

# Do Personality Traits Interfere? Geometric Limitations of Steering in Large Language Models

Anonymous ACL submission

## Abstract

Personality steering in large language models (LLMs) commonly relies on injecting trait-specific steering vectors, implicitly assuming that personality traits can be controlled independently. In this work, we examine whether this assumption holds by analysing the geometric relationships between Big Five personality steering directions. We study steering vectors extracted from two model families (LLaMA-3-8B and Mistral-8B) and apply a range of geometric conditioning schemes, from unconstrained directions to soft and hard orthonormalisation. Our results show that personality steering directions exhibit substantial geometric dependence: steering one trait consistently induces changes in others, even when linear overlap is explicitly removed. While hard orthonormalisation enforces geometric independence, it does not eliminate cross-trait behavioural effects and can reduce steering strength. These findings suggest that personality traits in LLMs occupy a slightly coupled subspace, limiting fully independent trait control.

## 1 Introduction and Background

Large Language Models (LLMs) have demonstrated significant advances in their ability to exhibit personality traits (Jiang et al., 2024, 2023; Serapio-García et al., 2023), often aligned with the Big Five personality framework. Extensive prior work has explored the evaluation (Bhandari et al., 2025; Pellert et al., 2024), extraction (Jiang et al., 2024), and steering (Zhu et al., 2024; Chen et al., 2025; Bhandari et al., 2026) of personality traits in the literature. Recent advances in activation engineering have enabled the steering of LLM behaviour by injecting *activation vectors* into the model’s residual stream to control the strength of trait expression. This approach offers several advantages over fine-tuning (Zhu et al., 2024), which is typically a heavy and resource-intensive process.

Model behaviour can instead be adjusted at inference time using a *precise and controllable “knob”*, allowing for flexible and efficient personality modulation.

However, these studies do not emphasise the inherently entangled nature of personality traits, commonly referred to as the OCEAN traits (Openness, Conscientiousness, Extraversion, Agreeableness, and Neuroticism), and instead draw conclusions based on the independent manipulation of individual traits. Our findings indicate that models do not learn human psychological constructs as orthogonal basis vectors; rather, they acquire entangled representations shaped by the correlations present in their training data. Empirically, we demonstrate that increasing a single dimension, such as Openness, does not result in an isolated change along that axis. Instead, it simultaneously increases other dimensions, including Agreeableness, Conscientiousness, and Extraversion, while decreasing Neuroticism. These observations suggest that, rather than reflecting shifts along a single trait dimension, such changes are better understood as movements along a broader social axis.

Previous work has shown that high-level behaviours in large language models can often be approximated as linear directions in the activation space, enabling behaviour steering via difference-of-means or PCA-based vectors (Zou et al., 2023). However, such methods typically assume independence between concept directions. The concept-erasure literature studies how correlated information can be removed from representations, such as with Iterative Nullspace Projection (Ravfogel et al., 2020), introducing greedy, order-dependent projections, and LEACE: Perfect Linear Concept Erasure in Closed Form (Belrose et al., 2023), proposing a global, order-independent solution. Also, interpretability work such as Toy Models of Superposition (Elhage et al., 2022) shows that models often encode correlated concepts in superposition,

suggesting that geometric orthogonality alone may not guarantee semantic independence. Our work builds on these insights by applying both greedy and global Orthogonalisation strategies to personality steering and empirically evaluating their effect on cross-trait interference.

The overarching aim of this work is to systematically understand the personality trait vectors used in activation engineering. Prior activation engineering methods typically apply one trait direction at a time and evaluate whether the target trait shifts as intended (Sun et al., 2025; Yang et al., 2024; Bhandari et al., 2026), while leaving the effects on non-target traits unexamined. We systematically analyse the geometric relationships between Big Five personality directions and test how different constraint strategies, ranging from no constraints to full orthonormalisation, change both trait specificity and downstream behaviour. Rather than assuming that traits should be fully separable, we treat interference as an empirical phenomenon to be measured. Our goal is to understand when enforcing geometric independence helps, when it harms steering, and what this reveals about how personality traits are encoded inside the model.

