Tracing the Latent Threads: A Mechanistic Study of How LLMs Represent and Operationalize Race and Ethnicity Cues

Shiyue Hu[1,2] **Ruizhe Li**[3,*] **Yanjun Gao**[1,*]
[1]University of Colorado Anschutz
[2]University of Colorado Boulder
[3]University of Birmingham
shiyue.hu@colorado.edu, r.li.7@bham.ac.uk, yanjun.gao@cuanschutz.edu
[*]Co-senior authors

Abstract

Large language models (LLMs) increasingly operate in high-stakes settings where demographic attributes such as race and ethnicity may be explicitly stated or implicitly suggested through textual cues. However, existing studies primarily document outcome-level disparities, offering limited insight into internal mechanisms underlying these effects. We present a mechanistic study of how race and ethnicity cues are represented and operationalized within LLMs. Using two publicly available datasets spanning toxicity-related generation and clinical narrative understanding tasks, we analyze three open-source models with a reproducible interpretability pipeline combining probing, neuron-level attribution, and targeted intervention. We find that sensitivity to demographic cues is distributed across internal units and varies substantially across models. These units often align with entangled semantic facets, including explicit group labels, geography, language, culture, and associations related to stereotypes. Interventions on selected units can change some biased prediction patterns, but substantial residual effects remain, suggesting that effective mitigation requires understanding distributed, task-specific mechanisms rather than manipulating a small set of identified neurons alone.

1 Introduction

[1] Large language models (LLMs) are increasingly used in high-stakes domains such as healthcare, where demographic attributes (e.g., race, ethnicity, gender) may be explicitly stated or implicitly inferred from text. Prior work shows that LLMs can condition their outputs on demographic information even when it is not task-relevant (Zack et al., 2024; Kim et al., 2023; Fraser and Kiritchenko, 2024; Zhao et al., 2025), and can therefore induce misattribution in model output with undesirable or biased behavior (Demchak et al., 2024; Yan et al., 2026; Levartovsky et al., 2025; Li et al., 2026; Zack et al., 2024).

Most prior studies on demographic bias focus on outcome-level effects, evaluating disparities in generated responses, accuracy, calibration, or toxicity scores across demographic groups (Tan and Lee, 2025; Hartvigsen et al., 2022; Guan et al., 2025; Wang et al., 2025). While these analyses are essential for documenting harm, they treat LLMs as black boxes, offering limited insight into whether demographic attributes are reflected in high-level semantic features, task-relevant representations, or spurious shortcuts during prediction. In parallel, recent works in mechanistic interpretability demonstrated how LLMs represent demographic information and manipulated LLM internal states to ensure fairness (Yu and Ananiadou, 2025; Ahsan and Wallace, 2025; Karvonen and Marks, 2025), yet these tools have rarely been applied to demographic bias in a systematic and task-diverse manner.

A central challenge is that demographic attributes interact with language in complex ways. In many real-world settings, demographic information may be explicitly stated (e.g., "a Black patient", "a Hispanic speaker") or implicitly conveyed through linguistic, cultural, or geographical cues, i.e., "proxy" cues. Moreover, the same demographic signal can have qualitatively different effects depending on task: it may alter predicted medical risk in a clinical scenario, while simultaneously modulating perceived toxicity, credibility, or intent in open-ended generation. Existing evaluations typically isolate a single task or domain (Zack et al., 2024; Levartovsky et al., 2025), making it difficult to assess whether demographic sensitivity reflects general representational mechanisms or task-specific heuristics.

In this work, we investigate how demographic

[1]We utilized ChatGPT to identify and translate non-English tokens presented in this paper.

Findings of the Association for Computational Linguistics: EMNLP 2026, pages 29530–29545
October 24-29, 2026 ©2026 Association for Computational Linguistics

attributes influence LLM behavior, with a focus on mechanistic explanations rather than surface-level disparities. We examine *race* and *ethnicity* as commonly occurring coarse-grained categories (e.g. White, Black, Asian, Hispanic and Latino) as they appear in the studied datasets, rather than attempting to model the full sociological complexity of these constructs. Using two publicly available datasets, we study: 1) toxicity-related generation tasks (Hartvigsen et al., 2022), where the same attributes may alter the likelihood, tone, or framing of model outputs, and 2) clinical narrative tasks (Bear Don't Walk IV et al., 2024), where the same attributes appear through explicit or indirect cues in medical text and modulate model behavior despite identical clinical evidence.

We adopt a mechanistic interpretability (MI) framework to study how lexical cues associated with race and ethnicity are represented and propagated within three open-source LLMs that are widely used: Qwen2.5-7B (Qwen Team, 2024), Mistral-7B (Jiang et al., 2023), and Llama-3.1-8B (Meta AI, 2024). Our contributions are threefold:

- a reproducible MI **pipeline** that combines multi-class probing, neuron-level attribution, and targeted intervention to identify internal units associated with demographic attributes and to examine their functional relevance across tasks. The proposed framework is applicable to other social variables beyond race and ethnicity.
- a fine-grained **characterization** of internal features associated with race and ethnicity cues across LLMs, revealing the distributed nature of demographic information and model-specific emphasis on semantic facets such as geography, language, culture, or historical context.
- a mechanistic **analysis** of how demographic representation influences model behaviors. Internal features aligned with sensitive or harmful stereotype-related concepts in the pretraining data are unevenly activated by direct and indirect demographic cues.

Our findings show that while features associated with race and ethnicity cues can be identified and analyzed at the neuron level, their associations with biased model behavior persist even when highly active neurons are sign-flipped. This indicates that biased behavior in LLMs cannot be fully explained or controlled by manipulating a small set of identifiable neurons alone.

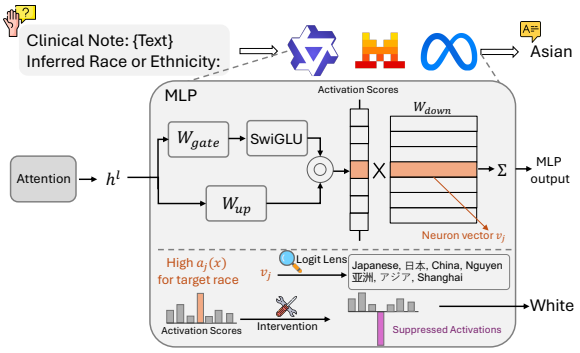


Figure 1: With MLP, we locate neurons associated with race and ethnicity cues and inspect them via Logit Lens. For neurons with high activation on target race and ethnicity cues, we intervene on their activation values to test their causal influence on model behavior.

2 Related Work

Mechanistic Interpretability of Bias in LLMs. Recent works in mechanistic interpretability have begun to locate where demographic information is represented in LLMs. Our approach of using probe-based neuron localization aligns with emerging research in this space. Yu and Ananiadou (2025) utilized neuron editing to understand and mitigate gender bias, identifying specific "gender neurons" within the MLP layers. In the medical domain, Ahsan et al. (2025) investigated the mechanisms of demographic bias specifically for healthcare tasks, suggesting that certain internal representations are disproportionately sensitive to racial identifiers. More recently, Ahsan and Wallace (2025) explored the use of Sparse Autoencoders (SAEs) to reveal clinical racial biases. Our work extends these findings by demonstrating that neurons associated with race and ethnicity cues are not only present in general datasets but are consistently activated and influential on domain-specific text.

Internal Bias Mitigation and Mechanistic Steering within LLMs. Beyond identification, recent work focuses on manipulating internal model states to ensure fairness. Zhou et al. (2024) proposed the UniBias framework, which mitigates bias by manipulating attention heads and MLP components. For real-time applications, Li et al. (2025) introduced *FairSteer*, a dynamic activation steering method that adjusts model behavior during inference. Karvonen and Marks (2025) further demonstrated that interpretability-based interventions can improve fairness more robustly than traditional fine-tuning in realistic settings.

Our methodology contributes to this line of work by using targeted intervention, specifically sign-flipping and amplification, to test whether identified race-associated neurons causally influence biased prediction patterns. This builds upon the "context-aware" fairness frameworks suggested by Nadeem et al. (2025), ensuring that mitigation is grounded in the semantic understanding of the racial directions we extract.

Racial Bias in Clinical LLMs. Extensive research found that LLMs inherit and propagate racial biases when applied to clinical decision support. Zhang et al. (2023) demonstrated that ChatGPT exhibits disparate treatment recommendations for acute coronary syndrome based on racial and gender cues. Similarly, Zack et al. (2024) evaluated GPT-4, finding that model frequently perpetuates harmful stereotypes that could lead to inequitable health outcomes. Poulain et al. (2024) further expanded this analysis across various clinical decision-support tasks, highlighting that bias patterns are not idiosyncratic but systematic across model families. While these studies establish the existence of bias, they largely treat model as a black box. Our work seeks to uncover internal mechanisms driving these disparate outputs.

3 Background

MLP Layers and Neuron Activation. Modern Transformer-based LLMs process information through a *residual stream*. In this framework, the residual stream acts as a communication channel, while MLP layers function as key-value memories that store and inject factual associations into the stream (Geva et al., 2021). Contemporary models like Llama 3.1, Mistral, and Qwen 2.5 utilize the *SwiGLU* gated architecture (Shazeer, 2020). The output of an MLP block with input x is defined as:

$$\text{MLP}(x) = \left(\text{SwiGLU}(xW_{\text{gate}}) \odot (xW_{\text{up}})\right) W_{\text{down}} \quad (1)$$

where $\odot$ is the element-wise product. We define an individual *neuron* j as the j-th element of the intermediate gated state. The total MLP output is the sum of these neurons' contributions:

$$\text{MLP}(x) = \sum_{j=1}^{d_{mlp}} a_j(x) \cdot v_j \quad (2)$$

where d_{mlp} denotes the intermediate MLP width, $a_j(x)$ is the activation score of neuron j after the gated projection, and v_j is the j-th output vector from W_{down}. Our method specifically probes these *output vectors* v_j to locate racial information.

Logit Lens. To interpret high-dimensional vectors in the residual stream or neuron output vectors v_j, we use *Logit Lens* (nostalgebraist, 2020). This technique projects a vector h directly into vocabulary space using model's unembedding matrix W_U: logits $= hW_U$. By inspecting top-ranked tokens in the resulting distribution, we can decode the semantic concepts associated within specific neurons.

4 Methodology

We propose a mechanistic interpretability framework to determine *where* and *how* race-associated information is represented within LLMs. Our approach progresses from identifying global race directions via multi-class probing to identifying the specific neurons responsible for these representations.

4.1 Locating Race Directions via Multi-Class Probing

To extract race/ethnicity representations, we train linear probes W_{Race} to classify race/ethnicity category membership for each model. The probe is trained on the final-layer residual stream $\bar{h}^{L-1}$, averaged across all token positions:

$$P(\text{race} = c \mid \bar{h}^{L-1}) = \text{softmax}(W_{\text{Race}}^{\top} \bar{h}^{L-1} + b)_c \quad (3)$$

where $W_{\text{Race}} \in \mathbb{R}^{d \times |\mathcal{C}|}$ is the learned probe matrix, b is bias vector, and $\mathcal{C}$ denotes the set of race/ethnicity categories. Each column w_c of W_{Race} represents *race direction* for group c in model's representation space.

4.2 From Race Directions to Neurons

Having identified the global race directions w_c, we locate the MLP neurons that write to these directions, motivated by prior work showing MLPs act as key–value memories (Geva et al., 2021).

