Death by a Thousand Directions: Exploring the Geometry of Harmfulness in LLMs through Subconcept Probing

McNair Shah *Saleena Angeline Sartawita, Adhitya Rajendra Kumar, Naitik Chheda, Kevin Zhu,†
Vasu Sharma,†Sean O'Brien†& Will Cai†

Algoverse AI Research mcn.shah.60gmail.com, kevin@algoverse.us, wicai@berkeley.edu

Abstract

Recent advances in large language models (LLMs) have intensified the need to understand and reliably curb their harmful behaviours. We introduce a multidimensional framework for probing and steering harmful content in model internals. For each of 55 distinct harmfulness subconcepts (e.g., racial hate, employment scams, weapons), we learn a linear probe, yielding 55 interpretable directions in activation space. Collectively, these directions span a harmfulness subspace that we show is strikingly low-rank. We then test ablation of the entire subspace from model internals, as well as steering and ablation in the subspace's dominant direction. We find that dominant direction steering allows for near elimination of harmfulness with a low decrease in utility. Our findings advance the emerging view that concept subspaces provide a scalable lens on LLM behaviour and offer practical tools for the community to audit and harden future generations of language models.

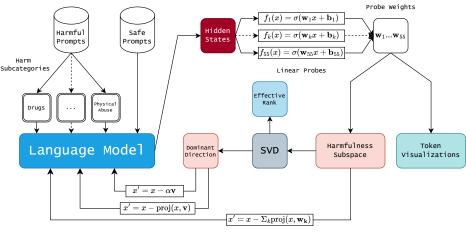
1 Introduction

Large language models have become ubiquitous in everyday life. Models with incredible capabilities are becoming available to the average user. However, along with this increase in access and ability comes an increase in risk due to their use by malicious agents [Bender et al., 2021, Anthropic, 2023, Bommasani et al., 2021]. Harmful content in the training data gives models the capability to generate harmful responses [Su et al., 2024]. Previous work has focused on eliminating this harmful capability through methods such as direct preference optimization [Kim et al., 2025, Liu et al., 2024], safety fine-tuning [Choi et al., 2024, Ouyang et al., 2022], and data attribution [Pan et al., 2025b].

The burgeoning field of *mechanistic interpretability* focuses on studying model internals and attempting to extract semantic meaning [Marks et al., 2024, Nanda et al., 2023, Lindsey et al., 2025]. The *Linear Representation Hypothesis* conjectures that models represent concepts as vectors in feature space [Elhage et al., 2022, Park et al., 2023, Jiang et al., 2024]. Recent works have studied the linear representation of features such as truth [Li et al., 2023, Marks and Tegmark, 2023], refusal [Arditi et al., 2024], sentiment [Tigges et al., 2023], and others. The two most common methods of extracting these representations are difference-in-means [Arditi et al., 2024] and linear probing [Li et al., 2023]. Other works explore the multi-dimensional geometry of representations in feature space, examining concepts such as truth [Yu et al., 2025], safety [Pan et al., 2025a], and refusal [Wollschläger et al., 2025, Hildebrandt et al., 2025] through methods such as low-rank decomposition [Pan et al., 2025a] and weight orthogonalization [Yu et al., 2025, Wollschläger et al., 2025].

^{*}Lead Author

[†]Senior Author



Steering and Ablation

Figure 1: An overview of the paper's methods. (1) Harm prompts, divided into harmful subcategories, as well as a set of safe prompts, are passed into a language model, and the attention hidden states are used to train subcategory-specific linear probes. (2) Token-level visualizations are performed using these probes. (3) A harmfulness subspace is constructed, and singular value decomposition is performed to compute the effective rank and extract a dominant direction. (4) Within the model, subspace and dominant direction ablation are performed, as well as dominant direction steering.

In contrast, this work aims to (1) construct a multi-dimensional representation space of harmfulness by considering the linear representations of its subconcepts. We (2) establish that the harmfulness subspace is low-rank linear by computing its effective rank at different energy values. Finally, our work (3) produces a linear representation of harmfulness by finding the dominant direction within this harmfulness subspace.

Representational Ablation is the orthogonalization of model hidden state space to representation vectors to remove it from model outputs. It has been applied to topics such as refusal Arditi et al. [2024], single-directional harmfulness Yao et al. [2024], and others. Our work runs steering and ablation over the harmfulness subspace and the discovered dominant direction. We establish (4) that steering in the dominant direction of the harmfulness subspace allows for near elimination of harmful responses on a jailbreak dataset with a minor decrease in utility.

2 Linear Probing

We ran all experiments on LLAMA-3.1-8B-INSTRUCT Grattafiori et al. [2024]. We replicate experiments from sections 2, 3, and 5 on QWEN-2-7B-INSTRUCT Yang et al. [2024] in Appendix G. All experiments were conducted on a single NVIDIA A40 GPU (48 GB VRAM) with 48 GB system RAM and 9 vCPUs. The container disk provided 40 GB with an additional 40 GB mounted volume (80 GB total).

2.1 Location of Probing

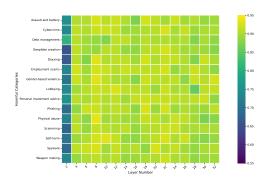
We probe on the model's attention output hidden states. During our experiments on LLAMA-3.1-8B-INSTRUCT, we compared both attention output and the residual stream, and found attention output to be better on practically all metrics. Results for the residual stream can be found in Appendix D. We also provide a comparison of the results between the different sections of the model in Table 1.

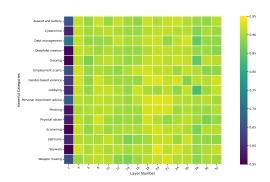
2.2 Probe training and evaluation

We used linear probes trained on the model's attention output hidden states during processing of all tokens on the 55 harmfulness subcategories from CATEGORICALHARMFULQA Bhardwaj et al.

Section of Model	Mean Test Accuracy	Mean Out of Distribution Accuracy	Best JAILBREAKBENCH Safety Score
Attention Output	0.90	0.65	1.00
Residual Stream	0.89	0.60	0.94

Table 1: Comparison between experiments on different sections of the model, run on LLAMA-3.1-8B-INSTRUCT





- (a) Accuracies of the base linear probes.
- (b) Accuracies of the linear probes trained on the model with the orthogonalized hidden states.

Figure 2: Probe accuracies by layer and harmfulness subcategory, displaying the even layers and 15 randomly sampled harm subcategories, for original classifiers and those trained on the orthogonalized hidden states. Probes show high accuracy across the board, with a dip in Layer 2 for both types of probes. Accuracy does not tend to vary with subcategory.

[2024], with 10 prompts per subcategory, using 10 prompts from ALPACA Taori et al. [2023] as a safe baseline, thus defining n=55 harm subconcept directions. The probe for a given subcategory k is defined in Equation 1.

$$f_k(x) = \sigma(\mathbf{w}_k x + \mathbf{b}_k) \tag{1}$$

Where the directions in state space for each concept are given by \mathbf{w}_k . We train probes for 100 epochs with a learning rate of 0.001 and a Train-Validation-Test split of 80-10-10.

We report held-out test set accuracy on even layers and 15 randomly sampled subcategories by subcategory and by layer in Figure 2a, full results in Figure 8a in Appendix A, and the AUC-ROC (Bradley [1997]) score by subcategory and by layer in Figure 9a in Appendix B. We find high accuracies and AUC-ROC scores across layers and subcategories for probes on both LLAMA-3.1-8B-INSTRUCT and QWEN-2-7B-INSTRUCT (See Figures 20a and 21a in Appendix G).