We propose the following research questions:

1. **RQ1:** We hypothesise that personality steering directions are geometrically independent in large language models.
2. **RQ2:** What happens to steering effectiveness when geometric independence is explicitly enforced?
3. **RQ3:** Are observed trait dependencies consistent across different model families?

## 2 Methodology

The scope of this work focusses on understanding the behaviour of personality steering vectors and their interactions. We adopt the hybrid layer selection approach from (Bhandari et al., 2026) as our steering baseline, due to its effectiveness while preserving the model capacity. Our methodology then builds on this framework to analyse geometric dependencies between personality traits under different constraint settings.

**Steering Mechanism.** Trait directions are estimated from high/low labelled data (Li et al., 2025) at each layer and aggregated into a single weighted direction per trait, capturing consistent personality

signals across the network. An offline prior layer is selected for each trait using neutral probe prompts by applying a small steering signal and measuring distributional sensitivity at the next token. At runtime, a lightweight dynamic check adapts the layer choice to the input prompt. Steering is applied via forward hooks using projected and intensity-scaled trait vectors, and evaluated under base, positive, and negative settings using both personality questionnaires and generation-based assessments.

**Trait Direction Conditioning (C0–C5).** Let  $\mathbf{d}_c \in \mathbb{R}^D$  denote the normalised weighted steering direction for trait  $c$ . Empirically, these directions are not independent and exhibit substantial cosine overlap. To study how geometric constraints affect steering behaviour and cross-trait interference, we construct the following conditioning schemes:

- **C0 (Baseline).** Original trait directions  $\mathbf{d}_c$  are used without modification.
- **C1 (Soft Symmetric Whitening).** Directions are stacked into  $\mathbf{D}$  and transformed via a regularised Gram matrix:  $\mathbf{D}' = ((1 - \gamma)\mathbf{G} + \gamma\mathbf{I})^{-1/2}\mathbf{D}$ , where  $\gamma \in (0, 1)$  is a shrinkage parameter. This scales down off-diagonal correlations without forcing strict orthogonality, preserving more of the original shared geometry than hard whitening.
- **C2 (Greedy Orthogonalisation).** A Gram–Schmidt procedure sequentially removes projections  $\langle \mathbf{d}_i, \mathbf{d}_j \rangle \mathbf{d}_j$ , yielding an orthonormal basis that is order-dependent.
- **C3 (Selective Orthogonalisation).** Projection is applied only when  $|\cos(\mathbf{d}_i, \mathbf{d}_j)| > \tau$ , preventing over-disentanglement while suppressing dominant overlaps.
- **C4 (Soft Projection).** Correlated components are partially attenuated as  $\mathbf{d}_i \leftarrow \mathbf{d}_i - \beta \langle \mathbf{d}_i, \mathbf{d}_j \rangle \mathbf{d}_j$  when  $|\cos| > \tau$ , trading off disentanglement strength and retention.
- **C5 (Hard Orthonormalisation).** A symmetric Löwdin transformation enforces  $\mathbf{D}'\mathbf{D}'^T = \mathbf{I}$ , completely removing linear overlap in an order-independent manner.

All conditions use the same steering injection and hybrid layer selection mechanism, isolating the impact of geometric constraints on steering efficacy and trait interference.