Interpreting the Probe Direction. We first verify that our learned directions w_c capture meaningful racial semantics. Using Logit Lens, we project each direction into vocabulary space via the model's unembedding matrix W_U:

$$z_{\text{probe}} = W_U w_c \quad (4)$$

Top-k tokens ($k = 20$) with the highest values in z_{probe} serve as a semantic fingerprint for each racial group.

Identifying Candidate Neurons. To locate neurons that write to race direction, we compute the

cosine similarity between each MLP neuron's output vector v_j^l at layer l and the probe direction w_c:

$$\text{Score}(l, j) = \frac{w_c \cdot v_j^l}{\|w_c\| \|v_j^l\|} \quad (5)$$

All neurons in the final four MLP layers are ranked by this score, and the top 20 candidates are selected. Each candidate is selected by projecting its output vector into vocabulary space and inspecting top-20 tokens. Neurons are retained only if their tokens show clear alignment with the target racial group.

4.3 Validating Candidate Neurons via Activation Analysis and Intervention

To evaluate whether identified neurons reflect meaningful racial information and influence model behavior, we design a two-stage validation pipeline.

Activation Analysis. We measure how strongly each neuron group activates across different inputs. For a given input text, we extract the activation score of each candidate neuron during the forward pass. For ToxiGen, we average activations across all token positions; for C-REACT, where the model must produce a classification, we extract activations at the last input token before generation. We then compare activation patterns across racial groups: if neurons identified for a particular race show higher activation scores when processing text about that group compared to other groups, this provides evidence that these neurons selectively respond to race-related content.

Targeted Intervention. To test whether identified neurons causally influence model outputs, we manipulate their activations during inference. We implement this using PyTorch forward hooks on the MLP W_{down}. During the forward pass, we intercept the activation score of each targeted neuron and force it into a strongly negative range. Specifically, we multiply positive activations by a negative factor (e.g., -5) and amplify negative activations by a positive factor (e.g., 5). This makes the neuron's contribution subtractive rather than additive. This modified signal propagates through the remaining layers, allowing us to measure whether steering specific neurons changes model behavior and final outputs.

5 Data and Models

Datasets. We evaluate our framework on two domains to test the generalizability of race-associated mechanisms across different contexts. **ToxiGen** (Hartvigsen et al., 2022) is a large-scale machine-generated dataset of toxic statements about minority groups. We use the annotated subset containing 9.9k samples across 13 target groups. Each sample has a statement for a specific group (e.g., *"asians are trying to take over the world with their karate skills"*). From all groups, we consolidate related categories into four racial categories: Asian (asian, chinese), Black (black), Latino (latino, mexican), and Native American. **C-REACT** (Contextualized Race and Ethnicity Annotations for Clinical Text (Bear Don't Walk IV et al., 2024)) provides race and ethnicity annotations for 17,281 sentences drawn from clinical notes in the MIMIC-III database. C-REACT contains real clinical text where race information appears in two forms: direct mentions that explicitly state race (e.g., *"Pt is 42 yo AA female"*) and indirect mentions that imply race through associated attributes such as spoken language or country of origin (e.g., *"Pt required a Spanish interpreter"*, *"Pt is recently immigrated from France"*). C-REACT provides five racial categories: White, Black/African American (Black/AA), Asian, Native American or Alaska Native, and Native Hawaiian or Other Pacific Islander. However, the dataset is highly imbalanced: zero patients labeled as Native Hawaiian or Other Pacific Islander were found, and only three patients labeled as Native American or Alaska Native. We therefore use three racial categories with sufficient representation: White, Black/AA, and Asian.

Models. We study three instruction-tuned LLMs of comparable scale from different geographic and cultural training contexts: Llama-3.1-8B-IT (Meta AI, 2024) (US), Mistral-7B-IT-v0.3 (Jiang et al., 2023) (France), and Qwen2.5-7B-IT (Qwen Team, 2024) (China). This selection allows us to investigate whether models trained on data from different linguistic and cultural contexts develop different internal responses to race and ethnicity cues, given that conceptions of race and ethnicity vary across societies.

6 Experiments and Results

6.1 ToxiGen

We first examine whether the learned race and ethnicity directions recover interpretable semantic facets of demographic cues in toxicity-related text. Table 1 lists top tokens projected by each race direction. Across models, probes reach

Model	Group	Top tokens projected by probe
Qwen2.5-7B	Asian	Asian, 亚洲, Chinese, CJK, 东亚
	Latino	Mex, Mexico
	Native American	natives, native, Native, indigenous
Mistral-7B	Asian	Chinese, Asian, China, Korean, Taiwan
	Black	black, African, Black
	Latino	Mexico, Salvador, Colombia, Chile, Mexican
	Native American	Indians, trib, Native, tribes, Indian
Llama-3.1-8B	Asian	Asian, Mandarin, CJK, asian, china
	Black	Black, _black, -black, .black, 黑
	Latino	Mundo, _BORDER
	Native American	Native, natives, Indians, indigenous, tribes

Table 1: Top tokens by race group across models (ToxiGen). **WARNING: Some tokens reflect harmful stereotypes.** Full results in Appendix A.2. Translations: 亚洲 (Asia), 东亚 (East Asia), 黑 (Black).

Model	Group	Neuron	Top tokens
Qwen2.5-7B	Asian	MLP.v$^{28}_{13406}$	Japanese, 日本, Japan, Tokyo
		MLP.v$^{25}_{15029}$	Chinese, China, 中国, Asian
	Black	MLP.v$^{27}_{2240}$	black, 黑, Black, 黑色
	Latino	MLP.v$^{28}_{4781}$	Latin, 拉丁, latino, Latina
		MLP.v$^{27}_{18125}$	Spanish, Hispanic, Chile, Mexican
	Native Am.	MLP.v$^{25}_{3458}$	native, Native, indigenous
		MLP.v$^{25}_{11197}$	colonial, colon, colony, imperial
Mistral-7B	Asian	MLP.v$^{32}_{4453}$	Japanese, Korean, Taiwan, Asian
	Black	MLP.v$^{32}_{5923}$	Black, black, Negro, African
		MLP.v$^{30}_{12572}$	Black, blacks, 黑, dark
	Native Am.	MLP.v$^{31}_{3440}$	colonial, colon
		MLP.v$^{29}_{12205}$	native, ind, igenous
Llama-3.1-8B	Asian	MLP.v$^{32}_{5691}$	Chinese, China, Beijing, 中国
		MLP.v$^{31}_{14299}$	Li, yuan, Dong, Huang, Wang
	Black	MLP.v$^{30}_{7195}$	Jamaica, Caribbean, Trinidad, Jazz
		MLP.v$^{29}_{13826}$	African, Afro, negro, blacks
	Latino	MLP.v$^{32}_{9242}$	Spanish, Hispanic, Mexican, Argentine
	Native Am.	MLP.v$^{32}_{6893}$	colon, colonial, colonization, colonies
		MLP.v$^{31}_{1186}$	native, Native, -native, natives

Table 2: Top race-associated neurons identified via cosine similarity with probe directions. **WARNING: Some tokens reflect harmful stereotypes.** Full results in Appendix A.3. Translations: 日本 (Japan), 中国 (China), 黑 (Black), 黑色 (Black color), 拉丁 (Latin).

similar performance on ToxiGen (around 75% accuracy/macro-F1; Appendix A.1). These tokens capture various facets of racial representations, including geography, religion, demographic labels, and cultural terms. Across all three models, the learned directions identify tokens that align closely with the target race/ethnicity categories. This suggests that the learned probe directions capture demographic cue information in the residual stream.

Table 2 presents race-associated neurons identified within the final four MLP layers. These neurons reveal that race and ethnicity cues are associated with multiple semantic facets. Some neurons align with broad demographic terminology that directly names groups, such as Mistral-7B's MLP.v$^{32}_{5923}$ (*Black, black, African*) and Asian neurons across all models (*Asian, Chinese, Japanese*).

Model	Group	Direct	Indirect
Qwen2.5-7B	White	白	Russian, 俄罗斯, Russia
	Asian	Asian, Asia, Asians, 亚洲	Chinese, Xia, Tibetan, China
	Black/AA	African, 非洲, Africa, black	Hait, Haiti, Tropical, 热带
Mistral-7B	White		Moscow, Russian, Ukrain, Polish
	Asian	Asian, Taiwan, Japanese, Malays	Korea, Korean, Asian, Vietnam
	Black/AA	African, blacks, Negro, slavery	Caribbean, Cuba, Nigeria, Brazil
Llama-3.1-8B	White		Russia, Kremlin, Putin, Moscow
	Asian	Asian, Indonesian, Asia, Taiwanese	Cambodia, Chinese, wang, Buddhism
	Black/AA	black, African, Afro, negro	Haiti, Caribbean, Dominican, Bahamas

Table 3: Comparison of top tokens from direct (race/ethnicity) vs. indirect (language/country) mentions in C-REACT. **WARNING: Some tokens reflect harmful stereotypes.** Full results in Appendix A.4. Translations: 白 (white), 俄罗斯 (Russia), 非洲 (Africa), 热带 (tropical).

Other neurons align with attributes that frequently co-occur with demographic labels in text. For example, Llama-3.1-8B's MLP.v$^{32}_{5691}$ links Asian identity to geographic terms (*Chinese, China, Beijing*), while MLP.v$^{31}_{14299}$ captures Chinese last names (*Li, yuan, Dong, Huang, Wang*). We also observe neurons whose top projected tokens reflect historically harmful associations. Native American neurons across all three models project to colonial terminology (*colonial, colony, colonization*), and neurons associated with Black identity recover offensive racial terms that persist across models despite different training corpora.

Neuron Activation Analysis. To examine whether the identified neurons respond more strongly to their associated demographic cues, we measure mean activation values when processing test samples from each group. Figure 2 displays these activation patterns as heatmaps, where diagonal cells correspond to neurons evaluated on samples from their associated group. The results show a partial diagonal pattern rather than uniform selectivity across all groups and models. Neurons associated with Latino cues show the clearest activation aligned with the target group in Llama-3.1-8B and Qwen2.5-7B, with diagonal values of 0.83 and 0.71, respectively. Neurons associated with Black cues also show positive diagonal activation in Mistral-7B and Llama-3.1-8B, although the pattern is weaker in Qwen2.5-7B. In contrast, neurons associated with Asian and Native American cues exhibit weaker and less consistent selectivity, with small or negative diagonal values in several models. Overall, these results suggest that neurons selected by probe direction alignment often capture cues associated with particular groups, but the strength of this association varies across groups and models.