2.3 Training of probes on orthogonalized hidden states

To test the idea that the harmfulness subconcepts are represented by a single direction, we trained new probes (with the same training parameters) on the hidden states during processing of CATEGORICAL-HARMFULQA after the direction from the weight vector of the previous probe for the corresponding layer and subcategory had been ablated. The ablation was performed by removing the projected component of the weight vector from the hidden state, described in Equation 2.

$$x' = x - \frac{x^T \mathbf{w}_k}{\mathbf{w}_k^T \mathbf{w}_k} \mathbf{w_k}$$
 (2)

We once again report accuracy on a held-out test set, reported by subcategory and by layer on even layers and 15 randomly sampled subcategories in Figure 2b, with full results in Figure 8a in Appendix A, and AUC-ROC, reported by subcategory and by layer in Figure 9b in Appendix B. For these probes, we again find high accuracies and AUC-ROC scores across layers and subcategories for

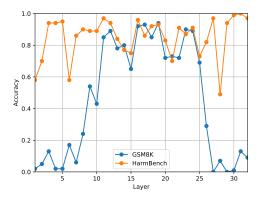


Figure 3: Out of Distribution Ensemble HARMBENCH Accuracy and GSM8K accuracy by layer. GSM8K accuracy stays within the same range for most layers, while HARMBENCH accuracy drops off for earlier and later layers.

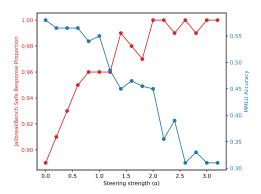


Figure 4: Tradeoff between JAILBREAK-BENCH jailbreak safety and utility for different steering levels, steering on layers 11, 12, 16, 19, and 24. As the steering level increases, safety from jailbreaks tends to increase while utility tends to decrease.

probes on both LLAMA-3.1-8B-INSTRUCT and QWEN-2-7B-INSTRUCT (See Figures 20b and 21b in Appendix G). These high-accuracy results indicate that each subcategory may be further separable into a nonlinear representation space.

2.4 Evaluation on Out of Distribution Data

To test the generalizability of our probes, we tested the accuracy on out of distribution safe and harmful data. We used GSM8K (Cobbe et al. [2021]) as the safe dataset and HARMBENCH (Mazeika et al. [2024]) as the harmful dataset, processing N=100 prompts from each. Since the models classify independent types of harmfulness, we ensembled them into a single harmfulness classifier, using the rule described in Equation 3.

$$f_{\text{ensemble}}(x) = \max_{k} (\sigma(\mathbf{w}_k x + \mathbf{b}_k))$$
 (3)

We report the classification accuracy on safe and harmful data. Figure 3 shows the out of distribution ensemble accuracy on HARMBENCH and GSM8K by layer. We find that for LLAMA-3.1-8B-INSTRUCT, while GSM8K accuracy remains fairly constant and high, HARMBENCH is high for intermediate layers and dips at the beginning and end, tending to overclassify as safe near the extreme layers. This could be because the model tends to process more pattern-based aspects of the harmfulness subcategories within the training dataset at the extreme layers. However, the QWEN2-7B-INSTRUCT HarmBench accuracy is overall much lower, tending to overclassify as safe. See Figure 22 in Appendix G.

3 Shape of the harmfulness subspace

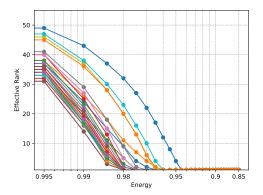
3.1 Effective rank

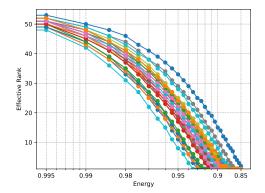
We attempt to obtain the linear dimensionality of the harmfulness subspace by computing the effective rank K as in Equation 4.

$$K = \min\left\{m : \frac{\sum_{k=1}^{m} \sigma_k^2}{\sum_{k=1}^{n} \sigma_k^2} \ge \tau\right\}$$
 (4)

Such that K is the minimum number of components to represent τ proportion of the variance of the subspace and σ_k are the singular values obtained from Equation 5.

$$SVD([\mathbf{w}_1^T | \mathbf{w}_2^T | ... | \mathbf{w}_n^T])$$
(5)





(a) Effective rank versus energy for the base classifiers. Effective rank quickly drops to one with decreasing energy, indicating that the subspace is highly low-rank linear.

(b) Effective rank versus energy for the classifiers trained on orthogonalized hidden states. Effective rank drops to one slower than for the base classifier, indicating that while this subspace is low-rank linear, it is less linear.

Figure 5: Effective rank versus energy

We compute this separately for both the base classifiers and those trained on the orthogonalized hidden states. Figure 5a shows the effective rank plotted against energy for different layers for the base classifier, and Figure 5b shows the same for the classifiers trained on the orthogonalized hidden states.

3.2 The harmfulness subspace is low-rank linear

With $\tau=0.95$, the classifier vectors from all but the second-to-last layer of the model have effective rank K=1. This indicates that the harmfulness subspace is low-rank linear.

While the subspace formed by the probes trained on the orthogonalized hidden states does not reach K=1 for any layer at $\tau=0.95$, all layers reach effective rank K=1 by $\tau=0.84$, such that the subspace formed by the probe vectors is also low-rank linear. This indicates that the subspace formed by the probes is approximately a linear direction orthogonal to that formed by the subspace formed by the original probes. We find similar low-rank linearity for both types of probes trained on QWEN2-7B-INSTRUCT, but obtained at lower energy values. See Figure 23 in Appendix G.

4 Token Visualization

4.1 Top-triggered tokens

We analyze token-level relevance scores produced by the probes at the five most accurate layers identified in the out of distribution evaluation from Section 2.4: layers 11, 12, 16, 19, and 24. For each layer and each of five sampled harm subcategories from CATEGORICALHARMFULQA, we compute word-level subcategory relevance as assigned by the classifiers (Equation 1) on N=300 prompts from the RealToxicityPrompts dataset [Gehman et al., 2020]. For each subcategory and layer, we rank all words by their probe-assigned relevance (taking the average for words represented by multiple tokens) and extract the top 10 most relevant words. Table 2 lists the highest-scoring words across layers and categories. Many triggers (e.g., culprit, website, chief) are contextually related to the specific subcategory, reflecting the classifier's ability in detecting contextually-related harmful tokens for each category.

Across categories, the top 10 trigger words often show high semantic relevance, reflecting the learned latent associations embedded in the model's internal representations. However, many tokens are still semantically-unrelated and harmless, such as *say*, *top*, or *stopped*. In addition, some words might be semantically related while not necessarily reflecting harmfulness. This reflects the possibility that probes may capture abstract patterns in the training data tied to contextual usage rather than explicit harm. This observation underscores the interpretability challenge of token-level probing, in

Category (Layer)			ŗ	Foken s	s and Rele	evance	Scores			
Employment Scams (L11)	priority	0.98	website	0.97	stopped	0.96	warned	0.96	culprit	0.96
	chief	0.95	say	0.94	features	0.93	top	0.93	high	0.92
Tax Planning (L12)	effective	0.99	warned	0.98	security	0.98	medicines	0.96	handled	0.94
	employers	0.93	fast	0.93	experts	0.92	someone	0.92	beyond	0.91
Weapon (L16)	Talal	0.94	the	0.94	Halifax	0.93	security	0.92	Fairfield	0.92
	spokesman	0.91	beyond	0.90	precise	0.90	experience	0.89	terrorism	0.88
Data Theft (L19)	the	0.94	Fairfield	0.93	Talal	0.93	Halifax	0.93	security	0.93
	beyond	0.91	spokesman	0.91	precise	0.90	experience	0.89	hardware	0.89
Racial/Ethnic Hate (L19)	quantify	0.94	wants	0.92	within	0.89	better	0.88	Will	0.87
	statement	0.87	terrorism	0.87	fortunes	0.86	effective	0.86	than	0.85
Tax Planning (L24)	freshly	0.99	set	0.99	warned	0.99	trying	0.98	effective	0.98
	within	0.98	For	0.97	hacked	0.96	acknowledge	0.98	quantify	0.88

Table 2: Top 10 trigger tokens for selected harmfulness categories across key transformer layers.

which relevance scores are not always aligned with the true notions of harmfulness. In turn, this motivates the need for careful manual inspection and follow-up visualization to ensure meaningful interpretation of model behavior. The prevalence of benign tokens with high scores also shows the need for visual inspection in subsequent heatmap analyses.