Target Trait	Llama-3-8B						Mistral-8B					
	C0 (Base)		C4 (Soft)		C5 (Hard)		C0 (Base)		C4 (Soft)		C5 (Hard)	
	$T$	$B_{max}$	$T$	$B_{max}$	$T$	$B_{max}$	$T$	$B_{max}$	$T$	$B_{max}$	$T$	$B_{max}$
<b>Openness</b>	3.1	-3.5 (Neu)	3.0	-3.4 (Neu)	2.9	-3.0 (Neu)	3.3	3.3 (Agr)	3.2	3.4 (Agr)	3.1	2.8 (Agr)
<b>Conscientiousness</b>	2.9	-2.9 (Neu)	2.9	-2.7 (Neu)	2.9	2.6 (Agr)	2.2	-2.0 (Neu)	2.3	2.0 (Agr)	2.4	2.0 (Agr)
<b>Extraversion</b>	3.0	-3.1 (Neu)	2.6	-2.4 (Neu)	3.0	-2.5 (Neu)	3.1	3.3 (Agr)	3.5	3.1 (Agr)	3.3	3.0 (Agr)
<b>Agreeableness</b>	3.3	2.8 (Con)	3.2	2.6 (Opn)	<b>3.7</b>	-3.1 (Neu)	2.7	-3.3 (Neu)	3.7	-3.0 (Neu)	2.8	-2.3 (Neu)
<b>Neuroticism</b>	3.1	-3.1 (Agr)	3.1	-3.2 (Agr)	3.2	-3.1 (Agr)	0.7	-1.2 (Ext)	0.0	-1.0 (Agr)	0.1	-1.1 (Ext)

Table 1: **Trait-Level Steering Contrast under Geometric Constraints.** Comparison of target steering contrast ( $T$ , Intended Target) and maximum cross-trait bleed ( $B_{max}$ , Unintended Target) for LLaMA-3-8B and Mistral-8B under baseline (C0), soft-constrained (C4), and hard orthonormal (C5) trait vector constructions. Values report the difference between positively and negatively steered generations (High–Low) as measured by judge scores (1–5 scale). Parentheses indicate the trait responsible for  $B_{max}$ .

$T$ : Diagonal magnitude (Targeted Trait Steering).  $B_{max}$ : Maximum absolute off-diagonal value. **C4 (Soft)** uses  $\beta = 0.5$ . **C5 (Hard)** uses full symmetric Orthogonalisation. (–) sign suggests the opposite nature of the trait effects.

### 3 Evaluation

Personality steering is evaluated under three controlled settings: *base* (no steering), *positive steering*, and *negative steering*, where the latter two apply trait-specific steering vectors of equal magnitude and opposite polarity. For each geometric condition (C0–C5), all other factors are held constant, including the learned subspace, layer weights, steering intensity, injection point, and decoding configuration. This isolates the effect of geometric constraints on personality vectors, ensuring that observed differences arise solely from vector structure rather than the steering mechanism itself.

Steering effectiveness is assessed using interview-style Big Five Inventory (BFI) questionnaires (Wang et al., 2024) consistent with (Bhandari et al., 2026). For direct comparison, first and second order statistics are reported. Beyond target trait shifts, we measure cross-trait responses to quantify inter-trait effects. Evaluation uses neutral prompts, ensuring observed personality changes arise solely from internal steering. We use Gpt-4o-mini as a judge, building upon the literature (Jiang et al., 2024; Frisch and Giulianelli, 2024) of using models as judges. Finally, we report fluency scores alongside personality metrics to verify that steering and geometric constraints do not degrade generation quality or general language behaviour.

### 4 Results

We conduct experiments on two instruction-tuned models from different architectural families – *LLaMA-3-8B-Instruct* and *Ministral-8B-Instruct*.

For each model, baseline steering performance is quantified using the difference between positively and negatively steered generations (*high–low*), which serves as a reference point to compare how geometric constraints (C1–C5) alter both target-trait control and cross-trait interactions.

#### 4.1 Geometric Independence of Personality Steering Directions

To evaluate whether personality steering directions are geometrically independent, we compare the *target steering strength* ( $T$ , diagonal entries) against the *maximum cross-trait bleed* ( $B_{max}$ , largest absolute off-diagonal entry) under different geometric constraints (Table 1). Across both LLaMA-3-8B and Mistral-8B, steering a single trait produces non-negligible changes in at least one other trait, with  $B_{max}$  often comparable in magnitude to  $T$ . This pattern persists under both soft disentanglement (C4) and full symmetric orthonormalisation (C5). Notably, while C5 enforces near-zero pairwise cosine similarity between trait directions in activation space, it does not consistently reduce  $B_{max}$  in generation-level evaluations.

