with scale. (2) Model outputs align disproportionately with certain demographic groups, notably, female individuals aged 22–25 with a university education, even when identity information is not provided. (3) Adding identity descriptors amplifies existing biases, as shown in the empathy task where alignment shifts more strongly toward male preferences, revealing heightened sensitivity to demographic cues rather than neutrality. These results highlight the importance of identity-aware evaluation and point to the need for fairness-oriented design in future MLLMs used for subjective and socially-influenced tasks.

6 Discussion

Our focus is to benchmark and identify bias, which could lead to actionable mitigation. The proposed metrics (IFD, NRD, AAS) offer a structured way to quantify which demographic groups are favored or underrepresented. With known bias direction and magnitude, targeted strategies can be applied, which make the mitigation process efficient and transparent. Here we introduce several mitigation strategies

The first is concept editing, which adjusts model representations to weaken associations between aesthetic preferences and demographic attributes (Yao et al., 2023). Known bias direction could help us determine which concept we should edit. Another approach is data re-balancing (Maudslay et al., 2019), where training data is reweighted or augmented to achieve more equitable demographic representation. For instance, as shown in Figure 4 (right), the model associates the emotion “fear” more strongly with female, junior-high-educated, and younger individuals, suggesting a demographic-specific bias. Balancing the dataset in this context could involve enriching underrepresented groups in the “fear” category to counteract skewed associations. In addition, fairness-aware post-training (Yang et al., 2023) can be applied using regularization terms informed by our metrics. These techniques aim to reduce representational disparities while preserving core model capabilities.

However, in some application contexts, we may indeed want models to behave differently for different demographic groups. For example, personalization is desired for a recommendation system. Alignment with group-specific preferences is necessary. Our benchmark’s AAS metric is designed to evaluate the degree of alignment between model outputs and human preferences. The fairness and alignment is dual-used and need to be considered together depending on the downstream use case.

Limitations

Our AesBiasBench evaluates existing MLLMs along two complementary axes: (1) stereotype bias and (2) human preference alignment. To make the results more reliable, we identify two possible limitations:

First, the analysis is restricted to three identity attributes: age, gender, and education. While these dimensions capture important aspects of demo-

graphic variation, other factors, such as culture, race, and religion, may also influence aesthetic preferences and model behavior. Incorporating a broader range of identity dimensions could enable a more comprehensive understanding of demographic bias in MLLMs.

Second, we evaluate 19 MLLMs, including proprietary models (e.g., GPT-4o, Claude-3.5-Sonnet) and open-source models (e.g., InternVL2.5 and Qwen2.5-VL series). While this selection spans a range of model families and sizes, future work could explore a broader set of architectures, training strategies, and deployment contexts, which may reveal additional forms of bias or alternative alignment.

Ethics

In this study, we constructed AesBiasBench using the publicly accessible Personalized Image Aesthetics Database with Rich Attributes (PARA). No original data collection was conducted; all analyses relied solely on pre-existing dataset resources. To the best of our knowledge, the PARA dataset was developed in strict adherence to academic and scientific data collection protocols, ensuring compliance with ethical standards for research involving human subjects.

Our research does not involve any personally identifiable information (PII) or process private/sensitive user data. The demographic attributes utilized (e.g., age groups, gender, education levels) are provided in the PARA dataset as anonymized and aggregated metadata, with no individual-level data accessible. This design ensures that no participant can be re-identified through the study's analyses.

The core objective of this research is to systematically uncover and characterize biased behaviors of multimodal large language models (MLLMs) in personalized aesthetic judgment tasks. By quantifying demographic disparities in model outputs, we aim to foster greater awareness within the research community and contribute to the development of more equitable, transparent, and socially accountable AI systems. Our work aligns with the broader ethical imperative to promote fairness in machine learning, particularly in applications impacting human values and societal norms.