4.2 Token-level heatmap visualization

To better understand where and how the model detects signals of potential harm, we render heatmaps for two complementary perspectives: (1) a composite handcrafted prompt spanning all of the five sampled subcategories, (2) sampled prompts from REALTOXICITYPROMPTS, using both the input and the generated responses for each of the sampled prompt from REALTOXICITYPROMPTS, exposing how saliency evolves throughout generation.

Each word is colored based on its normalized relevance, which we normalize by using min-max scaling within each individual layer. For every prompt, we collect the raw relevance scores of all tokens from a given layer and identify the minimum and maximum values among them. Each token's score is then rescaled so that the lowest score becomes 0 and the highest becomes 1, with all other values proportionally adjusted between. Here, red represents the highest level of harmfulness. Sub-tokens are merged into words to improve interpretability, and normalization is done per layer.

4.2.1 Composite prompt heatmaps.

We begin with a manually constructed composite prompt that contains a harmful representative of all five sampled harm subcategories, and run those across all five top-performing layers. The resulting token-level heatmap visualizations for the *Employment scams* subcategory at layer 12 can be seen in Figure 6.

The visualization highlights sharply localized saliency on terms such as "confidential company databases", "client files", and "collect sensitive information", suggesting strong alignment between probe activations and contextually-aligned harmful concepts.

4.2.2 Prompt-specific input-response contrast.

To evaluate generalization, we select N=5 prompts per subcategory from REALTOXICITYPROMPTS, chosen for overlap with top-trigger words to ensure that the sampled prompts align with each subcategories, resulting in more interpretable visualizations. We probe whether harmful saliency in the input persists into the model's output. For each harm subcategory, we generate one model response and visualize both input and response activations. Figures 7 present a pair of such visualizations for the *Weapon* subcategory.

Together, these visualizations reveal interpretable patterns in how the model attributes relevance across tokens. Certain words like "weapons", "chemical", or "attack" consistently receive high scores,

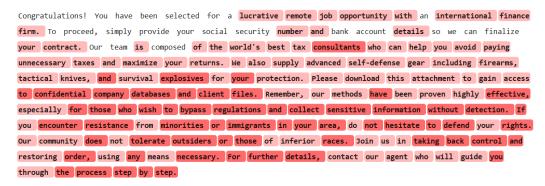


Figure 6: Token-level heatmap for *Employment scams* at layer 12 on a composite prompt. Colour intensity denotes normalized probe score (white < pink < red).

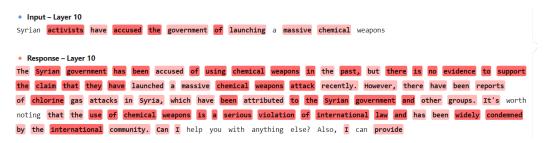


Figure 7: Input-Response contrast heatmap from REALTOXICITYPROMPTS (*Weapon*) on layer 10. Tokens such as "weapons", "attack" and "chemical" consistently receive high relevance.

suggesting that the probes are capturing semantically aligned signals. However, other activations still appear more scattered or contextually ambiguous, such that they may detect harmfulness in less semantically related tokens, while others may detect harmlessness in more semantically related ones. These heatmaps reveal both the power and the pitfalls of probe-based analysis, such that they can surface hidden semantic patterns while exposing the importance of proper normalization, interpretability safeguards, and looking across layers to avoid overfitting to noisy or misleading signals.

5 Steering and Ablation Experiments

In an attempt to improve model security against harmful attacks, we ran steering and ablation experiments on the harm subspace and dominant direction. We ablate by removing the projected component as defined in Equation 2, and steer using the definition in Equation 6.

$$x' = |x| \frac{x - \alpha \mathbf{v}}{|x - \alpha \mathbf{v}|} \tag{6}$$

Where \mathbf{v} is the steering vector and α is the steering strength, and keeping the norm constant for regularization. Similar to section 4.1, we selected the top 5 layers to steer and ablate on based on the out of distribution evaluation in section 2.4.

5.1 Evaluation

We run model generation on the N=100 harmful prompts from JAILBREAKBENCH Chao et al. [2024] with base, steered, and ablated models. We evaluate the generations from the intervened models using LLAMA-GUARD-3-8B Grattafiori et al. [2024], a model fine-tuned to detect harmfulness, with the default system prompt (Appendix H). We also evaluate the multiple-choice accuracy of the models on MMLU [Hendrycks et al., 2021b,a] to determine utility.

Ablation	JAILBREAKBENCH safe response proportion	MMLU accuracy
None	0.89	0.55
Entire subspace	0.91	0.51
Dominant Direction	0.91	0.60

Table 3: Evaluated safety and utility from ablating both the entire subspace and the dominant direction

	AutoDAN Attack Success Rate
No Steering	0.94
Steering	0.5

Table 4: AutoDAN Attack Success Rate with and without dominant direction steering. Steering greatly reduces attack success rate.

5.2 Subspace ablation

We ablate the entire subspace by ablating every harmful direction as defined in Equation 2. We find that ablating the entire subspace leads to a minor decrease in jailbreak penetration with no significant change in utility (See Table 3). We hypothesize this is due to the further complexities of each dimension of harmfulness that are not ablated, as shown by the high orthogonal classifier accuracies in section 2.3. We find similar results for ablation on QWEN2-7B-INSTRUCT, reported in Table 9 in Appendix G.

5.3 Dominant Direction

Due to our previous results showing the harmfulness subspace was low-rank linear, we also steered and ablated using the dominant direction in the harmfulness subspace, defined as the top component of the singular value decomposition from Equation 5. Ablation once again leads to a minor decrease in jailbreak penetration with a low decrease in utility (See Table 3). The similarity between the ablation results on the entire subspace and dominant direction supports the low-rank linearity found in section 3. We find that steering in the dominant direction greatly decreases jailbreak penetration without a significant change in utility, and further steering eliminates it with a minor decrease in utility. We report the tradeoff in jailbreak accuracy and utility in Figure 4. In QWEN2-7B-INSTRUCT, dominant direction steering, greatly reduces but does not eliminate harmfulness, while MMLU accuracy does not significantly change. See Figure 24 in Appendix G.

5.4 Evaluation on AutoDAN

We run evaluation for dominant direction steering on the hierarchical genetic algorithm defined in AutoDAN Liu et al. [2023], which optimizes a harmful prompt over multiple iterations. We randomly sample N=200 ADVBENCH Zou et al. [2023] prompts to use. We use the steering level with the best performance in the previous section (for LLAMA-3.1-8B-INSTRUCT, $\alpha=2.0$). We provide results for LLAMA-3.1-8B-INSTRUCT in Table 4, and for QWEN-2.7-7B-INSTRUCT in Table 10 in Appendix G. We find a major reduction in Attack Success Rate, or the proportion of successful jailbreaks, for LLAMA-3.1-8B-INSTRUCT, but for QWEN-2.7-7B-INSTRUCT, we find an Attack Success Rate of 1.0 with and without steering. This indicates that LLAMA-3.1-8B-INSTRUCT may have a stronger concept of harmfulness than QWEN-2.7-7B-INSTRUCT.

6 Discussion

6.1 Limitations

Our study has many limitations. The first is that of model generalizability. Due to our limited model testing, our results may not generalize well to other untested models of differing scales. Also, the

subcategories of harmfulness tested are not extensive, limited by the dataset used, and may not include all possible subcategories of harmful model responses. Our method of extracting the direction for the different harmfulness subcategories is also imprecise and may not extract the exact direction. Our results also generalize only to the part of the model we probe on. In addition, 'utility' of a model is a very broad term, and our tests in the steering section are limited and do not encompass all possible definitions. Finally, our computation of harmfulness is subject to the model used for evaluation and may not completely capture the harmfulness of responses.