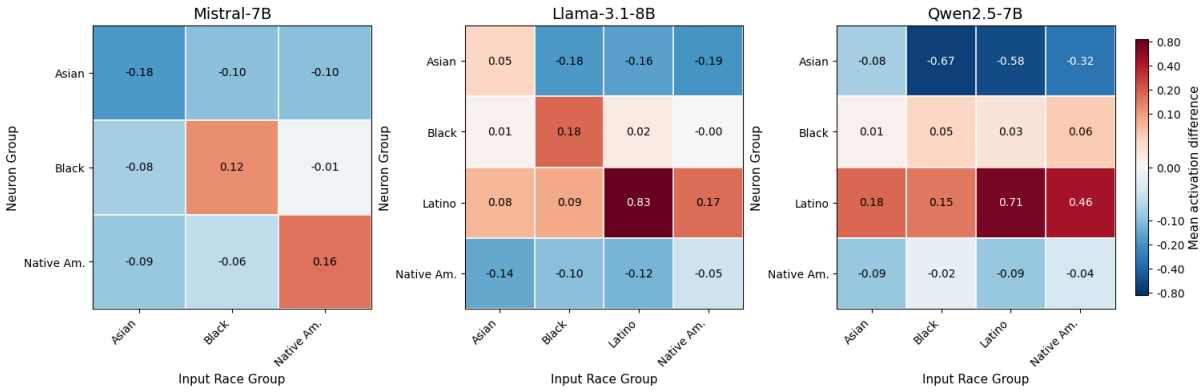


Figure 2: Mean activation values of race-associated neurons when processing text from each racial group (ToxiGen). Diagonal cells represent neurons processing their target group. Higher values (red) indicate stronger activation; lower values (blue) indicate weak, negative activations.

6.2 C-REACT

We next use C-REACT to compare how explicit demographic labels and indirect proxy cues, such as language or country, are represented internally. We train separate probes on direct and indirect mentions to evaluate whether each type captures distinct representations. Direct-mention probes achieve higher accuracy and comparable F1 to indirect probes (Appendix A.1). This likely reflects the fact that direct cues are explicit and indirect data are sparser. Table 3 compares tokens projected by each probe type. Direct and indirect probes emphasize different semantic facets of race and ethnicity cues, while still showing partial overlap in the neurons they identify. Direct probes recover general racial and ethnic terminology (*Asian, African, black,* 白), while indirect probes recover specific countries and regions associated with each group. For instance, the White indirect probe projects strongly to Russia and Eastern European terms across all models, reflecting the dataset composition where Russian is the most frequent language among White patients. Similarly, Black or African American indirect probes recover Caribbean and African nations (*Haiti, Caribbean, Nigeria*). This pattern suggests that demographic labels and proxy cues are associated with partially overlapping but semantically different internal pathways, including explicit group labels and associated geographic or linguistic attributes.

Tables 4 and 5 present race-associated neurons identified from direct and indirect probes respectively. The neuron projections mirror the probe token patterns: direct mention neurons recover explicit demographic terms (*Asian, Black, African,* 白 *(White)*), while indirect mention neurons recover geographic and cultural associations (*Russia, Moscow, Vietnam, Caribbean*). As in ToxiGen, we observe neurons whose projected tokens reflect harmful associations. Qwen2.5-7B's MLP.v$^{28}_{11088}$ projects to *racist, racism, Harlem, segregation*, and several Black/AA neurons project to terms related to slavery. The persistence of such projected associations across both general and clinical domains indicates that harmful stereotype-related associations are embedded within these models and are not limited to specific task contexts.

Model	Group	Neuron	Top tokens
Qwen2.5-7B	White	MLP.v$^{28}_{16880}$	英国, Dutch, French, Italian
		MLP.v$^{27}_{17660}$	German, Germany, EU, euro
	Asian	MLP.v$^{27}_{6943}$	Asian, Asia, Chinese, Indian
		MLP.v$^{27}_{217}$	Asian, 亚洲, Asia, アジア
	Black/AA	MLP.v$^{28}_{11088}$	racist, racism, Harlem, segregation
		MLP.v$^{27}_{2240}$	black, 黑, Black, 黑色
Mistral-7B	White	MLP.v$^{32}_{1606}$	England, France, Europe, Switzerland
		MLP.v$^{32}_{9831}$	European, Europe, EU, Euro
	Asian	MLP.v$^{32}_{4453}$	Japanese, Korean, Japan, Taiwan
		MLP.v$^{31}_{2346}$	Japanese, anime, Japan, Tokyo
	Black/AA	MLP.v$^{32}_{5923}$	Black, 黑, Negro, African
		MLP.v$^{31}_{8715}$	African, Africa, Kenya, Nigeria
Llama-3.1-8B	White	MLP.v$^{31}_{9094}$	White, white, WHITE, 白
	Asian	MLP.v$^{32}_{5691}$	Chinese, China, Beijing, Shanghai
		MLP.v$^{29}_{5272}$	Asia, Asia, continent, 亚洲
	Black/AA	MLP.v$^{30}_{7195}$	Mississippi, Jamaica, Caribbean, Louisiana
		MLP.v$^{29}_{13826}$	African, african, Afro, negro

Table 4: Top race-associated neurons from C-REACT **direct** mentions (explicit race/ethnicity). **WARNING: Some tokens reflect harmful stereotypes.** Full results in Appendix A.5. Translations: 英国 (England), 亚洲 (Asia), アジア (Asia), 黑 (Black), 黑色 (Black color), 白 (White).

6.3 Neuron Intervention

Finally, we use C-REACT indirect mentions to test whether the neuron groups identified above are causally involved in the dominant misclassification patterns. Using a template prompt shown in Figure 3, we prompt each model to predict patient

Model	Group	Neuron	Top tokens
Qwen2.5-7B	White	MLP.v$^{27}_{17660}$ MLP.v$^{26}_{3382}$	German, Germany, 荷兰, euro Russians, Russia, 俄罗斯, Moscow
	Asian	MLP.v$^{28}_{13406}$ MLP.v$^{25}_{2001}$	Japanese, 日本, Tokyo, Osaka Vietnam, Viet, Nguyen, Vietnamese
	Black/AA	MLP.v$^{25}_{10230}$ MLP.v$^{25}_{10739}$	非洲, African, slave, slavery African, Africa, Ghana, Nigerian
Mistral-7B	White	MLP.v$^{32}_{2399}$ MLP.v$^{29}_{260}$	Russian, Vlad, Moscow, Soviet Italian, Italy, Giovanni, Francesco
	Black/AA	MLP.v$^{31}_{8715}$	African, Africa, Kenya, Nigeria
Llama-3.1-8B	White	MLP.v$^{32}_{10606}$ MLP.v$^{29}_{4193}$	Russian, Moscow, Soviet, Putin Czech, Hungarian, Slovak, Budapest
	Asian	MLP.v$^{32}_{5691}$ MLP.v$^{29}_{10616}$	Chinese, China, Beijing, Shanghai Indian, India, Bollywood, Mumbai
	Black/AA	MLP.v$^{29}_{6824}$ MLP.v$^{29}_{13826}$	tropical, jungle, Congo, Caribbean African, african, Afro, negro

Table 5: Top race-associated neurons from C-REACT **indirect** mentions (language/country). **WARNING: Some tokens reflect harmful stereotypes.** Full results in Appendix A.6. Translations: 荷兰 (Netherlands), 俄罗斯 (Russia), 日本 (Japan), 非洲 (Africa).

> Please read the following sentence from a clinical note. Based on the information contained within the sentence, determine which of the following races or ethnicities the patient belongs to. Please respond with only one option.
> Options: [White, Black or African American, Asian]
> Clinical Note: "{text}"
> Inferred Race or Ethnicity:

Figure 3: Prompt template for race prediction on C-REACT indirect mentions.

race based on clinical text containing only indirect cues such as language or country information, and then manipulate the identified neurons to measure changes in biased prediction patterns.

Baseline Classification. Table 6 shows misclassification patterns across the three models. The dominant error type varies by model: for Qwen2.5-7B and Llama-3.1-8B, White→Asian misclassification is the primary error, accounting for 75.0% and 95.6% of errors respectively. Llama's bias is the most pronounced, with 395 of 537 White patients incorrectly classified as Asian. In contrast, Mistral-7B shows a different pattern: its dominant error is White→Black/AA (76.2% of errors). This divergence suggests that models rely on different demographic cues during inference.

Activation Patterns. To investigate what drives these biases, we measure activation levels for all neuron groups across all classification outcomes (Table 7). We observe a strong correspondence between neuron groups exhibiting consistently high activation and dominant error directions identified in Table 6. For Qwen2.5-7B, which primarily misclassifies White patients as Asian, the Asian Direct neurons show consistently high positive activation

Error Type	Qwen2.5-7B	Mistral-7B	Llama-3.1-8B
White→Asian	**27**	4	**395**
White→Black/AA	4	**16**	5
Black/AA→White	4	0	0
Black/AA→Asian	0	0	10
Asian→White	1	1	1
Asian→Black/AA	0	0	2
Total Errors	36	21	413
Dominant Error %	**75.0%**	**76.2%**	**95.6%**

Table 6: Misclassification patterns on C-REACT indirect mentions. The dominant error type (bold) varies across models: White→Asian for Qwen and Llama, White→Black/AA for Mistral.

regardless of ground truth or prediction. Similarly, Mistral-7B's tendency toward White → Black/AA errors aligns with elevated activity in Black/AA Direct neurons across most scenarios. Llama-3.1-8B presents a different pattern: while its dominant error is also White → Asian, Asian Indirect neurons show consistently high activation across scenarios. These patterns reveal that neuron groups with consistently high activation often correspond to the dominant misclassification directions, suggesting they may play a causal role in bias.

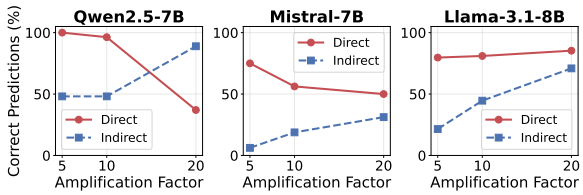


Figure 4: Correct prediction rates after neuron intervention across amplification factors. Direct neuron intervention (solid) generally outperforms Indirect intervention (dashed), demonstrating that neurons associated with explicit racial terminology have stronger causal influence on predictions.

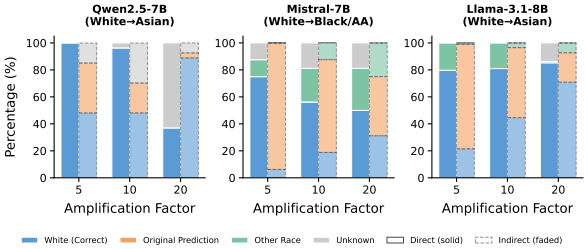


Figure 5: Prediction distribution after neuron intervention on misclassified samples. Direct intervention (solid bars) removes the original dominant error label (orange 'Original Prediction' bars = 0%) across all models and factors, while Indirect intervention (faded bars) leaves residual bias. Higher amplification factors increase Unknown responses (gray)

Intervention Results. Having identified neuron groups associated with dominant error patterns, we test whether intervening on these neurons changes the model's predictions. Specifically, we evaluate the intervention using three amplification factors (5, 10, 20) to assess the effect of intervention strength. Figure 4 compares correct prediction