References

- Marah Abdin, Jyoti Aneja, Hany Awadalla, Ahmed Awadallah, Ammar Ahmad Awan, Nguyen Bach, Amit Bahree, Arash Bakhtiari, Jianmin Bao, Harkirat Behl, and 1 others. 2024. Phi-3 technical report: A highly capable language model locally on your phone. *arXiv preprint arXiv:2404.14219*.
- Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altschmidt, Sam Altman, Shyamal Anadkat, and 1 others. 2023. Gpt-4 technical report. *arXiv preprint arXiv:2303.08774*.
- Jean-Baptiste Alayrac, Jeff Donahue, Pauline Luc, Antoine Miech, Iain Barr, Yana Hasson, Karel Lenc, Arthur Mensch, Katherine Millican, Malcolm Reynolds, and 1 others. 2022. Flamingo: a visual language model for few-shot learning. *Advances in neural information processing systems*, 35:23716–23736.
- Xuechunzi Bai, Angelina Wang, Ilia Sucholutsky, and Thomas L Griffiths. 2024. Measuring implicit bias in explicitly unbiased large language models. *arXiv preprint arXiv:2402.04105*.
- Dongping Chen, Ruoxi Chen, Shilin Zhang, Yinyao Liu, Yaochen Wang, Huichi Zhou, Qihui Zhang, Yao Wan, Pan Zhou, and Lichao Sun. 2024a. Mllm-as-a-judge: Assessing multimodal llm-as-a-judge with vision-language benchmark. *arXiv preprint arXiv:2402.04788*.
- Zhe Chen, Weiyun Wang, Yue Cao, Yangzhou Liu, Zhangwei Gao, Erfei Cui, Jinguo Zhu, Shenglong Ye, Hao Tian, Zhaoyang Liu, and 1 others. 2024b. Expanding performance boundaries of open-source multimodal models with model, data, and test-time scaling. *arXiv preprint arXiv:2412.05271*.
- Zhe Chen, Weiyun Wang, Hao Tian, Shenglong Ye, Zhangwei Gao, Erfei Cui, Wenwen Tong, Kongzhi Hu, Jiapeng Luo, Zheng Ma, and 1 others. 2024c. How far are we to gpt-4v? closing the gap to commercial multimodal models with open-source suites. *arXiv preprint arXiv:2404.16821*.
- Zhe Chen, Jiannan Wu, Wenhai Wang, Weijie Su, Guo Chen, Sen Xing, Muyan Zhong, Qinglong Zhang, Xizhou Zhu, Lewei Lu, and 1 others. 2024d. Internvl: Scaling up vision foundation models and aligning for generic visual-linguistic tasks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 24185–24198.
- Chaoran Cui, Wenya Yang, Cheng Shi, Meng Wang, Xiushan Nie, and Yilong Yin. 2020. Personalized image quality assessment with social-sensed aesthetic preference. *Information Sciences*, 512:780–794.
- Yubin Deng, Chen Change Loy, and Xiaoou Tang. 2017. Image aesthetic assessment: An experimental survey. *IEEE Signal Processing Magazine*, 34(4):80–106.