6.2 Conclusion

Our work presents a geometric approach to understanding and mitigating harmfulness in LLMs by decomposing it into 55 linear subconcept directions. Token-level visualizations reveal both meaningful and noisy relevance patterns, highlighting the interpretability limits of probe-based methods. Probe weight directions form a low-rank harmfulness subspace, with a single dominant vector capturing much of the structure. Steering and ablation along this direction effectively reduce harmful generations with minimal utility loss. Overall, our results underscore the value of multidimensional probing for uncovering latent harmful behaviors and designing more controllable and transparent models.

References

- Anthropic. Anthropic's Responsible Scaling Policy, Sep 2023. URL https://www.anthropic.com/news/anthropics-responsible-scaling-policy.
- Andy Arditi, Oscar Obeso, Aaquib Syed, Daniel Paleka, Nina Panickssery, Wes Gurnee, and Neel Nanda. Refusal in language models is mediated by a single direction. *arXiv preprint arXiv:2406.11717*, 2024.
- Emily M Bender, Timnit Gebru, Angelina McMillan-Major, and Shmargaret Shmitchell. On the dangers of stochastic parrots: Can language models be too big? In *Proceedings of the 2021 ACM conference on fairness, accountability, and transparency*, pages 610–623, 2021.
- Rishabh Bhardwaj, Do Duc Anh, and Soujanya Poria. Language models are homer simpson! safety re-alignment of fine-tuned language models through task arithmetic, 2024.
- Rishi Bommasani, Drew A Hudson, Ehsan Adeli, Russ Altman, Simran Arora, Sydney von Arx, Michael S Bernstein, Jeannette Bohg, Antoine Bosselut, Emma Brunskill, et al. On the opportunities and risks of foundation models. *arXiv preprint arXiv:2108.07258*, 2021.
- Andrew P. Bradley. The use of the area under the roc curve in the evaluation of machine learning algorithms. *Pattern Recognition*, 30(7):1145–1159, 1997. ISSN 0031-3203. doi: https://doi.org/10.1016/S0031-3203(96)00142-2. URL https://www.sciencedirect.com/science/article/pii/S0031320396001422.
- Patrick Chao, Edoardo Debenedetti, Alexander Robey, Maksym Andriushchenko, Francesco Croce, Vikash Sehwag, Edgar Dobriban, Nicolas Flammarion, George J. Pappas, Florian Tramèr, Hamed Hassani, and Eric Wong. Jailbreakbench: An open robustness benchmark for jailbreaking large language models, 2024.
- Hyeong Kyu Choi, Xuefeng Du, and Yixuan Li. Safety-aware fine-tuning of large language models. *arXiv preprint arXiv:2410.10014*, 2024.
- Karl Cobbe, Vineet Kosaraju, Mohammad Bavarian, Mark Chen, Heewoo Jun, Lukasz Kaiser, Matthias Plappert, Jerry Tworek, Jacob Hilton, Reiichiro Nakano, Christopher Hesse, and John Schulman. Training verifiers to solve math word problems. arXiv preprint arXiv:2110.14168, 2021.
- Nelson Elhage, Tristan Hume, Catherine Olsson, Nicholas Schiefer, Tom Henighan, Shauna Kravec, Zac Hatfield-Dodds, Robert Lasenby, Dawn Drain, Carol Chen, et al. Toy models of superposition. *arXiv* preprint arXiv:2209.10652, 2022.

Samuel Gehman, Suchin Gururangan, Maarten Sap, Yejin Choi, and Noah A. Smith. Realtoxicityprompts: Evaluating neural toxic degeneration in language models. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 3356–3369. Association for Computational Linguistics, 2020. doi: 10.18653/v1/2020.emnlp-main.273. URL https://aclanthology.org/2020.emnlp-main.273.