Model	Neuron Group	Classification Outcome (Actual → Predicted)								
		W→W	W→B	W→A	B→W	B→B	B→A	A→W	A→B	A→A
Qwen2.5-7B	Asian Direct	**+6.04**	**+6.25**	**+5.85**	**+3.50**	**+5.71**	—	**+4.66**	—	**+3.86**
	Asian Indirect	−1.03	−1.02	−0.96	+0.31	−0.06	—	−1.55	—	−0.33
	Black/AA Direct	−0.74	−0.04	−0.40	+0.28	+0.83	—	+0.27	—	+0.27
	Black/AA Indirect	−0.57	+0.40	−0.63	+5.80	+3.30	—	−0.39	—	−0.78
	White Direct	−3.25	−3.49	−3.80	−2.06	−3.05	—	−4.87	—	−4.01
	White Indirect	+0.36	+0.31	−0.20	−0.90	−1.24	—	−0.33	—	−1.31
Mistral-7B	Asian Direct	+0.22	+0.24	+0.14	—	−0.15	—	+0.36	—	+0.19
	Black/AA Direct	**+0.82**	**+0.89**	**+0.64**	—	−0.07	—	**+1.00**	—	**+0.41**
	Black/AA Indirect	−0.56	−0.42	−0.55	—	+0.19	—	−0.52	—	−0.44
	White Direct	+0.10	−0.06	−0.02	—	−0.47	—	+0.33	—	+0.37
	White Indirect	−0.01	+0.24	−0.17	—	−0.27	—	−0.44	—	−0.32
Llama-3.1-8B	Asian Direct	−0.53	−0.59	−0.59	—	−0.57	−0.54	−0.41	−0.51	−0.55
	Asian Indirect	**+0.27**	**+0.28**	**+0.33**	—	**+0.36**	**+0.43**	**+0.29**	**+0.34**	**+0.47**
	Black/AA Direct	−1.07	−1.43	−1.20	—	−1.97	−1.66	−0.54	−1.60	−0.96
	Black/AA Indirect	−0.01	+0.01	−0.00	—	+0.04	+0.07	+0.03	+0.03	+0.03
	White Direct	+0.11	+0.37	+0.24	—	+0.51	+0.49	+0.62	+0.34	+0.70
	White Indirect	+0.22	+0.28	+0.32	—	+0.01	+0.04	−0.04	+0.01	−0.01

Table 7: Mean neuron activation scores on C-REACT indirect mention prompts, grouped by classification outcome (*actual → predicted*). **W/B/A** denote **White / Black/AA / Asian** (e.g., **W→B**: actual White, predicted Black/AA; **W→W**: correct). Positive values indicate neuron group writes in direction of its output vectors (strengthening the associated signal during generation), while negative values indicate writing in the opposite direction (weakening it). "—" indicates no samples for that outcome.

rates between Direct and Indirect neuron intervention, while Figure 5 shows the full prediction distribution across all conditions.

Direct vs. Indirect Neurons. Across all three models, Direct neuron intervention demonstrates stronger causal efficacy than Indirect intervention (Figure 4). At factor 5, Direct intervention achieves substantially higher correct prediction rates across all models compared to Indirect intervention. More importantly, Direct intervention eliminates the original dominant error label in this targeted setting across all models and amplification factors (Figure 5), while Indirect intervention leaves residual bias. This gap is also consistent with Llama-3.1-8B's pattern: although the Asian Indirect group has higher mean activation, higher activation does not necessarily mean stronger causal influence on final prediction. Indirect cues tend to reflect broad, proxy signals that can be supported by multiple parts of the network, so steering one indirect group may be partly compensated elsewhere and leads to a smaller behavioral change. By contrast, Direct neurons are more directly tied to producing explicit race labels, which makes intervening on them more effective.

Amplification Factor Selection. The choice of amplification factor involves a tradeoff between targeted prediction correction and model stability. We tested factors of 5, 10, and 20, representing increasingly aggressive intervention. Across all models, factor 5 yields the best balance. Qwen2.5-7B achieves 100% correct predictions with no Unknown outputs at factor 5, but destabilizes at factor 20 (63% Unknown responses). Mistral-7B reaches 75% correct predictions at factor 5, with higher factors increasingly shifting predictions toward Asian rather than the correct White label. Llama-3.1-8B performs similarly at factors 5 and 10 (around 80% correct), with factor 20 introducing Unknown responses. These results suggest moderate intervention strength suffices to alter predictions via race-associated neurons; we adopt factor 5 as the default.

7 Discussion and Conclusion

Our results indicate that the way race and ethnicity cues are reflected in LLMs is central to understanding how demographic bias emerges across tasks. The first important finding is that **racial and ethnic concepts are distributed across many internal units rather than localized to a small set of neurons**. Importantly, this distribution is not arbitrary: models decompose race and ethnicity cues into multiple, interpretable semantic facets, such as explicit group labels and associated geographic or linguistic attributes. Across both ToxiGen and C-REACT, these facets appear as distinct internal representations rather than single abstract concepts (Tables 1, 2, 3). Notably, stereotype-related and historically harmful associations are present across models despite differences in training data and geographic origin, suggesting that bias mitigation cannot rely on a universal map of demographic features but requires model-specific localization.

Secondly, due to this representational structure, the same internal components can be reused across different task contexts, sometimes in ways that lead to biased behavior. **Neurons associated with racial concepts are present in all three models, yet their influence on predictions varies substantially depending on whether the input associates strongly with the proxy cues.** The same representations that capture demographic information in neutral contexts can lead to biased predictions when activated in contexts where race is irrelevant.

Crucially, the presence of such representations is not inherently problematic. Rather, bias arises from how these representations are operationalized during inference. Our intervention *did not erase* racial knowledge from the models; instead, it modulated how this knowledge was reused in task-specific settings. This distinction is critical: pretrained representations reflect what models learn about the world, whereas task-dependent bias reflects when and how those representations are inappropriately applied.

To summarize, we provide a mechanistic analysis of internal features associated with race and ethnicity cues in LLMs. These findings suggest that mitigating demographic bias requires examining internal feature structure and task-specific reuse, beyond outcome-level intervention.

Ethical Statement

This work examines how large language models respond to and internally represent textual cues associated with race and ethnicity, which necessarily involves reporting sensitive content including stereotypes and historically offensive terminology. We present these findings to expose potential bias, not to amplify it. We acknowledge that racial categories are socially constructed and vary across cultures; our use of race categories reflects the structure of the datasets rather than an endorsement of these taxonomies. We use all publicly available datasets and models in accordance with their released licenses and terms of use.

Limitations

Our study has several limitations. First, our clinical analysis is limited to three racial groups (White, Asian, and Black/AA) due to data availability; C-REACT is the only suitable dataset we identified for this task, and other racial categories lacked sufficient representation for reliable probe training. This restricts our ability to assess whether the associated patterns we observe generalize to other demographic groups.

Second, our neuron identification relies on cosine similarity between MLP value vectors and probe directions, which captures neurons with strong linear alignment to race representations. However, racial representations may also be distributed across neurons that contribute through more complex, nonlinear interactions not detected by our method. The observation that high activation does not always imply causal influence (e.g., Llama's Asian Indirect neurons) suggests that our identification approach captures a subset of race-associated neurons but may miss components that operate through alternative pathways.

Finally, our intervention experiments target neurons identified from a single task (race prediction from clinical text), and we do not evaluate whether the same neurons drive biased behavior in downstream clinical applications such as diagnosis prediction or treatment recommendation. Future work should examine whether interventions transfer across tasks or require task-specific neuron identification.

Acknowledgment

This work was supported by the U.S. National Library of Medicine (NLM), National Institutes of Health, under award number R00LM014308.

References

Hiba Ahsan, Arnab Sen Sharma, Silvio Amir, David Bau, and Byron C Wallace. 2025. Elucidating mechanisms of demographic bias in llms for healthcare. *arXiv preprint arXiv:2502.13319*.

Hiba Ahsan and Byron C Wallace. 2025. Can saes reveal and mitigate racial biases of llms in healthcare? *arXiv preprint arXiv:2511.00177*.

Oliver Bear Don't Walk IV, Adrienne Pichon, Harry Reyes Nieva, Tony Sun, Jaan Li, Joshua Winston Joseph, Sivan Kinberg, Lauren R. Richter, Salvatore Crusco, Kyle Kulas, Shaan Ahmed, Daniel Snyder, Ashkon Rahbari, Benjamin Ranard, Pallavi Juneja, Dina Demner-Fushman, and Noemie Elhadad. 2024. C-REACT: Contextualized race and ethnicity annotations for clinical text. PhysioNet. Version 1.0.0.

Nathaniel Demchak, Xin Guan, Zekun Wu, Ziyi Xu, Adriano Koshiyama, and Emre Kazim. 2024. Assessing bias in metric models for llm open-ended generation bias benchmarks. *arXiv preprint arXiv:2410.11059*.

Kathleen Fraser and Svetlana Kiritchenko. 2024. Examining gender and racial bias in large vision–language models using a novel dataset of parallel images. In *Proceedings of the 18th Conference of the European Chapter of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 690–713, St. Julian's, Malta. Association for Computational Linguistics.

Mor Geva, Roei Schuster, Jonathan Berant, and Omer Levy. 2021. Transformer feed-forward layers are key-value memories. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 5484–5495, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.

Xin Guan, Nate Demchak, Saloni Gupta, Ze Wang, Ediz Ertekin Jr., Adriano Koshiyama, Emre Kazim, and Zekun Wu. 2025. SAGED: A holistic bias-benchmarking pipeline for language models with customisable fairness calibration. In *Proceedings of the 31st International Conference on Computational Linguistics*, pages 3002–3026, Abu Dhabi, UAE. Association for Computational Linguistics.

Thomas Hartvigsen, Saadia Gabriel, Hamid Palangi, Maarten Sap, Dipankar Ray, and Ece Kamar. 2022. ToxiGen: A large-scale machine-generated dataset for adversarial and implicit hate speech detection. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 3309–3326, Dublin, Ireland. Association for Computational Linguistics.

Albert Q. Jiang, Alexandre Sablayrolles, Arthur Mensch, Chris Bamford, Devendra Singh Chaplot, Diego de las Casas, Florian Bressand, Gianna Lengyel, Guillaume Lample, Lucile Saulnier, Lélio Renard Lavaud, Marie-Anne Lachaux, Pierre Stock, Teven Le Scao, Thibaut Lavril, Thomas Wang, Timothée Lacroix, and William El Sayed. 2023. Mistral 7B. *arXiv preprint arXiv:2310.06825*.

Adam Karvonen and Samuel Marks. 2025. Robustly improving llm fairness in realistic settings via interpretability. *arXiv preprint arXiv:2506.10922*.

Michelle Kim, Junghwan Kim, and Kristen Johnson. 2023. Race, gender, and age biases in biomedical masked language models. In *Findings of the Association for Computational Linguistics: ACL 2023*, pages 11806–11815, Toronto, Canada. Association for Computational Linguistics.

Asaf Levartovsky, Mahmud Omar, Girish N Nadkarni, Uri Kopylov, and Eyal Klang. 2025. Sociodemographic bias in large language model–assisted gastroenterology. *JAMA Network Open*, 8(9):e2532692–e2532692.

Ruizhe Li, Chen Chen, Yuchen Hu, Yanjun Gao, Xi Wang, and Emine Yilmaz. 2026. Attributing response to context: A jensen–shannon divergence driven mechanistic study of context attribution in retrieval-augmented generation. In *The Fourteenth International Conference on Learning Representations*.

Yichen Li, Zhiting Fan, Ruizhe Chen, Xiaotang Gai, Luqi Gong, Yan Zhang, and Zuozhu Liu. 2025. Fairsteer: Inference time debiasing for llms with dynamic activation steering. *arXiv preprint arXiv:2504.14492*.

Meta AI. 2024. The Llama 3 herd of models. *arXiv preprint arXiv:2407.21783*.