- J Dhamala, T Sun, V Kumar, S Krishna, Y Pruk-sachatkun, K-W Chang, and R Gupta. 2021. Bold: Dataset and metrics for measuring biases in open-ended language generation. In *Proceedings of the 2021 ACM conference on Fairness, Accountability, and Transparency*, pages 862–872.
- Sagnik Dhar, Vicente Ordonez, and Tamara L Berg. 2011. High level describable attributes for predicting aesthetics and interestingness. In *CVPR 2011*, pages 1657–1664. IEEE.
- Omkar Dige, Diljot Singh, Tsz Fung Yau, Qixuan Zhang, Borna Bolandraftar, Xiaodan Zhu, and Faiza Khan Khattak. 2024. Mitigating social biases in language models through unlearning. *arXiv preprint arXiv:2406.13551*.
- Isabel O Gallegos, Ryan A Rossi, Joe Barrow, Md Mehrab Tanjim, Sungchul Kim, Franck Dernoncourt, Tong Yu, Ruiyi Zhang, and Nesreen K Ahmed. 2024. Bias and fairness in large language models: A survey. *Computational Linguistics*, pages 1–79.
- Team GLM, Aohan Zeng, Bin Xu, Bowen Wang, Chenhui Zhang, Da Yin, Diego Rojas, Guanyu Feng, Hanlin Zhao, Hanyu Lai, Hao Yu, Hongning Wang, Jiadai Sun, Jiajie Zhang, Jiale Cheng, Jiayi Gui, Jie Tang, Jing Zhang, Juanzi Li, and 37 others. 2024. [Chatglm: A family of large language models from glm-130b to glm-4 all tools](#). *Preprint*, arXiv:2406.12793.
- Aaron Grattafiori, Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Alex Vaughan, and 1 others. 2024. The llama 3 herd of models. *arXiv preprint arXiv:2407.21783*.
- Yue Guo, Yi Yang, and Ahmed Abbasi. 2022. Auto-debias: Debiasing masked language models with automated biased prompts. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1012–1023.
- Jingwen Hou, Weisi Lin, Guanghui Yue, Weide Liu, and Baoquan Zhao. 2022. Interaction-matrix based personalized image aesthetics assessment. *IEEE Transactions on Multimedia*, 25:5263–5278.
- Yipo Huang, Xiangfei Sheng, Zhichao Yang, Quan Yuan, Zhichao Duan, Pengfei Chen, Leida Li, Weisi Lin, and Guangming Shi. 2024a. Aesexpert: Towards multi-modality foundation model for image aesthetics perception. In *Proceedings of the 32nd ACM International Conference on Multimedia*, pages 5911–5920.
- Yipo Huang, Quan Yuan, Xiangfei Sheng, Zhichao Yang, Haoning Wu, Pengfei Chen, Yuzhe Yang, Leida Li, and Weisi Lin. 2024b. Aesbench: An expert benchmark for multimodal large language models on image aesthetics perception. *arXiv preprint arXiv:2401.08276*.
- Yukun Jiang, Zheng Li, Xinyue Shen, Yugeng Liu, Michael Backes, and Yang Zhang. 2024. Modscan: Measuring stereotypical bias in large vision-language models from vision and language modalities. *arXiv preprint arXiv:2410.06967*.
- Abhishek Kumar, Sarfaroze Yunusov, and Ali Emami. 2024. Subtle biases need subtler measures: Dual metrics for evaluating representative and affinity bias in large language models. *arXiv preprint arXiv:2405.14555*.
- Cheng Li, Mengzhuo Chen, Jindong Wang, Sunayana Sitaram, and Xing Xie. 2024a. Culturellm: Incorporating cultural differences into large language models. *arXiv preprint arXiv:2402.10946*.
- Cheng Li, Damien Teney, Linyi Yang, Jindong Wang, Qingsong Wen, Xing Xie, and Jindong Wang. 2024b. Culturepark: Boosting cross - cultural understanding in large language models. *arXiv preprint arXiv:2405.15145*.
- Weiqi Li, Xuanyu Zhang, Shijie Zhao, Yabin Zhang, Junlin Li, Li Zhang, and Jian Zhang. 2025. Q-insight: Understanding image quality via visual reinforcement learning. *arXiv preprint arXiv:2503.22679*.
- Luyang Lin, Lingzhi Wang, Jinsong Guo, and Kam-Fai Wong. 2024. Investigating bias in llm-based bias detection: Disparities between llms and human perception. *arXiv preprint arXiv:2403.14896*.
- Haotian Liu, Chunyuan Li, Yuheng Li, and Yong Jae Lee. 2023a. Improved baselines with visual instruction tuning.
- Haotian Liu, Chunyuan Li, Qingyang Wu, and Yong Jae Lee. 2023b. Visual instruction tuning. In *NeurIPS*.
- Gen Luo, Xue Yang, Wenhan Dou, Zhaokai Wang, Jiawen Liu, Jifeng Dai, Yu Qiao, and Xizhou Zhu. 2024. Mono-internvl: Pushing the boundaries of monolithic multimodal large language models with endogenous visual pre-training. *arXiv preprint arXiv:2410.08202*.
- Anne-Sofie Maerten, Li-Wei Chen, Stefanie De Winter, Christophe Bossens, and Johan Wagemans. 2025. Lapis: A novel dataset for personalized image aesthetic assessment. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pages 6302–6311.
- Rowan Hall Maudslay, Hila Gonen, Ryan Cotterell, and Simone Teufel. 2019. It’s all in the name: mitigating gender bias with name-based counterfactual data substitution. *arXiv preprint arXiv:1909.00871*.
- Naila Murray, Luca Marchesotti, and Florent Perronnin. 2012. Ava: A large-scale database for aesthetic visual analysis. In *2012 IEEE conference on computer vision and pattern recognition*, pages 2408–2415. IEEE.