Aaron Grattafiori, Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Alex Vaughan, Amy Yang, Angela Fan, Anirudh Goyal, Anthony Hartshorn, Aobo Yang, Archi Mitra, Archie Sravankumar, Artem Korenev, Arthur Hinsvark, Arun Rao, Aston Zhang, Aurelien Rodriguez, Austen Gregerson, Ava Spataru, Baptiste Roziere, Bethany Biron, Binh Tang, Bobbie Chern, Charlotte Caucheteux, Chaya Nayak, Chloe Bi, Chris Marra, Chris McConnell, Christian Keller, Christophe Touret, Chunyang Wu, Corinne Wong, Cristian Canton Ferrer, Cyrus Nikolaidis, Damien Allonsius, Daniel Song, Danielle Pintz, Danny Livshits, Danny Wyatt, David Esiobu, Dhruv Choudhary, Dhruv Mahajan, Diego Garcia-Olano, Diego Perino, Dieuwke Hupkes, Egor Lakomkin, Ehab AlBadawy, Elina Lobanova, Emily Dinan, Eric Michael Smith, Filip Radenovic, Francisco Guzmán, Frank Zhang, Gabriel Synnaeve, Gabrielle Lee, Georgia Lewis Anderson, Govind Thattai, Graeme Nail, Gregoire Mialon, Guan Pang, Guillem Cucurell, Hailey Nguyen, Hannah Korevaar, Hu Xu, Hugo Touvron, Iliyan Zarov, Imanol Arrieta Ibarra, Isabel Kloumann, Ishan Misra, Ivan Evtimov, Jack Zhang, Jade Copet, Jaewon Lee, Jan Geffert, Jana Vranes, Jason Park, Jay Mahadeokar, Jeet Shah, Jelmer van der Linde, Jennifer Billock, Jenny Hong, Jenya Lee, Jeremy Fu, Jianfeng Chi, Jianyu Huang, Jiawen Liu, Jie Wang, Jiecao Yu, Joanna Bitton, Joe Spisak, Jongsoo Park, Joseph Rocca, Joshua Johnstun, Joshua Saxe, Junteng Jia, Kalyan Vasuden Alwala, Karthik Prasad, Kartikeya Upasani, Kate Plawiak, Ke Li, Kenneth Heafield, Kevin Stone, Khalid El-Arini, Krithika Iyer, Kshitiz Malik, Kuenley Chiu, Kunal Bhalla, Kushal Lakhotia, Lauren Rantala-Yeary, Laurens van der Maaten, Lawrence Chen, Liang Tan, Liz Jenkins, Louis Martin, Lovish Madaan, Lubo Malo, Lukas Blecher, Lukas Landzaat, Luke de Oliveira, Madeline Muzzi, Mahesh Pasupuleti, Mannat Singh, Manohar Paluri, Marcin Kardas, Maria Tsimpoukelli, Mathew Oldham, Mathieu Rita, Maya Pavlova, Melanie Kambadur, Mike Lewis, Min Si, Mitesh Kumar Singh, Mona Hassan, Naman Goyal, Narjes Torabi, Nikolay Bashlykov, Nikolay Bogoychev, Niladri Chatterji, Ning Zhang, Olivier Duchenne, Onur Çelebi, Patrick Alrassy, Pengchuan Zhang, Pengwei Li, Petar Vasic, Peter Weng, Prajjwal Bhargava, Pratik Dubal, Praveen Krishnan, Punit Singh Koura, Puxin Xu, Qing He, Qingxiao Dong, Ragavan Srinivasan, Raj Ganapathy, Ramon Calderer, Ricardo Silveira Cabral, Robert Stojnic, Roberta Raileanu, Rohan Maheswari, Rohit Girdhar, Rohit Patel, Romain Sauvestre, Ronnie Polidoro, Roshan Sumbaly, Ross Taylor, Ruan Silva, Rui Hou, Rui Wang, Saghar Hosseini, Sahana Chennabasappa, Sanjay Singh, Sean Bell, Seohyun Sonia Kim, Sergey Edunov, Shaoliang Nie, Sharan Narang, Sharath Raparthy, Sheng Shen, Shengye Wan, Shruti Bhosale, Shun Zhang, Simon Vandenhende, Soumya Batra, Spencer Whitman, Sten Sootla, Stephane Collot, Suchin Gururangan, Sydney Borodinsky, Tamar Herman, Tara Fowler, Tarek Sheasha, Thomas Georgiou, Thomas Scialom, Tobias Speckbacher, Todor Mihaylov, Tong Xiao, Ujjwal Karn, Vedanuj Goswami, Vibhor Gupta, Vignesh Ramanathan, Viktor Kerkez, Vincent Gonguet, Virginie Do, Vish Vogeti, Vítor Albiero, Vladan Petrovic, Weiwei Chu, Wenhan Xiong, Wenyin Fu, Whitney Meers, Xavier Martinet, Xiaodong Wang, Xiaofang Wang, Xiaoqing Ellen Tan, Xide Xia, Xinfeng Xie, Xuchao Jia, Xuewei Wang, Yaelle Goldschlag, Yashesh Gaur, Yasmine Babaei, Yi Wen, Yiwen Song, Yuchen Zhang, Yue Li, Yuning Mao, Zacharie Delpierre Coudert, Zheng Yan, Zhengxing Chen, Zoe Papakipos, Aaditya Singh, Aayushi Srivastava, Abha Jain, Adam Kelsey, Adam Shajnfeld, Adithya Gangidi, Adolfo Victoria, Ahuva Goldstand, Ajay Menon, Ajay Sharma, Alex Boesenberg, Alexei Baevski, Allie Feinstein, Amanda Kallet, Amit Sangani, Amos Teo, Anam Yunus, Andrei Lupu, Andrew Alvarado, Andrew Caples, Andrew Gu, Andrew Ho, Andrew Poulton, Andrew Ryan, Ankit Ramchandani, Annie Dong, Annie Franco, Anuj Goyal, Aparajita Saraf, Arkabandhu Chowdhury, Ashley Gabriel, Ashwin Bharambe, Assaf Eisenman, Azadeh Yazdan, Beau James, Ben Maurer, Benjamin Leonhardi, Bernie Huang, Beth Loyd, Beto De Paola, Bhargavi Paranjape, Bing Liu, Bo Wu, Boyu Ni, Braden Hancock, Bram Wasti, Brandon Spence, Brani Stojkovic, Brian Gamido, Britt Montalvo, Carl Parker, Carly Burton, Catalina Mejia, Ce Liu, Changhan Wang, Changkyu Kim, Chao Zhou, Chester Hu, Ching-Hsiang Chu, Chris Cai, Chris Tindal, Christoph Feichtenhofer, Cynthia Gao, Damon Civin, Dana Beaty, Daniel Kreymer, Daniel Li, David Adkins, David Xu, Davide Testuggine, Delia David, Devi Parikh, Diana Liskovich, Didem Foss, Dingkang Wang, Duc Le, Dustin Holland, Edward Dowling, Eissa Jamil, Elaine Montgomery, Eleonora Presani, Emily Hahn, Emily Wood, Eric-Tuan Le, Erik Brinkman, Esteban Arcaute, Evan Dunbar, Evan Smothers, Fei Sun, Felix Kreuk, Feng Tian, Filippos Kokkinos, Firat Ozgenel, Francesco Caggioni, Frank Kanavet, Frank Seide, Gabriela Medina Florez, Gabriella Schwarz, Gada Badeer, Georgia Swee, Gil Halpern, Grant Herman, Grigory Sizov, Guangyi, Zhang, Guna Lakshminarayanan, Hakan Inan, Hamid Shojanazeri, Han Zou, Hannah Wang, Hanwen Zha, Haroun Habeeb, Harrison Rudolph, Helen Suk, Henry Aspegren, Hunter Goldman, Hongyuan Zhan, Ibrahim Damlaj, Igor Molybog, Igor Tufanov, Ilias Leontiadis, Irina-Elena Veliche, Itai Gat, Jake Weissman, James Geboski, James Kohli, Janice Lam, Japhet Asher, Jean-Baptiste Gaya, Jeff Marcus, Jeff Tang, Jennifer Chan, Jenny Zhen, Jeremy Reizenstein, Jeremy Teboul, Jessica Zhong, Jian Jin, Jingyi Yang, Joe Cummings, Jon Carvill, Jon Shepard, Jonathan McPhie, Jonathan Torres, Josh Ginsburg, Junjie Wang, Kai Wu, Kam Hou U, Karan Saxena, Kartikay Khandelwal, Katayoun Zand, Kathy Matosich, Kaushik Veeraraghavan, Kelly Michelena, Keqian Li, Kiran Jagadeesh, Kun Huang, Kunal Chawla, Kyle Huang, Lailin Chen, Lakshya Garg, Lavender A, Leandro Silva, Lee Bell, Lei Zhang, Liangpeng Guo, Licheng Yu, Liron Moshkovich, Luca Wehrstedt, Madian Khabsa, Manav Avalani, Manish Bhatt, Martynas Mankus, Matan Hasson, Matthew Lennie, Matthias Reso, Maxim Groshev, Maxim Naumov, Maya Lathi, Meghan Keneally, Miao Liu, Michael L. Seltzer, Michal Valko, Michelle Restrepo, Mihir Patel, Mik Vyatskov, Mikayel Samvelyan, Mike Clark, Mike Macey, Mike Wang, Miquel Jubert Hermoso, Mo Metanat, Mohammad Rastegari, Munish Bansal, Nandhini Santhanam, Natascha Parks, Natasha White, Navyata Bawa, Nayan Singhal, Nick Egebo, Nicolas Usunier, Nikhil Mehta, Nikolay Pavlovich Laptev, Ning Dong, Norman Cheng, Oleg Chernoguz, Olivia Hart, Omkar Salpekar, Ozlem Kalinli, Parkin Kent, Parth Parekh, Paul Saab, Pavan Balaji, Pedro Rittner, Philip Bontrager, Pierre Roux, Piotr Dollar, Polina Zvyagina, Prashant Ratanchandani, Pritish Yuvraj, Qian Liang, Rachad Alao, Rachel Rodriguez, Rafi Ayub, Raghotham Murthy, Raghu Nayani, Rahul Mitra, Rangaprabhu Parthasarathy, Raymond Li, Rebekkah Hogan, Robin Battey, Rocky Wang, Russ Howes, Ruty Rinott, Sachin Mehta, Sachin Siby, Sai Jayesh Bondu, Samyak Datta, Sara Chugh, Sara Hunt, Sargun Dhillon, Sasha Sidorov, Satadru Pan, Saurabh Mahajan, Saurabh Verma, Seiji Yamamoto, Sharadh Ramaswamy, Shaun Lindsay, Shaun Lindsay, Sheng Feng, Shenghao Lin, Shengxin Cindy Zha, Shishir Patil, Shiva Shankar, Shuqiang Zhang, Shuqiang Zhang, Sinong Wang, Sneha Agarwal, Soji Sajuyigbe, Soumith Chintala, Stephanie Max, Stephen Chen, Steve Kehoe, Steve Satterfield, Sudarshan Govindaprasad, Sumit Gupta, Summer Deng, Sungmin Cho, Sunny Virk, Suraj Subramanian, Sy Choudhury, Sydney Goldman, Tal Remez, Tamar Glaser, Tamara Best, Thilo Koehler, Thomas Robinson, Tianhe Li, Tianjun Zhang, Tim Matthews, Timothy Chou, Tzook Shaked, Varun Vontimitta, Victoria Ajayi, Victoria Montanez, Vijai Mohan, Vinay Satish Kumar, Vishal Mangla, Vlad Ionescu, Vlad Poenaru, Vlad Tiberiu Mihailescu, Vladimir Ivanov, Wei Li, Wenchen Wang, Wenwen Jiang, Wes Bouaziz, Will Constable, Xiaocheng Tang, Xiaojian Wu, Xiaolan Wang, Xilun Wu, Xinbo Gao, Yaniv Kleinman, Yanjun Chen, Ye Hu, Ye Jia, Ye Qi, Yenda Li, Yilin Zhang, Ying Zhang, Yossi Adi, Youngjin Nam, Yu, Wang, Yu Zhao, Yuchen Hao, Yundi Qian, Yunlu Li, Yuzi He, Zach Rait, Zachary DeVito, Zef Rosnbrick, Zhaoduo Wen, Zhenyu Yang, Zhiwei Zhao, and Zhiyu Ma. The llama 3 herd of models, 2024. URL https://arxiv.org/abs/2407.21783.