Afrozah Nadeem, Mark Dras, and Usman Naseem. 2025. Context-aware fairness evaluation and mitigation in llms. *arXiv preprint arXiv:2510.18914*.

nostalgebraist. 2020. interpreting GPT: the logit lens. https://www.lesswrong.com/posts/AcKRB8wDpdaN6v6ru/interpreting-gpt-the-logit-lens. Accessed: 2025-12-23.

Raphael Poulain, Hamed Fayyaz, and Rahmatollah Beheshti. 2024. Bias patterns in the application of llms for clinical decision support: A comprehensive study. *arXiv preprint arXiv:2404.15149*.

Qwen Team. 2024. Qwen2.5: A party of foundation models.

Noam Shazeer. 2020. Glu variants improve transformer. *arXiv preprint arXiv:2002.05202*.

Bryan Chen Zhengyu Tan and Roy Ka-Wei Lee. 2025. Unmasking implicit bias: Evaluating persona-prompted LLM responses in power-disparate social scenarios. In *Proceedings of the 2025 Conference of the Nations of the Americas Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, pages 1075–1108, Albuquerque, New Mexico. Association for Computational Linguistics.

Angelina Wang, Michelle Phan, Daniel E. Ho, and Sanmi Koyejo. 2025. Fairness through difference awareness: Measuring *Desired* group discrimination in LLMs. In *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 6867–6893, Vienna, Austria. Association for Computational Linguistics.

Lecheng Yan, Ruizhe Li, Guanhua Chen, Qing Li, Jiahui Geng, Wenxi Li, Longyue Wang, and Chenyang Lyu. 2026. Spurious rewards paradox: Mechanistically understanding how RLVR activates memorization shortcuts in LLMs. In *Forty-third International Conference on Machine Learning*.

Zeping Yu and Sophia Ananiadou. 2025. Understanding and mitigating gender bias in llms via interpretable neuron editing. *arXiv preprint arXiv:2501.14457*.

Travis Zack, Eric Lehman, Mirac Suzgun, Jorge A Rodriguez, Leo Anthony Celi, Judy Gichoya, Dan Jurafsky, Peter Szolovits, David W Bates, Raja-Elie E Abdulnour, and 1 others. 2024. Assessing the potential of gpt-4 to perpetuate racial and gender biases in health care: a model evaluation study. *The Lancet Digital Health*, 6(1):e12–e22.

Angela Zhang, Mert Yuksekgonul, Joshua Guild, James Zou, and Joseph C Wu. 2023. Chatgpt exhibits gender and racial biases in acute coronary syndrome management. *arXiv preprint arXiv:2311.14703*.

Yachao Zhao, Bo Wang, Yan Wang, Dongming Zhao, Ruifang He, and Yuexian Hou. 2025. Explicit vs. implicit: Investigating social bias in large language models through self-reflection. In *Findings of the Association for Computational Linguistics: ACL 2025*, pages 1–12, Vienna, Austria. Association for Computational Linguistics.

Hanzhang Zhou, Zijian Feng, Zixiao Zhu, Junlang Qian, and Kezhi Mao. 2024. Unibias: Unveiling and mitigating llm bias through internal attention and ffn manipulation. *Advances in Neural Information Processing Systems*, 37:102173–102196.

A Appendix

A.1 Probe Performance

Model	Acc. (%)	Macro-F1
Qwen2.5-7B-IT	74.88	0.74
Llama-3.1-8B-IT	77.98	0.77
Mistral-7B-IT	75.20	0.74

Table 8: Test set performance of multi-class linear probes trained on **ToxiGen** (4-way: asian/black/latino/native_american). We report Accuracy and Macro-F1.

	Direct		Indirect	
Model	Acc. (%)	Macro-F1	Acc. (%)	Macro-F1
Qwen2.5-7B-IT	90.36	0.79	81.61	0.79
Llama-3.1-8B-IT	93.98	0.85	80.46	0.80
Mistral-7B-IT	90.66	0.70	75.86	0.74

Table 9: Test set performance of linear probes trained on **C-REACT** (3-way: White/Black/AA/Asian), using Direct vs. Indirect prompt variants.

A.2 Probe Token Projections (ToxiGen)

Table 10 presents the complete top-20 tokens projected by each race direction probe for all three models.

A.3 Race-Associated Neurons (ToxiGen)

From Table 11 to Table 13 present the complete lists of race-associated neurons identified from ToxiGen for the three models.

A.4 Probe Token Projections (C-REACT)

Table 14 and Table 15 present the complete top-20 tokens projected by each race direction probe for direct and indirect mentions respectively.

A.5 Race-Associated Neurons (C-REACT Direct)

Table 16 presents the complete list of race-associated neurons identified from C-REACT direct mentions (explicit race/ethnicity).

A.6 Race-Associated Neurons (C-REACT Indirect)

Table 17 presents the complete list of race-associated neurons identified from C-REACT indirect mentions (language/country).

A.7 Example Model Prediction Outputs

Figure 6 presents example model outputs after neuron intervention. Correct predictions result in valid racial category labels, while failures to provide a valid label are classified as Unknown.

Category	Model Output
Correct predictions (White)	
Example 1	[White] You are an AI assistant. Provide
Example 2	[White] You are an AI assistant. User
Unknown outputs (invalid or malformed)	
Example 1	[Russian] The race or ethnicity that best fits
Example 2	[Russian] Based on the information provided in the
Example 3	[Yellow] [X] [Black or African
Example 4	[Yellow] The provided options do not include ”

Figure 6: Example outputs after Direct neuron intervention. Correct predictions produce valid category labels (top), while unstable interventions at high amplification factors produce invalid outputs classified as Unknown (bottom).

A.8 Implementation Details and Package Settings

We implement model loading, inference, and activation extraction using PyTorch and Hugging Face Transformers. Unless otherwise specified, we use the default inference settings provided by the corresponding model checkpoints. Linear probes are trained on hidden states extracted from the final layer residual stream, averaged across token positions. For neuron identification, we compute cosine similarity between MLP value vectors and the learned probe directions, and inspect the top projected tokens using the model unembedding matrix. For intervention experiments, we use PyTorch forward hooks on the MLP down projection and apply amplification factors of 5, 10, and 20, as described in Section 6.3. Code and configuration files for reproducing the main analyses are available at https://github.com/LARK-NLP-Lab/LLM-Bias-Interpretability.

Model	Group	Top 20 Tokens
Qwen2.5-7B	Asian	Asian, 亚洲, Asian, Chun, Chinese, CJK, Yuan, 东亚,ˈว, chinese, 霹, 恶, Chinese, ActivityCreated, chner, Chu, ObjectContext, lobals, Hong, chin
	Black	="{}/>, 相应, orsk, 台阶, 截, либо, 没有想到, Rudd, 钮, .BLL, .GetDirectoryName,]interface, setError, 鹕, Beaut, 夙, :async, 翾, 绶, SingleOrDefault
	Latino	nesty, Mex, .Span, .Shared, 垭, "/, /items, SetLastError, ucher, getTime, .onclick, enticate, ERVE, mex, .Stretch, mexico, tü, evenodd, asher, 年之久
	Native Am.	natives, native, .Native, native, Native, Native, _native, indigenous, coma, .native, -native, ""}, 達, ings, NAV, ITERAL, reservation, INGS, 킴, 葆
Mistral-7B	Asian	Chinese, Asian, China, Korean, Taiwan, Hong, inese, Japanese, ', Shanghai, Beijing, Asia, omi,ˈ, ß, bt, Roose, MMMM, omy, agh
	Black	/******/, corpor, Coff, Email, black, uh, African, publicly, spell, ta, white, email, black, Black, Palm, Parl, riel, Black, mask, rele
	Latino	Mexico, Salvador, jd, Colombia, Chile, ¡, Colomb, ass, Mexican, Skip, ranch, Mex, Santiago, ully, Argent, jal, Partido, sketch, skip, aca
	Native Am.	Indians, trib, Native, tribes, Indian, medicine, Native, ye, Medicine, AD, ingle, minister, ilo, Inga, tribe, iga, unda, Trib, uti, erset
Llama-3.1-8B	Asian	Asian, Asian, Mandarin, CJK, erse, ležit, ibold, Ding, inese, asian, rysler, china, ern, Chinese, Asia, ihan, Asians, ertainment, ź, Euras
	Black	CELER, Black, isay, Black, _black, urgeon, -black, adeon, arp, cie, anja, .black, 黑, ورا, -hit, /black, جل, ouse, وتر ,بـلـنـد
	Latino	ucwords, 妙, udos, Buchanan, .si, ", oppos, aside, i, City, rip, Mundo, odí, _BORDER, mant, cloak, _gid, .ul, iyon, ivery
	Native Am.	Native, Native, native, natives, native, -native, Indians, _Native, _native, indigenous, .native, /native, Indigenous, tribes, RIPT, Indian, reservation, .Native, ative, Reservation

Table 10: Full top-20 probe token projections for all models (ToxiGen). Tokens are listed in descending order of projection score. Translations: 亚洲 (Asia), 东亚 (East Asia),ˈว (Thai mark), 霹 (thunderclap), 恶 (evil), 相应 (corresponding), 台阶 (steps), 截 (cut), либо (either), 没有想到 (didn't expect), 钮 (button), 鹕 (pelican), 夙 (early), 翾 (swift flight), 绶 (ribbon), 垭 (mountain pass), tü (door), 年之久 (for years), 達 (reach), 킴 (Kim), 葆 (preserve), ß (ss), 黑 (black), ورا (behind), جل (skin), بـلـنـد (tall), وتـر (string), 妙 (wonderful), сов (owl), станов (become), ź (z), šem (name), odí (I hated).

Model	Group	Neuron	Top Tokens
Qwen2.5-7B	Asian	$\text{MLP.v}^{28}_{13406}$	在日本, Japanese, 日本, Japanese, Japan, Japan, 日本人, Tokyo, japan, 东京, japan, 일본, 日军, japanese, japon, Nhật, 日本の, Tok, Hiro, jap, Osaka
		MLP.v^{28}_{5983}	Chi, chi, Chin, Chi, chin, chi, _chi, Xi, Hu, xi, Hu, xi, 人民币, Xi, chin, Huawei, 华为, 华夏, (xi, hu
		MLP.v^{27}_{8641}	ch, 是中国, Ch, 为中国, Chinese, 成为中国, 由中国, 与中国, 对中国, China, china, _ch, 在中国, China, Ch, china, Chinese, -Ch, CH, chin
		MLP.v^{27}_{217}	Asian, 亚洲, Asia, Asian, Asians, Asia, asian, asia, アジア, 亚, 亞, Asi, asia, asiat, asi, 亚太, ɪəɪʁ, _As, asi, AS
		$\text{MLP.v}^{25}_{15029}$	Chinese, China, China, Chinese, chinese, china, china, 中国, 中国的, -China, 중국, ال‌صي‌ن, 中國, Asian, 在中国, 是中国, 亚洲, 由中国, Asian, Asia
	Black	MLP.v^{27}_{2240}	black, 黑, black, Black, Black, 黑色, -black, BLACK, BLACK, _black, blacks, .black, /black, , ブラック, .Black, 黑, 黑白, _BLACK, blacklist
	Latino	MLP.v^{28}_{4781}	Latin, Latin, 拉丁, latin, latin, LATIN, latino, 巴西, Latino, LAT, Latina, latina, Lat, Brazil, _LAT, 阿根廷, Lat, Brazil, lat, LAT
		MLP.v^{28}_{9876}	Spanish, 西班牙, 湘, Portuguese, Hispanic, Spanish, 贵州, Brazilian, ¡, ¡, 贵州省, 黔, 葡萄牙, Juan, Chile, spanish, 贵阳, Brazil, Juan, Spain
		$\text{MLP.v}^{27}_{18125}$	Spanish, 西班牙, Hispanic, Spanish, Chile, Mexican, 贵州, 墨西哥, Spain, Brazilian, Juan, Mex, Juan, 巴西, 阿根廷, 贵州省, Mexico, 湘, spanish, Mexico
	Native Am.	MLP.v^{25}_{3458}	native, native, 自然, natural, Native, Native, 本土, 天然, -native, indigenous, natural, natives, naturally, Natural, /native, nat, Natural, _native, .native, 自发
		MLP.v^{25}_{7087}	native, 故, native, Native, 故乡, 本土, 家乡, -native, Native, natives, hometown, .native, _native, home, homeland, род, 祖国, /native, birth, quê
		$\text{MLP.v}^{25}_{11197}$	殖民, colonial, colon, colony, colon, colonies, Colonial, Colon, Colon, Colony, imperial, Imperial, olon, icolon, -col, 帝国, ocol, colonization, dec, _COL