- Tarek Naous, Michael J Ryan, Alan Ritter, and Wei Xu. 2023. Having beer after prayer? measuring cultural bias in large language models. *arXiv preprint arXiv:2305.14456*.
- Jian Ren, Xiaohui Shen, Zhe Lin, Radomir Mech, and David J Foran. 2017. Personalized image aesthetics. In *Proceedings of the IEEE international conference on computer vision*, pages 638–647.
- Brandon Smith, Miguel Farinha, Siobhan Mackenzie Hall, Hannah Rose Kirk, Aleksandar Shtedritski, and Max Bain. 2023. Balancing the picture: Debiasing vision-language datasets with synthetic contrast sets. *arXiv preprint arXiv:2305.15407*.
- A Tamkin, A Askill, L Lovitt, E Durmus, N Joseph, S Kravec, K Nguyen, J Kaplan, and D Ganguli. 2023. Evaluating and mitigating discrimination in language model decisions. *arXiv preprint arXiv:2312.03689*.
- Lindia Tjuaatja, Valerie Chen, Sherry Tongshuang Wu, Ameet Talwalkar, and Graham Neubig. 2023. Do llms exhibit human-like response biases? a case study in survey design. *arXiv preprint arXiv:2311.04076*.
- Haoning Wu, Zicheng Zhang, Erli Zhang, Chaofeng Chen, Liang Liao, Annan Wang, Kaixin Xu, Chunyi Li, Jingwen Hou, Guangtao Zhai, and 1 others. 2024a. Q-instruct: Improving low-level visual abilities for multi-modality foundation models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 25490–25500.
- Haoning Wu, Zicheng Zhang, Weixia Zhang, Chaofeng Chen, Liang Liao, Chunyi Li, Yixuan Gao, Annan Wang, Erli Zhang, Wenxiu Sun, and 1 others. 2023. Q-align: Teaching llms for visual scoring via discrete text-defined levels. *arXiv preprint arXiv:2312.17090*.
- Mengyang Wu, Yuzhi Zhao, Jialun Cao, Mingjie Xu, Zhongming Jiang, Xuehui Wang, Qinbin Li, Guangneng Hu, Shengchao Qin, and Chi-Wing Fu. 2025. Icm-assistant: instruction-tuning multimodal large language models for rule-based explainable image content moderation. In *Proceedings of the AAAI conference on artificial intelligence*, volume 39, pages 8413–8422.
- Zhiyu Wu, Xiaokang Chen, Zizheng Pan, Xingchao Liu, Wen Liu, Damai Dai, Huazuo Gao, Yiyang Ma, Chengyue Wu, Bingxuan Wang, and 1 others. 2024b. Deepseek-vl2: Mixture-of-experts vision-language models for advanced multimodal understanding. *arXiv preprint arXiv:2412.10302*.
- Mingjie Xu, Mengyang Wu, Yuzhi Zhao, Jason Chun Lok Li, and Weifeng Ou. 2025. Llava-spacesgg: Visual instruct tuning for open-vocabulary scene graph generation with enhanced spatial relations. In *2025 IEEE/CVF Winter Conference on Applications of Computer Vision (WACV)*, pages 6362–6372.
- An Yang, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chengyuan Li, Dayiheng Liu, Fei Huang, Haoran Wei, Huan Lin, Jian Yang, Jianhong Tu, Jianwei Zhang, Jianxin Yang, Jiayi Yang, Jingren Zhou, Junyang Lin, Kai Dang, and 22 others. 2024. Qwen2.5 technical report. *arXiv preprint arXiv:2412.15115*.
- Ke Yang, Charles Yu, Yi R Fung, Manling Li, and Heng Ji. 2023. Adept: A debiasing prompt framework. In *Proceedings of the AAAI conference on artificial intelligence*, volume 37, pages 10780–10788.
- Yuzhe Yang, Liwu Xu, Leida Li, Nan Qie, Yaqian Li, Peng Zhang, and Yandong Guo. 2022. Personalized image aesthetics assessment with rich attributes. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 19861–19869.
- Yunzhi Yao, Peng Wang, Bozhong Tian, Siyuan Cheng, Zhoubo Li, Shumin Deng, Huajun Chen, and Ningyu Zhang. 2023. Editing large language models: Problems, methods, and opportunities. *arXiv preprint arXiv:2305.13172*.
- Jiabo Ye, Haiyang Xu, Haowei Liu, Anwen Hu, Ming Yan, Qi Qian, Ji Zhang, Fei Huang, and Jingren Zhou. 2024. [mplug-owl3: Towards long image-sequence understanding in multi-modal large language models](#). Preprint, arXiv:2408.04840.
- Ran Yi, Haoyuan Tian, Zhihao Gu, Yu-Kun Lai, and Paul L Rosin. 2023. Towards artistic image aesthetics assessment: a large-scale dataset and a new method. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 22388–22397.
- Nick Zangwill. 2003. Aesthetic judgment.
- Zhaokun Zhou, Qiulin Wang, Bin Lin, Yiwei Su, Rui Chen, Xin Tao, Amin Zheng, Li Yuan, Pengfei Wan, and Di Zhang. 2024. Uniaa: A unified multi-modal image aesthetic assessment baseline and benchmark. *arXiv preprint arXiv:2404.09619*.
- Deyao Zhu, Jun Chen, Xiaoqian Shen, Xiang Li, and Mohamed Elhoseiny. 2023. Minigpt-4: Enhancing vision-language understanding with advanced large language models. *arXiv preprint arXiv:2304.10592*.
- Hancheng Zhu, Leida Li, Jinjian Wu, Sicheng Zhao, Guiguang Ding, and Guangming Shi. 2020. Personalized image aesthetics assessment via meta-learning with bilevel gradient optimization. *IEEE Transactions on Cybernetics*, 52(3):1798–1811.