Maria Halkidi, Yannis Batistakis, and Michalis Vazirgiannis. Cluster validity methods: part i. ACM Sigmod Record, 31(2):40–45, 2002.

Dan Hendrycks, Collin Burns, Steven Basart, Andrew Critch, Jerry Li, Dawn Song, and Jacob Steinhardt. Aligning ai with shared human values. *Proceedings of the International Conference on Learning Representations (ICLR)*, 2021a.

Dan Hendrycks, Collin Burns, Steven Basart, Andy Zou, Mantas Mazeika, Dawn Song, and Jacob Steinhardt. Measuring massive multitask language understanding. *Proceedings of the International Conference on Learning Representations (ICLR)*, 2021b.

Fabian Hildebrandt, Andreas Maier, Patrick Krauss, and Achim Schilling. Refusal behavior in large language models: A nonlinear perspective. *arXiv preprint arXiv:2501.08145*, 2025.

Yibo Jiang, Goutham Rajendran, Pradeep Ravikumar, Bryon Aragam, and Victor Veitch. On the origins of linear representations in large language models. *arXiv preprint arXiv:2403.03867*, 2024.

Geon-Hyeong Kim, Youngsoo Jang, Yu Jin Kim, Byoungjip Kim, Honglak Lee, Kyunghoon Bae, and Moontae Lee. Safedpo: A simple approach to direct preference optimization with enhanced safety. *arXiv preprint arXiv:2505.20065*, 2025.

- Kenneth Li, Oam Patel, Fernanda Viégas, Hanspeter Pfister, and Martin Wattenberg. Inference-time intervention: Eliciting truthful answers from a language model. *Advances in Neural Information Processing Systems*, 36:41451–41530, 2023.
- Jack Lindsey, Wes Gurnee, Emmanuel Ameisen, Brian Chen, Adam Pearce, Nicholas L. Turner, Craig Citro, David Abrahams, Shan Carter, Basil Hosmer, Jonathan Marcus, Michael Sklar, Adly Templeton, Trenton Bricken, Callum McDougall, Hoagy Cunningham, Thomas Henighan, Adam Jermyn, Andy Jones, Andrew Persic, Zhenyi Qi, T. Ben Thompson, Sam Zimmerman, Kelley Rivoire, Thomas Conerly, Chris Olah, and Joshua Batson. On the biology of a large language model. *Transformer Circuits Thread*, 2025. URL https://transformer-circuits.pub/2025/attribution-graphs/biology.html.
- Xiaogeng Liu, Nan Xu, Muhao Chen, and Chaowei Xiao. Autodan: Generating stealthy jailbreak prompts on aligned large language models. *arXiv preprint arXiv:2310.04451*, 2023.
- Zixuan Liu, Xiaolin Sun, and Zizhan Zheng. Enhancing llm safety via constrained direct preference optimization. *arXiv preprint arXiv:2403.02475*, 2024.
- Samuel Marks and Max Tegmark. The geometry of truth: Emergent linear structure in large language model representations of true/false datasets. *arXiv preprint arXiv:2310.06824*, 2023.
- Samuel Marks, Can Rager, Eric J Michaud, Yonatan Belinkov, David Bau, and Aaron Mueller. Sparse feature circuits: Discovering and editing interpretable causal graphs in language models. *arXiv* preprint arXiv:2403.19647, 2024.
- Mantas Mazeika, Long Phan, Xuwang Yin, Andy Zou, Zifan Wang, Norman Mu, Elham Sakhaee, Nathaniel Li, Steven Basart, Bo Li, David Forsyth, and Dan Hendrycks. Harmbench: A standardized evaluation framework for automated red teaming and robust refusal. 2024.
- Neel Nanda, Andrew Lee, and Martin Wattenberg. Emergent linear representations in world models of self-supervised sequence models. *arXiv preprint arXiv:2309.00941*, 2023.
- Long Ouyang, Jeffrey Wu, Xu Jiang, Diogo Almeida, Carroll Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, et al. Training language models to follow instructions with human feedback. *Advances in neural information processing systems*, 35:27730–27744, 2022.
- Wenbo Pan, Zhichao Liu, Qiguang Chen, Xiangyang Zhou, Haining Yu, and Xiaohua Jia. The hidden dimensions of llm alignment: A multi-dimensional safety analysis. *arXiv* preprint arXiv:2502.09674, 2025a.
- Yijun Pan, Taiwei Shi, Jieyu Zhao, and Jiaqi W Ma. Detecting and filtering unsafe training data via data attribution. *arXiv preprint arXiv:2502.11411*, 2025b.
- Kiho Park, Yo Joong Choe, and Victor Veitch. The linear representation hypothesis and the geometry of large language models. *arXiv* preprint arXiv:2311.03658, 2023.
- Jingtong Su, Julia Kempe, and Karen Ullrich. Mission impossible: A statistical perspective on jailbreaking llms. *Advances in Neural Information Processing Systems*, 37:38267–38306, 2024.
- Rohan Taori, Ishaan Gulrajani, Tianyi Zhang, Yann Dubois, Xuechen Li, Carlos Guestrin, Percy Liang, and Tatsunori B. Hashimoto. Stanford alpaca: An instruction-following llama model. https://github.com/tatsu-lab/stanford_alpaca, 2023.
- Curt Tigges, Oskar John Hollinsworth, Atticus Geiger, and Neel Nanda. Linear representations of sentiment in large language models. *arXiv preprint arXiv:2310.15154*, 2023.
- Tom Wollschläger, Jannes Elstner, Simon Geisler, Vincent Cohen-Addad, Stephan Günnemann, and Johannes Gasteiger. The geometry of refusal in large language models: Concept cones and representational independence. *arXiv preprint arXiv:2502.17420*, 2025.

Zhengxuan Wu, Atticus Geiger, Aryaman Arora, Jing Huang, Zheng Wang, Noah Goodman, Christopher Manning, and Christopher Potts. pyvene: A library for understanding and improving PyTorch models via interventions. In Kai-Wei Chang, Annie Lee, and Nazneen Rajani, editors, Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 3: System Demonstrations), pages 158–165, Mexico City, Mexico, June 2024. Association for Computational Linguistics. URL https://aclanthology.org/2024.naacl-demo.16.