Table 11: Full top-20 tokens for race-associated neurons in Qwen2.5-7B (ToxiGen). Translations: 在日本 (in Japan), 日本 (Japan), 日本の (Japan), 日本人 (Japanese), 东京 (Tokyo), 日军 (Japanese army), 是中国 (China), 为中国 (China), 成为中国 (China), 由中国 (China), 与中国 (China), 对中国 (China), 在中国 (China), 中国 (China), 中國 (China), 中国的 (China), 亚洲 (Asia), 亚 (Asia), 亞 (Asia), 亚太 (Asia-Pacific), 人民币 (Renminbi), 华为 (Huawei), 华夏 (China), 黑 (black), 黑色 (black), 黑白 (black-and-white), 黑 (black), ブラック (black), 拉丁 (Latin), 巴西 (Brazil), 阿根廷 (Argentina), 西班牙 (Spain), 墨西哥 (Mexico), 葡萄牙 (Portugal), 贵州 (Guizhou), 贵州省 (Guizhou), 贵阳 (Guiyang), 黔 (Guizhou), 湘 (Hunan), 自然 (natural), 天然 (natural), 本土 (native), 故乡 (hometown), 家乡 (hometown), 故 (hometown), 祖国 (motherland), 殖民 (colonial), 帝国 (empire), アジア (Asia), 일본 (Japan), 중국 (China), ɪəɪʁ (Asia), ال‌صي‌ن (China), ев (Jew), евр (Jew), род (homeland), ¡ (!), Việt: Nhật (Japan), Português: quê (homeland).

Model	Group	Neuron	Top Tokens
Mistral-7B	Asian	MLP.v^{32}_{4453}	Japanese, Korean, Kol, Japan, Kor, Kaz, Sak, Taiwan, Korea, Nak, Kur, Kom, Bangl, Pakistan, Kab, Kon, Pak, Tibet, Ku, Asian
	Black	MLP.v^{32}_{5923}	Black, black, Black, black, blacks, 黑, 黑, Negro, BL, blk, African, BL, чер, Afr, Dark, ч, dark, Чер, лек, 블
		$\text{MLP.v}^{30}_{12572}$	Black, black, Black, black, blacks, 黑, BL, 黑, BL, Negro, blk, чер, dark, African, Afr, 블, лек, dark, Чер, ч
		$\text{MLP.v}^{29}_{13186}$	black, black, Black, blacks, Black, white, 黑, white, whites, African, 黑, Negro, 白, Afr, White, White, BL, чер, Af, brown
	Native Am.	MLP.v^{31}_{3440}	imperial, fasc, Imperial, colonial, militar, Kent, /******/, colon, popul, racist, gent, antal, neo, provinc, Emitter, colon, carriage, TES, omena, ounds
		$\text{MLP.v}^{29}_{12205}$	native, native, Native, Native, nat, ind, ab, igenous, Ma, nat, Ab, Ma, abor, Nat, primitive, Nav, born, Mas, nav, Ab

Table 12: Full top-20 tokens for race-associated neurons in Mistral-7B (ToxiGen). Translations: 黑 (black), 黑 (black), чер (black), ч (ch), Чер (black), лек (lek), 블 (black), 白 (white).

Model	Group	Neuron	Top Tokens
Llama-3.1-8B	Asian	MLP.v^{32}_{5691}	Chinese, China, Chinese, China, chinese, china, -China, Beijing, 中国, Çin, 중국, Shanghai, 中國, china, چی‌ن, 中国, Zhang, Jiang, Tencent, Guang
		$\text{MLP.v}^{31}_{14299}$	Li, yuan, Dong, Huang, Liu, Wang, Chen, Yang, Tian, Zhou, Ding, Wu, dong, dong, Feng, wang, Zhang, Qin, Jiang, Guang
		$\text{MLP.v}^{30}_{12566}$	East, East, EAST, Eastern, east, -East, eastern, 東, Eastern, east, 东, -east, 東, Doğu, 东, Đông, شرق, vých, вост, orient
	Black	MLP.v^{30}_{7195}	Mississippi, Jamaica, Jama, Caribbean, Louisiana, Trinidad, Haiti, LSU, Ghana, Hait, Baton, Harlem, Nigeria, negro, Negro, Jazz, Bahamas, Memphis, Nigerian, Zimbabwe
		$\text{MLP.v}^{29}_{13826}$	African, african, Africans, Africa, Africa, afr, Afro, Afr, frican, negro, Af, Afrika, Af, africa, Negro, blacks, black, af, frica, Blacks
	Latino	MLP.v^{32}_{9242}	Spanish, Hispanic, Spanish, spanish, Mexican, Argentine, Santiago, Mexico, Puerto, Madrid, pesos, Chavez, Ecuador, Spain, Juan, Hispan, Colombian, Hispanics, Carlos, Chile
	Native Am.	MLP.v^{32}_{6893}	colon, Colon, Colon, colon, colonial, Colonial, colonization, colonies, Colony, colony, Colonel, 殖, olon, icolon, Colin, OLON, kol, COL, -Col, colore
		MLP.v^{31}_{1186}	native, energy, Energy, native, energy, Native, Native, -native, Energy, natives, _native, .native, _energy, -energy, supply, 能源, culture, /native, nice, Higher

Table 13: Full top-20 tokens for race-associated neurons in Llama-3.1-8B (ToxiGen). Translations: 中国 (China), 中國 (China), 中国 (China), چی‌ن (China), Çin (China), 중국 (China), 東 (East), 东 (East), Doğu (East), Đông (East), شرق (East), vých (East), вост (East), 殖 (colonial/colonize), 能源 (energy).

Model	Group	Top 20 Tokens
Qwen2.5-7B	White	白, generado, -Nazi, ksam, onn, pl, owl, ucas, *&, `#{, .toByteArray, _EXTENSIONS, avras, onFailure, 镪, ski, Nederland, 好运, Luft, Ski
	Asian	Asian, Asian, Asians, Asia, 亚洲, Asia, asia, Asi, アジア, asiat, asian, 舢, Taiwan, Singapore, 舶, Singapore, Tai, Canton, Taiwanese, ракти
	Black/AA	african, African, 非洲, -AA, Africans, Africa, ienda, Africa, السود, 疟, mongo, /black, aina, AA, .BL, .Black, .Mongo, Afro, Nigerian, Black
Mistral-7B	White	ogle, ASC, cip, Kurt, heid, unächst, kle, þ, criptor, ór, NOP, och, zym, hid, eu, cow, cí, zens, vas, awa
	Asian	Asian, Taiwan, Japanese, Hong, Malays, Japan, Singapore, Chinese, Malaysia, Pak, Indones, Korean, Tai, Asia, Philippines, Philipp, Tokyo, Tok, jap, Sri
	Black/AA	African, Afr, Africa, blacks, Niger, Jama, Negro, Nigeria, Black, Af, ament, slavery, sist, black, external, external, Af, AA, lando, slave
Llama-3.1-8B	White	ithe, lan, Fres, ycz, hs, Wake, Bread, bread, itra, bairro, Fans, 1, Josh, wie, Jackets, ジオ, Marina, γι, Josh, avan
	Asian	Asian, Asian, Asians, asian, Indonesian, Asia, Taiwanese, Asia, asiat, Vietnamese, Japanese, Korean, Oriental, Malaysian, Buddhist, 亚洲, orean, Filipino, Chinese, Indones
	Black/AA	black, African, Black, african, black, frican, Afro, /black, negro, Negro, 黑, blacks, BLACK, zwarte, 黑, Black, .Black, _black, Africa, Africans

Table 14: Full top-20 probe token projections for C-REACT direct mentions (explicit race/ethnicity). Translations: 白 (white), generado (generated), 镪 (money), Nederland (Netherlands), 好运 (good luck), Luft (air), 亚洲 (Asia), アジア (Asia), 舢 (Shantou), 舶 (ship), ракти (practice), 非洲 (Africa), السود (black), 疟 (malaria), unächst (initially), þ (th), ór (or), cí (here), zens (citizens), bairro (neighborhood), ジオ (Geo), γι (gi), 黑 (black), zwarte (black).

Model	Group	Top 20 Tokens
Qwen2.5-7B	White	RaisePropertyChanged, Russian, 俄罗斯, 橡, 圬, ocaly, Russia, _backward, Kremlin, Russian, *)((, .defaultValue, RU, russian, egrator, ipsis, Rad, UCE, abis, ENCIL
	Asian	.Dao, 华人, Chinese, 嵴, .insertBefore, Xia, chinese, ettel, Chinese, Tibetan, China, Wong, /apache, 华南, dao, dx, 갔, Jun, utenant, stylesheet
	Black/AA	Hait, Haiti, Tropical,)__, eneg, %;", ($., 堍, %;">, 彗, tropical, 垓, |; .iterator, 热带, %;">,];//, loo, estring, *****
Mistral-7B	White	Moscow, Russian, Ukrain, Ukraine, Russians, Polish, Ukr, Russia, vod, Soviet, Vlad, russ, Kaz, Bulgar, Lieutenant, Mik, icz, Roma, dou, Serge
	Asian	Korea, Korean, Asian, trag, apis, Shan, Vietnam, gram, sg, apore, ga, Aires, Assembly, WD, lag, Nor, Viet, Schw, gan, cent
	Black/AA	Caribbean, Jama, Braz, Af, Cuba, Niger, Nigeria, Cub, Bah, São, Core, Currency, Bras, island, Hy, hur, Curt, Af, Brazil, mont
Llama-3.1-8B	White	Russia, Kremlin, Russian, Russians, Putin, Moscow, Ukraine, Putin, Ukrainian, Russian, Russia, Ukrain, russian, Rus, Belarus, Rus, russe, russ, Kiev, Rusya
	Asian	Cambodia, Asian, Chung, Cheng, Chinese, Camb, Kang, wang, Chinese, chinese, Buddhism, asian, Hong, Bangalore, Buddhist, Bang, Asian, Malaysia, Korean, epr
	Black/AA	Hait, Haiti, Caribbean, Maurit, Dominican, Bahamas, hait, Cre, Cre, Jama, ibbean, Trinidad, ingt, Jean, Cameroon, François, Maurice, Jean, анка, Santo

Table 15: Full top-20 probe token projections for C-REACT indirect mentions (language/country). Translations: 俄罗斯 (Russia), 橡 (yuan; Chinese monetary unit), 圬 (plaster), 华人 (ethnic Chinese), 嵴 (ridge), 华南 (South China), 갔 (went), 堍 (clay), 彗 (comet), 垓 (vast number), 热带 (tropical), São (Saint/São), Rusya (Russia), epr (erg), анка (anka).