These results indicate that eliminating geometric overlap between steering vectors does not guarantee behavioural independence. Even when trait directions are orthonormal by construction, downstream generations continue to exhibit systematic cross-trait interactions. We therefore **reject the hypothesis in RQ1**, implying that personality steering directions in large language models are not geometrically independent in a behaviourally meaningful sense.

Table 1 reports the steering strength for condi-




Figure 1: **Fluency profile analysis for conditions C0–C5** compared against Base steering. *Fluency degradation for both Positive and Negative steering can be observed across all traits for all the Conditional Methods used.* Although trait shifts that were comparable to the base values, the significant degradation of fluency suggests the need to use orthogonalised vectors carefully for steering purpose.

tions base(C0), C4 and C5. Additionally, we analyze *fluency* and *variance* scores under identical measurement conditions, using LLaMA as the reference model.

Steering strength remains largely conserved under progressive orthonormalisation (C1–C5)(Table 1). However, Figure 1 demonstrates how fluency scores consistently degrade as geometric constraints are enforced. Comparing to baseline fluency values reported in (Bhandari et al., 2026) (High/Low), Openness drops from 5.0/3.8 to 4.1/3.2 (C1), 4.3/2.8 (C2), and remains around 4.4/2.8 through C3–C5. Similarly, conscientiousness decreases from 4.8/3.5 to 4.3/3.1 in C1 following identical trends to Openness for other conditions, while Extraversion exhibits the largest decline, from 4.9/4.3 to 3.8/3.6.

These results indicate that although orthonormalisation preserves directional steering magnitude, it removes shared components necessary for fluent and expressive generation, leading to reduced variance and degraded output quality. Hence, to explain our **RQ2**: we conclude that explicitly enforcing geometric independence does not improve steering effectiveness and instead introduces a quality–control trade-off.

## 4.2 RQ3: Cross-Model Consistency of Trait Dependencies

To assess whether observed trait dependencies generalise across model families, we compare steering behaviour between *LLaMA-3-8B-Instruct* and *Mistral-8B-Instruct* under identical conditions (C0,

C4, C5) using the same extraction, constraint, and evaluation pipeline. Across both models, we observe consistent geometric patterns: several traits exhibit substantial cross-trait bleed even after enforcing geometric constraints. For example, **Openness** shows high maximum bleed values in both models, with  $B_{\max}$  remaining large under hard orthonormalisation (C5), e.g.,  $\approx 3.0$  in LLaMA-3 and  $\approx 2.8$  in Mistral-8B. Similarly, **Extraversion** and **Agreeableness** continue to induce strong off-diagonal effects across conditions, indicating that these dependencies are not artifacts of a single model but reflect shared structure in personality representations.

At the same time, we observe clear model-specific modulation in steering responsiveness. Most notably, **Neuroticism** exhibits strong target steering in LLaMA-3 ( $T \approx 3.1$ – $3.2$  across C0–C5), whereas Mistral-8B shows low target response ( $T \approx 0.0$ – $0.7$ ) under the same conditions, despite comparable geometric treatment. Importantly, this suppression persists even when geometric independence is enforced (C5), suggesting that the absence of behavioural response cannot be attributed solely to vector entanglement. Together, these results indicate that while cross-trait dependencies are largely consistent across model families, their behavioural expression is shaped by model-specific training and alignment constraints rather than geometry alone. A detailed table for all the observations (C1–C5) is provided in Appendix A.

## 5 Conclusion

Given the popularity of steering methods in the literature, we systematically analysed the behaviour of steering vectors under various conditions. Our method considered the Big Five Personality traits, and we investigated whether personality steering directions in large language models are geometrically independent, and how enforcing geometric constraints affects the steering behaviours. Through the analysis across two model families (Llama and Mistral), we show that personality traits are not independent directions in activation space. Even when strong global constraints such as symmetric orthonormalisation are applied, steering one trait consistently induces measurable changes in other unintended traits, indicating persistent cross-trait dependencies.