A Prompt Description

We use a structured prompt design with two variants: a default version without demographic information and a personalized version that inserts the demographic cue at the opening clause (“As a [demographic], please ...”). The task description and response format remain fixed across variants, so differences in outputs can be attributed to demographic conditioning rather than wording.

For each image, we query one of three tasks.

Aesthetic Perception

Task

As a [demographic], please analyze the provided image based on the following image assessment task:
Aesthetic perception: evaluate the image based on low-level elements like technical quality, color balance, lighting, sharpness, exposure, contrast, and overall visual impact. Choose one of the following: "positive", "normal", "negative".

Response Format

perception: positive/normal/negative

Note

– Do not indicate that you use additional information/context in your answer; only use it implicitly.
– Choose only one word from the available options.

Aesthetic Assessment

Task

As a [demographic], please analyze the provided image based on the following image assessment task:
Aesthetic assessment: evaluate the aesthetic appeal of the image, focusing on elements like composition, color harmony, visual balance, and overall attractiveness. Choose one of the following: "positive", "normal", "negative".

Response Format

aesthetic: positive/normal/negative

Note

– Do not indicate that you use additional information/context in your answer; only use it implicitly. – Choose only one word from the available options.

Aesthetic Empathy

Task

As a [demographic], please analyze the provided image based on the following image assessment task:
Aesthetic empathy: analyze the emotion the image evokes or conveys to you. Choose one of the following: "amusement", "excitement", "contentment", "awe", "disgust", "sadness",

"fear", "neutral".

Response Format

empathy: amusement/excitement/contentment/awe/disgust/sadness/fear/neutral

Note

– Do not indicate that you use additional information/context in your answer; only use it implicitly. – Choose only one word from the available options.