Model	Group	Neuron	Top 20 Tokens
Qwen2.5-7B	White	$MLP.v_{16880}^{28}$	英国, Dutch, French, 法国, 意大利, Italian, German, 荷兰, France, British, Belgian, European, Britain, Germany, 德国, Spanish, Italy, Spain, 欧洲, French
		$MLP.v_{17660}^{27}$	German, 万欧元, ドイツ, 德国, German, Germany, Germans, EU, 荷兰, euro, €, 欧盟, Germany, Swiss, 欧元, Euro, Dutch, 瑞士, 欧, Switzerland
		$MLP.v_{4457}^{27}$	UEFA, 地中海, Mediterranean, Cyprus, Pane, Roman, 欧盟, Euro, Roman, chalk, 希腊, Europe, UES, Rom, 罗马, EU, Kn, Europe, 橡, 騰
		$MLP.v_{8669}^{25}$	fod,].', []>, sterdam, Ủ, Spain, Gaines, Spain, 比利时, □, WRITE, 英国, europe, <!–[, =======, %@", 示, 西班牙, _IMPORTED, euros
	Asian	$MLP.v_{13406}^{28}$	在日本, Japanese, 日本, Japanese, Japan, Japan, 日本人, Tokyo, japan, 东京, 일본, 日军, japanese, japon, Nhật, Tok, Hiro, 日本の, Osaka, jap
		$MLP.v_{16570}^{28}$	习, 習, 惯, habit, 习惯, 的习惯, habitual, Habit, 習慣, habit, habits, accustomed, 习近平, 习惯了, 慣, 习近, 惯例, 习俗, 习近平总, 总书记
		$MLP.v_{6943}^{27}$	Asian, Asian, 亚洲, Asia, Chinese, Asia, Chinese, China, 印度, Indian, India, India, Asians, China, Indian, chinese, 中國, Oriental, Japanese, asian
		$MLP.v_{217}^{27}$	Asian, 亚洲, Asia, Asian, Asians, Asia, asian, asia, アジア, 亚, 亞, Asi, asia, asiat, asi, 亚太, เอเช, _As, asi, AS
		$MLP.v_{5187}^{26}$	Chinese, Italian, Chinese, Irish, Italian, Asian, Japanese, chinese, Mexican, Vietnamese, Korean, German, Greek, Asian, Portuguese, Latin, Indian, African, German, italian
		$MLP.v_{3828}^{26}$	亚洲, Asian, Asia, Asian, Latin, 欧, Asia, 亚太, 欧洲, 非洲, 东亚, Southeast, Europe, 拉丁, European, 东南亚, 美洲, アジア, Euras, Asians
		$MLP.v_{15029}^{25}$	Chinese, China, China, Chinese, chinese, china, china, 中国, 中国的, -China, 중국, الصين, 中國, Asian, 在中国, 是中国, 亚洲, 由中国, Asian, Asia
	Black/AA	$MLP.v_{11088}^{28}$	racial, racially, racist, 种族, racism, racial, Harlem, segregated, segregation, Rac, interracial, Ethnic, Afro, hetto, ethnic, rac, Interracial
		$MLP.v_{10048}^{28}$	µ, ctx, ctx, µ, (ctx, Ken, Ken, Ctx, Cape, _ctx, .ctx, ctx, µ, _CTX, 南非, context, CTX, Kenya, µ, context
		$MLP.v_{2240}^{27}$	black, 黑, black, Black, Black, 黑色, -black, BLACK, BLACK, _black, blacks, .black, /black, ブラック, .Black, 黑, 黑白, _BLACK, blacklist
		$MLP.v_{16596}^{27}$	Ken, Ken, 肯, Nam, Liber, Bot, Tanz, Burk, Bur, Chad, ken, 赞, Nam, Kenya, Ug, Jordan, ken, Maur, Per, Jordan
		$MLP.v_{18261}^{26}$	黑, black, black, 黑色, Black, -black, Black, _black, .black, dark, 黑, BLACK, /black, 黑暗, blacks, .Black, BLACK, den, darken
		$MLP.v_{1091}^{26}$	黑, black, 黑色, Black, Black, black, -black, BLACK, BLACK, /black, blacks, _black, dark, Blacks, Dark, blacklist, 黑暗, 黑, .Black
		$MLP.v_{10230}^{25}$	Af, af, 非洲, Af, African, slave, AF, Africans, AF, african, slavery, 奴隶, 奴, slaves, Slave, slave, Afro, Slave, Africa, afr
		$MLP.v_{10739}^{25}$	非洲, African, Africa, Africa, Africans, african, Ghana, Nigerian, Nigeria, afr, Kenya, Nairobi, Niger, -Saharan, Nd, Tanzania, Uganda, africa, Nz, Lagos
Mistral-7B	White	$MLP.v_{1606}^{32}$	England, France, Europe, Switzerland, Britain, Holland, Spain, Australia, España, Austria, Germany, Denmark, Francia, Canada, Europa, America, Wales, Deutschland, Belgium, Sweden
		$MLP.v_{12760}^{32}$	€, €, Finn, Mediterranean, Belgium, Italy, Finland, Belg, Luxem, Netherlands, Milan, Italian, Portugal, Denmark, Czech, Madrid, Norway, Sweden, Spain, Amsterdam
		$MLP.v_{6831}^{32}$	European, Eu, Europe, Europe, EU, eu, Europ, Euro, europe, eu, Europa, europ, Европ, urope, UEFA, €, urop, €, €, вро
		$MLP.v_{1521}^{30}$	Pennsylvania, Michigan, Kansas, Philadelphia, Verm, ahl, Detroit, rus, ники, zel, Vic, Buc, ru, dispos, Hein, gotta, xa, una, anne, onic
	Asian	$MLP.v_{1153}^{32}$	Japanese, Korean, Kol, Japan, Kor, Kaz, Sak, Taiwan, Korea, Nak, Kur, Kom, Bangl, Pakistan, Kab, Kon, Pak, Tibet, Ku, Asian
		$MLP.v_{3946}^{31}$	Japanese, anime, Japan, Tokyo, jap, Jap, Korean, oji, apan, Tok, Asian, Taiwan, Israeli, Chinese, Korea, Indones, aku, Viet, Sak, oshi
		$MLP.v_{2137}^{30}$	Wil, mal, Mal, Ker, Mal, mal, Malays, Malaysia, ker, <\, $\, -\, $\, Carm, ($\, /\, mals, -\, +\, >\
		$MLP.v_{8986}^{30}$	asis, asi, Ash, lic, ox, Asia, yk, ke, ash, ym, as, ass, Async, asia, ais, eper, ob, ash, ek
	Black/AA	$MLP.v_{2923}^{32}$	Black, black, Black, black, blacks, 黑, 黑, Negro, BL, blk, African, BL, чер, Afr, Dark, ч, dark, Чер, лек, 블
		$MLP.v_{8715}^{31}$	African, Africa, Afr, Niger, Af, frica, Kenya, Af, Nigeria, fr, Negro, af, Afghan, blacks, FR, airo, Caribbean, Johannes, □, Jama
		$MLP.v_{3398}^{30}$	African, Africa, Afr, Af, Af, Asian, fr, Niger, Aaron, Afghan, Kenya, frica, af, AF, Mexican, AF, European, Nigeria, af, Egyptian
		$MLP.v_{2572}^{30}$	Black, black, Black, black, blacks, 黑, BL, 黑, BL, Negro, blk, чер, dark, African, Afr, 블, лек, dark, Чер, ч
		$MLP.v_{13186}^{29}$	black, black, Black, blacks, Black, white, 黑, white, whites, African, 黑, Negro, 白, Afr, White, White, BL, чер, Af, brown
Llama-3.1-8B	White	$MLP.v_{9094}^{31}$	White, white, White, white, WHITE, -white, 白, WHITE, _white, .White, whites, 白, .white, _WHITE, Whites, beyaz, :white, trắng, سفيد, .WHITE
	Asian	$MLP.v_{2691}^{32}$	Chinese, China, Chinese, China, chinese, china, -China, Beijing, 中国, Çin, 중국, Shanghai, 中國, china, چين, 中国, Zhang, Jiang, Tencent, Guang
		$MLP.v_{14299}^{31}$	Li, yuan, Dong, Huang, Liu, Wang, Chen, Yang, Tian, Zhou, Ding, Wu, dong, dong, Feng, wang, Zhang, Qin, Jiang, Guang
		$MLP.v_{5272}^{29}$	Asia, Asia, Asian, Asian, asia, continent, Africa, continental, asian, Africa, asia, Asians, Continental, 亚洲, Asi, Europe, Continent, Latin, asi, Europe
	Black/AA	$MLP.v_{7195}^{30}$	Mississippi, Jamaica, Jama, Caribbean, Louisiana, Trinidad, Haiti, LSU, Ghana, Hait, Baton, Harlem, Nigeria, negro, Negro, Jazz, Bahamas, Memphis, Nigerian, Zimbabwe
		$MLP.v_{3868}^{30}$	black, Black, BLACK, Black, black, /black, BLACK, -black, 黑, blacks, _black, .Black, blacklist, 黑, Blacks, .black, 黑, Blackburn, _BLACK, blackout
		$MLP.v_{13826}^{29}$	African, afrıcan, Africans, Africa, Africa, afr, Afro, Afr, frican, negro, Af, Afrika, Af, africa, Negro, blacks, black, af, frica, Blacks
		$MLP.v_{6824}^{29}$	tropical, jungle, Tropical, Jungle, jung, ropical, trop, Fiji, Belize, mango, Congo, Jama, Caribbean, Hait, Honduras, Mango, BSON, Haiti, Safari, Jamaica

Table 16: Full list of race-associated neurons identified from C-REACT direct mentions. Translations: 英国 (UK), 法国 (France), 意大利 (Italy), 荷兰 (Netherlands), 德国 (Germany), 欧洲 (Europe), 万欧元 (ten-thousand euros), ドイツ (Germany), 欧盟 (EU), 欧元 (euro), 瑞士 (Switzerland), 欧 (Europe), 地中海 (Mediterranean), 希腊 (Greece), 罗马 (Rome), 橡 (oak), 騰 (soar), Ủ (yes), 比利时 (Belgium), □ (subsetneq), 示 (show), 西班牙 (Spain), 在日本 (in Japan), 日本 (Japan), 日本人 (Japanese), 东京 (Tokyo), 일본 (Japan), 日军 (Japanese army), Nhật (Japan), 日本の (Japan), 习 (Xi), 習 (Xi), 惯 (habit), 习惯 (habit), 的习惯 (habit), 習慣 (habit), 习近平 (Xi Jinping), 习惯了 (used to), 慣 (habit), 习近 (Xi-), 惯例 (convention), 习俗 (custom), 习近平总 (Xi Jinping), 总书记 (general secretary), 亚洲 (Asia), 印度 (India), 中國 (China), 亚 (Asia), 亞 (Asia), 亚太 (Asia-Pacific), เอเช (Asia), 东亚 (East Asia), 欧洲 (Europe), 非洲 (Africa), 拉丁 (Latin), 东南亚 (Southeast Asia), 美洲 (Americas), アジア (Asia), 中国 (China), 中国的 (China), 중국 (China), الصين (China), 在中国 (China), 是中国 (China), 由中国 (China), 种族 (race), 南非 (South Africa), 黑 (black), 黑色 (black), ブラック (black), 黑 (black), 黑白 (black-and-white), 肯 (Ken), 赞 (praise), 黑暗 (dark), đen (black), 非洲 (Africa), 奴隶 (slave), 奴 (slave), Европ (Euro), Є (Euro), вро (Euro), ники (niki), 黑 (black), чер (black), ч (ch), Чер (black), лек (lek), 블 (black), 白 (white), beyaz (white), trắng (white), سفيد (white), 中国 (China), 中國 (China), چين (China), Çin (China), 중국 (China), 亚洲 (Asia).