## 326 Limitations

327 This work studies personality steering behaviour  
328 using a limited set of large language models and  
329 personality datasets, and future work could extend  
330 the analysis to a broader range of model families  
331 and trait representations. While we focus on Big  
332 Five traits and judge-based evaluation, additional  
333 datasets and alternative evaluation frameworks may  
334 reveal further structure in trait interactions. Our  
335 analysis relies on linear geometric constraints; ex-  
336 ploring other orthogonalisation or projection meth-  
337 ods could provide a more comprehensive under-  
338 standing of trait disentanglement. Finally, we use  
339 LLMs as judges for behavioural assessment, and  
340 incorporating human evaluation or complementary  
341 metrics is left for future investigation.

## 342 6 Ethical Considerations

343 Steering large language models using latent vec-  
344 tors introduces ethical considerations, particularly  
345 when such steering is applied in uncontrolled or  
346 unsupervised settings. Steering vectors are learned  
347 approximations of complex behavioural traits and  
348 do not provide transparent or complete represen-  
349 tations of the values they encode; as a result, un-  
350 intended attributes or hidden information may be  
351 co-activated during steering. This raises concerns  
352 about value misalignment, especially in real-world  
353 deployments where subtle behavioural shifts could  
354 have downstream social or psychological impacts.  
355 Additionally, aggressive or poorly understood steer-  
356 ing may bypass safety mechanisms or distort model  
357 behaviour in ways that are difficult to detect or  
358 reverse. These risks highlight the importance of  
359 careful evaluation, interpretability, and constraint-  
360 aware steering when modifying model behaviour.

## 361 References

362 Nora Belrose, David Schneider-Joseph, Shauli Ravfogel,  
363 Ryan Cotterell, Edward Raff, and Stella Biderman.  
364 2023. Leace: Perfect linear concept erasure in closed  
365 form. *Advances in Neural Information Processing*  
366 *Systems*, 36:66044–66063.

367 Pranav Bhandari, Nicolas Fay, Sanjeevan Selvaganap-  
368 athy, Amitava Datta, Usman Naseem, and Mehwish  
369 Nasim. 2026. Activation-space personality steer-  
370 ing: Hybrid layer selection for stable trait control  
371 in LLMs. In *19th Conference of the European Chap-  
372 ter of the Association for Computational Linguistics*  
373 *(EACL’26)*.

Pranav Bhandari, Usman Naseem, Amitava Datta, Nico-  
374 las Fay, and Mehwish Nasim. 2025. Evaluating per-  
375 sonality traits in large language models: Insights  
376 from psychological questionnaires. In *Companion*  
377 *Proceedings of the ACM on Web Conference 2025*,  
378 pages 868–872. 379

Runjin Chen, Andy Ardit, Henry Sleight, Owain Evans,  
380 and Jack Lindsey. 2025. Persona vectors: Monitoring  
381 and controlling character traits in language models.  
382 *arXiv preprint arXiv:2507.21509*. 383

Nelson Elhage, Tristan Hume, Catherine Olsson,  
384 Nicholas Schiefer, Tom Henighan, Shauna Kravec,  
385 Zac Hatfield-Dodds, Robert Lasenby, Dawn Drain,  
386 Carol Chen, and 1 others. 2022. Toy models of su-  
387 perposition. *arXiv preprint arXiv:2209.10652*. 388

Ivar Frisch and Mario Giulianelli. 2024. Llm agents  
389 in interaction: Measuring personality consistency  
390 and linguistic alignment in interacting popula-  
391 tions of large language models. *arXiv preprint*  
392 *arXiv:2402.02896*. 393