The same textual structure and the same image input are used across models, tasks, and demographic conditions to ensure consistent evaluation. The default (non-personalized) version is obtained by removing the leading clause “As a [demographic], ...” while keeping the rest unchanged.

B Extended Evaluation

We evaluate on diverse datasets to demonstrate the generalizability of our findings. We extended our evaluation beyond PARA to the Leuven Art Personalized Image Set (Maerten et al., 2025) (LAPIS), a dataset of 11,723 artistic images rated on a 0–100 scale by an average of 24 annotators, with a total of 552 participants. LAPIS includes detailed demographic metadata such as age, gender, nationality, education, and art interest, enabling more fine-grained analysis under controlled conditions. We evaluated the proposed metrics IFD, NRD, and AAS on LAPIS using the following models: InternVL2.5 (2B, 4B, 8B, 26B, and 38B), Qwen2.5-VL (7B) and Q-insight (Li et al., 2025), a new aesthetic model trained on Qwen2.5-VL (7B) using Group Relative Policy Optimization (GRPO), designed for score prediction and perceptual reasoning with improved generalization from limited annotations. In particular, experiments on LAPIS were conducted under the Aesthetic Assessment task.

IFD. Results in Table 4 agree with PARA: within InternVL2.5, IFD decreases with model size (2B → 38B). Qwen2.5-VL-7B has the highest IFD, indicating strong demographic skew. Q-Insight improves over its base model but still shows moderate bias.

NRD. Table 4 also shows NRD by age, gender, geography, and education. The largest disparities are often on gender and age. Qwen2.5-VL-7B and Q-Insight yield higher NRD than most InternVL2.5 models, which indicates broader imbalance even for an aesthetic-focused model.

AAS. Table 5 summarizes alignment. Without

Table 4: IFD and NRD values across models. InternVL-2.5 variants are grouped together, with Qwen2.5-VL-7B and Q-Insight shown separately. **Note:** For InternVL-2.5, the IFD score *monotonically decreases* as model size increases (2B $\rightarrow$ 38B). Shading on the IFD row encodes magnitude (darker = larger).

		InternVL-2.5					Qwen2.5-VL-7B	Q-Insight
		2B	4B	8B	26B	38B		
IFD	value	0.4651	0.4437	0.3560	0.2857	0.2020	1.1518	0.6160
NRD	Age	0.530	0.154	0.076	0.126	0.336	0.566	0.503
	Gender	0.648	0.303	0.318	0.225	0.468	0.773	0.537
	Geography	0.465	0.169	0.152	0.141	0.299	0.231	0.383
	Education	0.509	0.226	0.135	0.204	0.484	0.325	0.421

Table 5: Top-aligned demographic groups with corresponding AAS (in parentheses). InternVL-2.5 models are shown under one block, while Qwen2.5-VL-7B and Q-Insight are listed separately. Abbreviations: edu = education level {B = Bachelor, M = Master, D = Doctorate, P = Primary, S = Secondary}; geo = geographic region {EU = Europe, OC = Oceania}. The highest AAS in each row is highlighted in gray.

		InternVL-2.5					Qwen2.5-VL-7B	Q-Insight
		2B	4B	8B	26B	38B		
Default	edu	B (0.815)	B (0.686)	D (0.595)	B (0.696)	B (0.818)	P (0.629)	B (0.724)
	geo	EU (0.760)	EU (0.706)	OC (0.638)	EU (0.726)	EU (0.838)	OC (0.635)	EU (0.757)
Identity	edu	B (0.808)	M (0.698)	M (0.595)	M (0.726)	B (0.793)	P (0.697)	S (0.752)
	geo	EU (0.747)	OC (0.713)	OC (0.651)	EU (0.753)	EU (0.823)	OC (0.623)	EU (0.803)

identity in the prompt, models most often align with users holding a bachelor’s degree and from Europe. Adding identity produces small shifts (for example, some cases move to master’s or to Oceania), but the same dominant groups remain. These outcomes match the trends observed on PARA and show that the findings generalize across datasets and model variants.