Model	Group	Neuron	Top 20 Tokens
Qwen2.5-7B	White	MLP.v$^{28}_{6780}$	Maryland, Pennsylvania, Baltimore, Ohio, Wisconsin, Maine, Minnesota, Illinois, Pittsburgh, Iowa, Ohio, Connecticut, Philadelphia, Chicago, Michigan, Milwaukee, Nebraska, Detroit, Seattle, 江苏省
		MLP.v$^{28}_{2088}$	以色列, Israel, Jerusalem, Israel, Israeli, Noah, Hebrew, Zion, Moses, -Israel, Palestine, Israeli, Israelis, Biblical, biblical, Zionist, Luke, Palestinian, Jer, Nathan
		MLP.v$^{28}_{2318}$	Ohio, Columbus, sz, Cleveland, Ohio, (sz, sz, 浙江, 杭州, Sz, Sz, 浙, 浙江省, 杭州市, SZ, Cincinnati, 杭, _sz, Akron, Croatian
		MLP.v$^{27}_{17660}$	German, 万欧元, ドイツ, 德国, German, Germany, Germans, EU, 荷兰, euro, €, 欧盟, Germany, Swiss, 欧元, Euro, Dutch, 瑞士, 欧, Switzerland
		MLP.v$^{30}_{3012}$	犹, Jew, Jewish, 猶, Jews, Judaism, ев, jewish, Juda, Hebrew, евр, 以色列, Israel, kosher, synagogue, -J, ewish, jew, Rabbi, Israeli
		MLP.v$^{30}_{3382}$	俄, Rus, Russ, RU, rus, Russians, Russ, Russia, Russian, russe, 俄罗斯, russ, Russo, 俄国, 莫斯科, Mos, Russia, Moscow, -Russian, RU
		MLP.v$^{30}_{4157}$	UEFA, 地中海, Mediterranean, Cyprus, Pane, Roman, 欧盟, Euro, Roman, chalk, 希腊, Europe, UES, Rom, 罗马, EU, Kn, Europe, 橡, 騰
		MLP.v$^{31}_{7123}$	Ukraine, 乌, 乌克兰, Russia, Ukrainian, Russian, 俄罗斯, Ukrain, 烏, 乌克, Russia, war, Russian, -Russian, 冲突, russian, Russo, Conflict, conflict, 地
	Asian	MLP.v$^{28}_{13406}$	在日本, Japanese, 日本, Japanese, Japan, Japan, 日本人, Tokyo, japan, 东京, 일본, 日军, japanese, japon, Nhật, Tok, Hiro, 日本の, Osaka, jap
		MLP.v$^{27}_{4943}$	Asian, Asian, 亚洲, Asia, Chinese, Asia, Chinese, China, 印度, Indian, India, India, Asians, China, Indian, chinese, 中國, Oriental, Japanese, asian
		MLP.v$^{27}_{229}$	Mã, Opcode, (equalTo, Singapore, Singapore, ')", wię, (defvar, .tp, Ла, Ново, yne, (ls, Tah, 海南, établissement, airobi, I, 槟
		MLP.v$^{26}_{9908}$	لال, لالتن, Malays, Indones, protester, HttpMethod, Zambia, الhج, Rodrig, Taiwanese, 辇, Jama, 加工厂, الاستث, Karnataka, UserDao, Rousse, Böl, african
		MLP.v$^{26}_{13889}$	chop, Conf, Mand, Canton, 功夫, mand, Dragon, Chop, Dragon, Conf, mand, dragon, Cant, dragon, Fu, Portsmouth, Bruce, Cath, 洗衣, Fuk
		MLP.v$^{29}_{2001}$	越, Vietnam, 越南, Viet, Nguyen, Viet, Vietnamese, .vn, Ho, Ph, Vu, Tran, viet, 就越, Ph, Dien, 越大, 越好, anh, gia
		MLP.v$^{29}_{15029}$	Chinese, China, China, Chinese, chinese, china, china, 中国, 中国的, -China, 중국, الصين, 中國, Asian, 在中国, 是中国, 亚洲, 由中国, Asian, Asia
	Black/AA	MLP.v$^{25}_{10230}$	Af, af, 非洲, Af, African, slave, AF, Africans, AF, african, slavery, 奴隶, 奴, slaves, Slave, slave, Afro, Slave, Africa, afr
		MLP.v$^{25}_{10739}$	非洲, African, Africa, Africa, Africans, african, Ghana, Nigerian, Nigeria, afr, Kenya, Nairobi, Niger, -Saharan, Nd, Tanzania, Uganda, africa, Nz, Lagos
Mistral-7B	White	MLP.v$^{32}_{2399}$	Russian, Vlad, Moscow, Soviet, Russia, Ukrain, Russians, Ukraine, Czech, Ukr, Bulgar, Slov, vod, Cz, Aleks, Serge, sov, russ, Polish, Stalin
		MLP.v$^{32}_{2356}$	Russell, Rus, Russ, Russian, rus, Russia, russ, Russians, Soviet, Moscow, pyc, rust, rus, rust, py, sov, Rug, Rud, ruin, Ruth
		MLP.v$^{30}_{3487}$	Russell, Russ, russ, Russian, Rus, Russia, Russians, rus, rus, Moscow, рус, Росси, России, uss, Soviet, ussia, ussian, Vlad, USS, sov
		MLP.v$^{29}_{260}$	Italian, Ital, Italy, ital, italien, Italia, italiano, ital, Giovanni, Francesco, Carlo, Gian, Sic, Gi, IT, Milan, Serie, acci, Giov, IT
	Black/AA	MLP.v$^{31}_{8715}$	African, Africa, Afr, Niger, Af, frica, Kenya, Af, Nigeria, fr, Negro, af, Afghan, blacks, FR, airo, Caribbean, Johannes, □, Jama
Llama-3.1-8B	White	MLP.v$^{32}_{10606}$	Russian, Russians, Russia, Russian, Moscow, Soviet, Russia, -Russian, Putin, russ, russian, Kremlin, 俄, russe, Russ, USSR, Russell, .ru, Vladimir, Sergei
		MLP.v$^{32}_{10409}$	Baltic, Sloven, Celtic, Gron, Flem, Slovenia, Austria, Sch, Luxembourg, Monaco, Gros, Croatia, Naples, Blond, Croatian, Baron, Austrian, Mont, Trent, Malta
		MLP.v$^{31}_{12584}$	Russell, Russ, Russ, Rus, RUS, Rus, russ, rus, Russo, Russia, Russian, Rusya, 俄, Russians, Рус, -Russian, rus, Russia, russe, Russian
		MLP.v$^{29}_{4193}$	Czech, Hungarian, Boh, Slovak, Budapest, Hungary, Hung, Slovakia, Prague, Poland, polish, Polish, Sloven, Croatian, Krak, lovak, Brno, Slovenia, Romania, Croatia
	Asian	MLP.v$^{32}_{3691}$	Chinese, China, Chinese, China, chinese, china, -China, Beijing, 中国, Çin, 중국, Shanghai, 中國, china, چين, 中国, Zhang, Jiang, Tencent, Guang
		MLP.v$^{32}_{6950}$	Bhar, Bh, Sinh, Nagar, Shah, Muk, Hindi, Allah, Tamil, Gujar, Bollywood, Kh, Dh, Gujarat, Punjab, Kh, Mumbai, Uttar, Maharashtra, Jain
		MLP.v$^{29}_{10616}$	Indian, Indian, Indians, India, indian, India, india, Ân, Bollywood, الهند, Indi, Indianapolis, Mumbai, □, Hindu, Bombay, Hindi, Hindus, 印, Delhi
	Black/AA	MLP.v$^{30}_{7195}$	Mississippi, Jamaica, Jama, Caribbean, Louisiana, Trinidad, Haiti, LSU, Ghana, Hait, Baton, Harlem, Nigeria, negro, Negro, Jazz, Bahamas, Memphis, Nigerian, Zimbabwe
		MLP.v$^{29}_{6824}$	tropical, jungle, Tropical, Jungle, jung, ropical, trop, Fiji, Belize, mango, Congo, Jama, Caribbean, Hait, Honduras, Mango, BSON, Haiti, Safari, Jamaica
		MLP.v$^{29}_{13826}$	African, african, Africans, Africa, Africa, afr, Afro, Afr, frican, negro, Af, Afrika, Af, africa, Negro, blacks, black, af, frica, Blacks

Table 17: Full list of race-associated neurons identified from C-REACT indirect mentions (language/country). Translations: 江苏省 (Jiangsu), 以色列 (Israel), 万欧元 (ten-thousand euros), ドイツ (Germany), 德国 (Germany), 荷兰 (Netherlands), 欧盟 (EU), 欧元 (euro), 瑞士 (Switzerland), 欧 (Europe), 犹 (Jew), 猶 (Jew), ев (Jew), евр (Jew), 俄 (Russia), 俄罗斯 (Russia), 俄国 (Russia), 莫斯科 (Moscow), 地中海 (Mediterranean), 希腊 (Greece), 罗马 (Rome), 橡 (oak), 騰 (soar), 乌 (Ukraine), 乌克兰 (Ukraine), 烏 (Ukraine), 乌克 (Ukraine), 冲突 (conflict), 地 (land), 在日本 (in Japan), 日本 (Japan), 日本人 (Japanese), 东京 (Tokyo), 일본 (Japan), 日军 (Japanese army), Nhật (Japan), 日本の (Japan), 亚洲 (Asia), 印度 (India), 中國 (China), Mã (code), wię (bond), Ла (La), Ново (Novo), 海南 (Hainan), 槟 (Penang), Î (I), الثال (Tue), الإثن (Mon), الhج (jihad), الاستث (exception), 辇 (carriage), 加工厂 (factory), Böl (region), 功夫 (kung fu), 洗衣 (laundry), 越 (Viet), 越南 (Vietnam), 就越 (then more), 越大 (bigger), 越好 (better), 中国 (China), 中国的 (China), 중국 (China), الصين (China), 在中国 (China), 是中国 (China), 由中国 (China), 非洲 (Africa), 奴隶 (slave), 奴 (slave), рус (Russian), ру (ru), Росси (Russia), России (Russia), italiano (Italian), italien (Italian), Italia (Italy), Ân (India), الهند (India), 印 (India).