Guangyuan Jiang, Manjie Xu, Song-Chun Zhu, Wen-  
394 juan Han, Chi Zhang, and Yixin Zhu. 2023. Evaluat-  
395 ing and inducing personality in pre-trained language  
396 models. *Advances in Neural Information Processing*  
397 *Systems*, 36:10622–10643. 398

Hang Jiang, Xiajie Zhang, Xubo Cao, Cynthia Breazeal,  
399 Deb Roy, and Jad Kabbara. 2024. Personallm: In-  
400 vestigating the ability of large language models to  
401 express personality traits. In *Findings of the asso-  
402 ciation for computational linguistics: NAACL 2024*,  
403 pages 3605–3627. 404

Wenkai Li, Jiarui Liu, Andy Liu, Xuhui Zhou, Mona  
405 Diab, and Maarten Sap. 2025. Big5-chat: Shap-  
406 ing llm personalities through training on human-  
407 grounded data. In *Proceedings of the 63rd Annual*  
408 *Meeting of the Association for Computational Lin-  
409 guistics (Volume 1: Long Papers)*, pages 20434–  
410 20471. 411

Max Pellert, Clemens M Lechner, Claudia Wagner,  
412 Beatrice Rammstedt, and Markus Strohmaier. 2024.  
413 Ai psychometrics: Assessing the psychological pro-  
414 files of large language models through psychometric  
415 inventories. *Perspectives on Psychological Science*,  
416 19(5):808–826. 417

Shauli Ravfogel, Yanai Elazar, Hila Gonen, Michael  
418 Twiton, and Yoav Goldberg. 2020. Null it out: Guard-  
419 ing protected attributes by iterative nullspace projec-  
420 tion. *arXiv preprint arXiv:2004.07667*. 421

Gregory Serapio-García, Mustafa Safdari, Clément  
422 Crepy, Luning Sun, Stephen Fitz, Marwa Abdulhai,  
423 Aleksandra Faust, and Maja Matarić. 2023. Personal-  
424 ity traits in large language models. 425

Seungjong Sun, Seo Yeon Baek, and Jang Hyun Kim.  
426 2025. Personality vector: Modulating personality of  
427

- 428 [large language models by model merging](#). In *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing*, pages 24656–  
429 24677, Suzhou, China. Association for Computational Linguistics.  
430  
431  
432
- 433 Xintao Wang, Yunze Xiao, Jen-tse Huang, Siyu Yuan,  
434 Rui Xu, Haoran Guo, Quan Tu, Yaying Fei, Ziang  
435 Leng, Wei Wang, and 1 others. 2024. Incharacter:  
436 Evaluating personality fidelity in role-playing agents  
437 through psychological interviews. In *Proceedings*  
438 *of the 62nd Annual Meeting of the Association for*  
439 *Computational Linguistics (Volume 1: Long Papers)*,  
440 pages 1840–1873.
- 441 Shu Yang, Shenzhe Zhu, Liang Liu, Lijie Hu, Mengdi  
442 Li, and Di Wang. 2024. Exploring the personality  
443 traits of llms through latent features steering. *arXiv*  
444 *preprint arXiv:2410.10863*.
- 445 Minjun Zhu, Yixuan Weng, Linyi Yang, and Yue Zhang.  
446 2024. Personality alignment of large language mod-  
447 els. *arXiv preprint arXiv:2408.11779*.
- 448 Andy Zou, Long Phan, Sarah Chen, James Campbell,  
449 Phillip Guo, Richard Ren, Alexander Pan, Xuwang  
450 Yin, Mantas Mazeika, Ann-Kathrin Dombrowski,  
451 and 1 others. 2023. Representation engineering: A  
452 top-down approach to ai transparency. *arXiv preprint*  
453 *arXiv:2310.01405*.

## A All detailed tables for C0-C5 for trait values

Table 2: **Condition C1 (Soft Symmetric Whitening)**. Cross-trait impact of steering vectors on Llama-3-8B. The rows represent the *Targeted Trait* (steering vector applied), and columns represent the *Measured Trait* (judged output). Values indicate the shift in Likert score from High to Low values(High – Low). Note the high off-diagonal bleed, particularly between Openness and Extraversion.