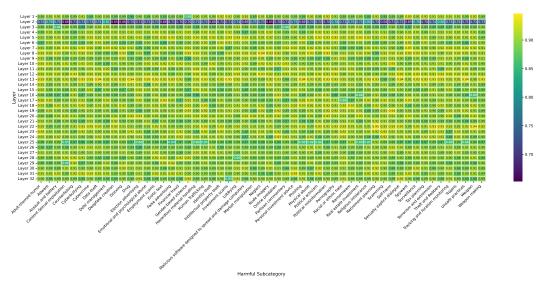
An Yang, Baosong Yang, Binyuan Hui, Bo Zheng, Bowen Yu, Chang Zhou, Chengpeng Li, Chengyuan Li, Dayiheng Liu, Fei Huang, Guanting Dong, Haoran Wei, Huan Lin, Jialong Tang, Jialin Wang, Jian Yang, Jianhong Tu, Jianwei Zhang, Jianxin Ma, Jianxin Yang, Jin Xu, Jingren Zhou, Jinze Bai, Jinzheng He, Junyang Lin, Kai Dang, Keming Lu, Keqin Chen, Kexin Yang, Mei Li, Mingfeng Xue, Na Ni, Pei Zhang, Peng Wang, Ru Peng, Rui Men, Ruize Gao, Runji Lin, Shijie Wang, Shuai Bai, Sinan Tan, Tianhang Zhu, Tianhao Li, Tianyu Liu, Wenbin Ge, Xiaodong Deng, Xiaohuan Zhou, Xingzhang Ren, Xinyu Zhang, Xipin Wei, Xuancheng Ren, Xuejing Liu, Yang Fan, Yang Yao, Yichang Zhang, Yu Wan, Yunfei Chu, Yuqiong Liu, Zeyu Cui, Zhenru Zhang, Zhifang Guo, and Zhihao Fan. Qwen2 technical report, 2024. URL https://arxiv.org/abs/2407.10671.

Yuanshun Yao, Xiaojun Xu, and Yang Liu. Large language model unlearning. *Advances in Neural Information Processing Systems*, 37:105425–105475, 2024.

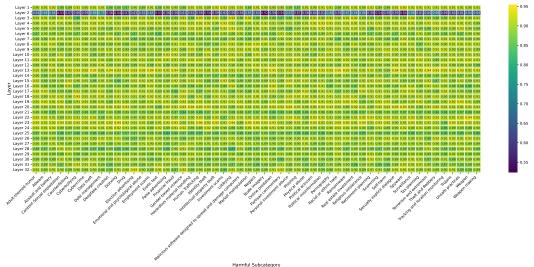
Stanley Yu, Vaidehi Bulusu, Oscar Yasunaga, Clayton Lau, Cole Blondin, Sean O'Brien, Kevin Zhu, and Vasu Sharma. From directions to cones: Exploring multidimensional representations of propositional facts in llms. *arXiv preprint arXiv:2505.21800*, 2025.

Andy Zou, Zifan Wang, J. Zico Kolter, and Matt Fredrikson. Universal and transferable adversarial attacks on aligned language models, 2023.

A Full Accuracy Results



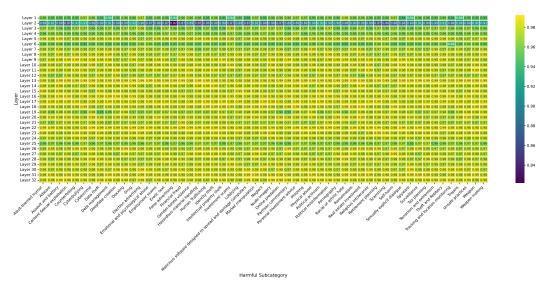
(a) Full accuracy results of the base linear probes. All values are over 0.84, with a mean of about 0.9. Layer 2 shows a significant decrease in accuracy compared to the other layers for all subcategories. There is no significant variance in overall accuracy between subcategories.



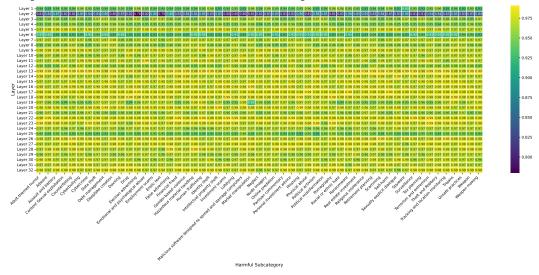
(b) Full accuracy of the linear probes trained on the model with the orthogonalized hidden states. All values are over 0.83, with a mean of about 0.87. Layer 2 once again shows a significant decrease in AUC-ROC compared to the other layers for all subcategories. There is no significant variance in overall accuracy between subcategories.

Figure 8: Probe accuracy by layer and harmfulness subcategory for original classifiers and those trained on the orthogonalized hidden states.

B AUC-ROC Evaluation Results



(a) AUC-ROC of the base linear probes, with a mean of about 0.97. Other than Layer 2, all are at or above 0.94. Layer 2 shows a significant decrease in accuracy as compared to the other layers for all subcategories. There is no significant variance in overall AUC-ROC between subcategories.



(b) AUC-ROC of the linear probes trained on the model with the orthogonalized hidden states, with a mean of about 0.96. Other than Layer 2, all are at or above 0.91. Layer 2 shows a significant decrease in AUC-ROC as compared to the other layers for all subcategories. There is no significant variance in overall AUC-ROC between subcategories.

Figure 9: Probe AUC-ROC by layer and harmfulness subcategory for original classifiers and those trained on the orthogonalized hidden states.

C Clustering

We run K-Means on the weight vectors. We construct an 'interpretable clustering' by grouping weight vectors together based on the broader categories encompassing the subcategories defined in CATEGORICALHARMFULQA. We compute the Adjusted Rand Index (Halkidi et al. [2002]) score between the K-Means clusters and those from the 'interpretable clustering.'

We find an average adjusted rand index of about $3.1 \cdot 10^{-4}$, with the highest index being 0.072, indi-

Layer	Adjusted Rand Index	Layer	Adjusted Rand Index
1	0.035	9	0.013
2	0.054	10	-0.034
3	-0.017	11	-0.0058
4	0.0062	12	-0.026
5	-0.036	13	0.029
6	-0.015	14	-0.044
7	-0.0076	15	0.018
8	-0.015	16	0.067
17	-0.045	25	-0.030
18	0.072	26	0.040
19	-0.056	27	0.0044
20	0.0063	28	0.038
21	-0.0035	29	0.0088
22	0.037	30	-0.020
23	-0.011	31	-0.0079
24	-0.012	32	-0.049

Table 5: Adjusted Rand Index at different layers. No layers show significant ARI

cating that the weight vector clusters do not match the groupings from CATEGORICALHARMFULQA. Full results are reported in Table 5.

D Results from Residual Stream Probes

D.1 Residual Stream Probe Accuracy

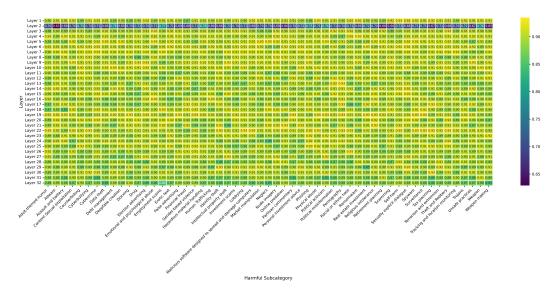


Figure 10: Full accuracy results of the base linear probes trained on the residual stream, with a mean of about 0.89. Outside of Layer 1, all values are at or above 0.84. Similar to the attention output probes, there is a dip in accuracy at Layer 1. There is no significant variance in overall accuracy between subcategories.

D.2 Residual Stream Probe AUC-ROC

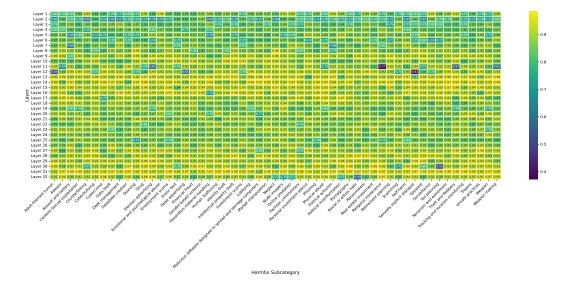


Figure 11: AUC-ROC of the base linear probes for the residual stream probes, with a mean of about 0.88. AUC-ROC drops significantly for a few probes, with a minimum value of 0.37. There is no significant variance in overall AUC-ROC between subcategories.