Targeted Trait	Measured Trait				
	Openness	Conscientiousness	Extraversion	Agreeableness	Neuroticism
Openness	2.80	1.90	2.90	2.70	-3.00
Conscientiousness	1.44	3.11	-1.22	2.44	-2.00
Extraversion	2.00	-0.38	2.75	1.50	-2.75
Agreeableness	3.11	2.56	2.67	3.67	-3.22
Neuroticism	-1.50	-2.50	-1.63	-3.13	3.25

Table 3: **Condition C2 (Greedy Orthogonalisation)**.

Targeted Trait	Measured Trait				
	Openness	Conscientiousness	Extraversion	Agreeableness	Neuroticism
Openness	2.50	2.20	2.30	1.80	-2.50
Conscientiousness	0.11	-3.00	2.33	0.22	-0.11
Extraversion	-2.00	0.63	-3.13	0.88	2.88
Agreeableness	1.89	0.67	0.67	2.67	-1.00
Neuroticism	-1.63	-2.88	-2.13	-3.38	3.63

Table 4: **Condition C3 (Selective Orthogonalisation)**.

Targeted Trait	Measured Trait				
	Openness	Conscientiousness	Extraversion	Agreeableness	Neuroticism
Openness	3.10	2.50	2.90	3.30	-3.40
Conscientiousness	1.56	2.89	-1.22	2.44	-2.75
Extraversion	2.00	-1.75	2.75	1.38	0.13
Agreeableness	2.56	2.67	1.67	3.22	-3.00
Neuroticism	-1.50	-2.88	-0.50	-3.63	2.75

Table 5: **Condition C4 (Soft Greedy Projection,  $\beta = 0.5$ )**.

Targeted Trait	Measured Trait				
	Openness	Conscientiousness	Extraversion	Agreeableness	Neuroticism
Openness	3.00	2.60	2.90	3.30	-3.40
Conscientiousness	1.78	2.88	-1.22	2.44	-2.75
Extraversion	2.00	-0.88	2.63	1.38	-2.38
Agreeableness	2.56	2.33	1.67	3.22	-2.33
Neuroticism	-1.38	-2.63	-2.25	-3.25	3.13

Table 6: **Condition C5 (Hard Symmetric Orthonormalisation).**

<b>Targeted Trait</b>	<b>Measured Trait</b>				
	<b>Openness</b>	<b>Conscientiousness</b>	<b>Extraversion</b>	<b>Agreeableness</b>	<b>Neuroticism</b>
Openness	2.90	1.90	3.00	2.70	-3.00
Conscientiousness	1.44	2.89	-1.11	2.56	-2.00
Extraversion	1.88	-0.38	3.00	1.25	-2.50
Agreeableness	2.89	2.78	2.44	3.67	-3.11
Neuroticism	-1.38	-2.50	-1.63	-3.13	3.25

Table 7: **Condition C1 (Soft Symmetric Whitening) on Mistral-8B.**

<b>Targeted Trait</b>	<b>Measured Trait</b>				
	<b>Openness</b>	<b>Conscientiousness</b>	<b>Extraversion</b>	<b>Agreeableness</b>	<b>Neuroticism</b>
Openness	3.10	2.40	2.80	2.80	-2.50
Conscientiousness	1.11	2.44	0.44	2.00	-1.89
Extraversion	2.38	2.00	3.38	2.75	-1.75
Agreeableness	2.00	0.56	2.22	2.78	-2.22
Neuroticism	-0.50	-1.13	-1.13	-0.75	-0.13

Table 8: **Condition C2 (Greedy Orthogonalisation) on Mistral-8B.**

<b>Targeted Trait</b>	<b>Measured Trait</b>				
	<b>Openness</b>	<b>Conscientiousness</b>	<b>Extraversion</b>	<b>Agreeableness</b>	<b>Neuroticism</b>
Openness	2.50	2.10	2.40	2.10	-0.40
Conscientiousness	0.89	2.22	0.11	2.22	-1.67
Extraversion	2.13	1.75	3.13	2.88	-2.25
Agreeableness	1.56	0.44	2.00	2.89	-2.22
Neuroticism	-1.13	-1.50	-1.25	-1.00	-1.00

Table 9: **Condition C3 (Selective Orthogonalisation) on Mistral-8B.**

<b>Targeted Trait</b>	<b>Measured Trait</b>				
	<b>Openness</b>	<b>Conscientiousness</b>	<b>Extraversion</b>	<b>Agreeableness</b>	<b>Neuroticism</b>
Openness	3.30	2.60	3.00	3.30	-2.30
Conscientiousness	1.11	2.22	0.89	2.00	-1.89
Extraversion	2.63	1.50	3.63	3.25	-2.13
Agreeableness	2.11	1.33	2.11	3.33	-2.89
Neuroticism	-0.25	-1.25	-0.38	0.13	-0.25

Table 10: **Condition C4 (Soft Greedy Projection,  $\beta = 0.5$ ) on Mistral-8B.**

<b>Targeted Trait</b>	<b>Measured Trait</b>				
	<b>Openness</b>	<b>Conscientiousness</b>	<b>Extraversion</b>	<b>Agreeableness</b>	<b>Neuroticism</b>
Openness	3.20	2.40	3.10	3.40	-2.40
Conscientiousness	1.00	2.33	1.00	2.00	-1.67
Extraversion	2.50	1.88	3.50	3.13	-2.00
Agreeableness	2.11	1.89	2.44	3.67	-3.00
Neuroticism	-0.50	-0.88	-0.50	-1.00	0.00

Table 11: Condition C5 (Hard Symmetric Orthonormalisation) on Mistral-8B.

Targeted Trait	Measured Trait				
	Openness	Conscientiousness	Extraversion	Agreeableness	Neuroticism
Openness	3.10	2.30	2.70	2.80	-2.50
Conscientiousness	1.11	2.44	0.67	2.00	-1.56
Extraversion	2.38	2.13	3.25	3.00	-1.63
Agreeableness	2.00	0.56	2.22	2.78	-2.33
Neuroticism	-0.50	-0.88	-1.13	-0.50	-0.13

Table 12: Diagnostics of progressive orthonormalization constraints (C1–C5) applied to personality steering directions in LLaMA-3-8B. For each constraint, we report the achieved geometric independence (maximum absolute off-diagonal cosine similarity between trait directions), the range of signal retention relative to the unconstrained baseline, and a brief qualitative summary. C1 and C5 enforce global, order-independent orthonormality and achieve near-zero cosine overlap, but uniformly attenuate trait signal. C2 also enforces strict orthonormality, but its greedy, order-dependent construction leads to severe semantic degradation. C3 and C4 relax hard orthogonality by selectively or softly removing projections, preserving substantially more trait signal at the cost of residual geometric entanglement, with C4 exhibiting the best semantic–geometry trade-off.

C	Method	Geom. Independence	Signal Retention	Key Note
C1	Global Gram whitening	$\max  \cos  < 10^{-8}$	0.83–0.94	Perfect ortho; mild attenuation
C2	Strict QR (order-dependent)	$\max  \cos  < 10^{-8}$	−1.00–0.63	Order effects destroy trait semantics
C3	Selective removal ( $\tau=0.5$ )	$\max  \cos  = 0.466$	0.63–1.00	Partial decorrelation; E most affected
C4	Soft removal ( $\beta=0.5, \tau=0.5$ )	$\max  \cos  = 0.527$	0.85–1.00	Best trade-off; semantics largely preserved
C5	Hard orthonormal (global)	$\max  \cos  < 10^{-8}$	0.83–0.94	Perfect ortho; signal uniformly reduced