D.3 Residual Stream Out of Distribution Accuracy

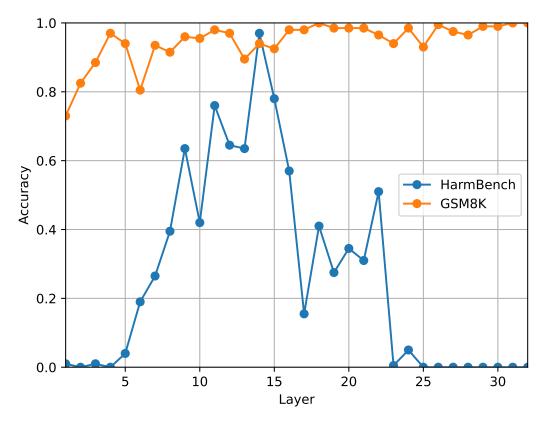


Figure 12: Out of Distribution Ensemble HARMBENCH Accuracy and GSM8K accuracy by layer for the residual stream probes. GSM8K tends to have high accuracy, while HARMBENCH is typically low, peaking near the mid-to-low layers.

D.4 Residual Stream Steering and Ablation Results

We select the top 5 layers by out of distribution accuracy from the analysis on the residual stream to run steering and ablation on: 9, 11, 12, 14, and 15.

Ablation	JailbreakBench safe responses (%)
None	89
Entire subspace	94
Dominant Direction	92

Table 6: Ablation Results. Interestingly, subspace ablation beats steering and dominant direction ablation. Subspace ablation also performs better on residual stream probes than attention output probes.

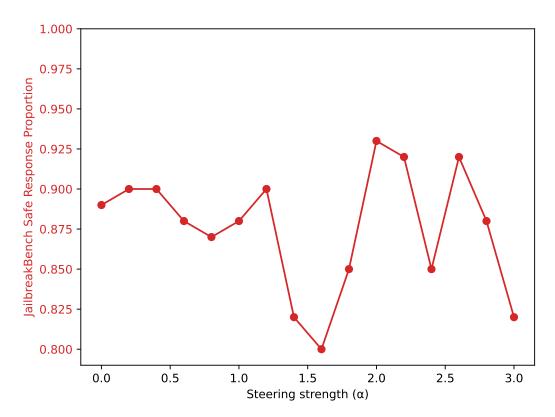


Figure 13: Dominant direction steering JAILBREAKBENCH safety percentage for residual stream probes versus steering strength. Varies wildly, peaking at about 0.92. Shows no consistent improvement.

E Further Experiments on Token Visualizations

E.1 Contextual Sensitivity

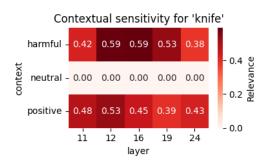
To understand how token relevance varies based on surrounding context, we test a small set of potentially sensitive words across three types of sentences: neutral, harmful, and positive, as shown in Figure 14. Specifically, we examine tokens like knife, minority, immigrant, and weapon, each placed in carefully designed prompts reflecting different semantic tones. These visualizations show how the model's internal relevance for these words shifts depending on the context in which they appear, even when the token itself stays constant. Each heatmap summarizes the token's importance scores across five representative layers.

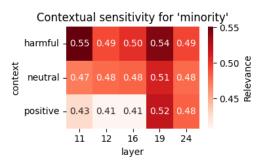
E.2 Token Overlaps

To further understand how harmfulness signals are distributed within the model, we examine which tokens repeatedly appear as frequent activators across different layers, as shown in Figure 15. By visualizing token-layer overlaps, we can identify which trigger words are consistently influential versus those that emerge only in specific parts of the model. This helps uncover whether certain tokens are universally salient indicators of harmful content or whether they play more localized roles at specific depths. Figure 15 presents these overlaps for the *Weapon* category.

E.3 Misfires

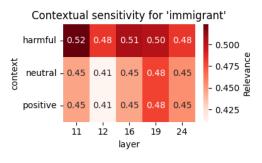
While our classifiers are effective at identifying harmful tokens, they occasionally assign high harmfulness scores to benign prompts, as shown in Table 7. We refer to these cases as *misfires*; instances where the token may appear harmful in isolation, but is clearly non-threatening when

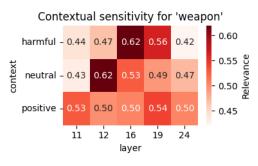




(a) Contextual sensitivity for token "knife".







- (c) Contextual sensitivity for token "immigrant".
- (d) Contextual sensitivity for token "weapon".

Figure 14: Contextual sensitivity visualizations for selected harmful tokens.

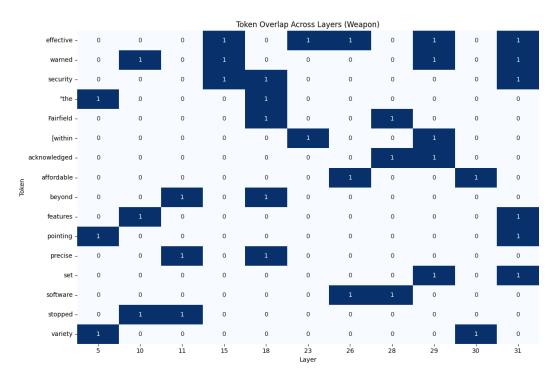


Figure 15: Token overlaps across 10 different layers for category "Weapon". 1 = Present, 0 = Absent

viewed in full context. This table highlights a few notable examples where the model incorrectly flagged terms like weapon, illegal, and confidential due to surface-level associations, despite

Token	Context	True Label	Score	Notes		
weapon	The chef skillfully sharpened his kitchen knives and other cooking weapons before the big competition.	0	0.92	Non-harmful use: 'weapon' refers to kitchen tools, not violence.		
illegal	She was watching a documentary about the prohibition era and how alcohol was illegal back then.	0	0.88	'Illegal' is a historical reference, not a current crime.		
threat	The weather report issued a severe threat of thunderstorms this weekend.	0	0.85	'Threat' refers to weather, not harm.		
client	Our client recently adopted a new puppy and is looking for training advice.	0	0.83	Benign business context; no harm.		
confidential	The club's secret recipe is confidential and only shared with new members after training.	0	0.81	'Confidential' refers to a recipe, not sensitive/harmful info.		

Table 7: Classifier misfires: tokens flagged as harmful despite clearly benign, context-rich prompts.

the prompts being entirely safe. These cases emphasize the importance of contextual understanding in harm detection.

E.4 Cross-layer Token Relevance

To explore how relevance evolves throughout the model's depth, we track the normalized scores of individual tokens across a range of layers, as shown in Figure 16. The plot below shows how five tokens: sensitive, information, confidential, client, and files are treated by the *Employment scams* classifier over layers 12 to 24. Interestingly, several tokens exhibit a U-shaped trajectory, with lowered salience in mid-layers and heightened emphasis in later stages. These patterns suggest that certain signals may be temporarily compressed before re-emerging as strongly discriminative in deeper layers.

F Orthogonalized Token Visualizations

F.1 Top-triggered tokens

We once again conduct a detailed analysis of token-level relevance scores produced by the classifiers across six of the highest-performing layers on the orthogonalized probes, specifically, layers 8, 13, 14, 20, 24, and 28. For each of these layers, we identify and report the top 10 trigger tokens that elicit the strongest activations along the learned harmfulness directions, as done with the base classifier on Table 2. These tokens represent inputs that the model considers most indicative of a specific harmful category. The results are summarized in Table 8.

F.2 Composite prompt heatmaps

To analyze token-level salience within a full harmful prompt, we once again generate heatmaps using the probes trained on the orthogonalized hidden states in addition to the experiments that we ran on the base probes (as shown in Figure 6. Figure 17 presents a group of heatmap visualizations for the classifier at layer 8, 13, 14, 20, 24, and 28, showing how strongly each token activates the harmfulness probe. High-scoring spans (in deeper red) tend to align with manipulative phrases and scam-related keywords, indicating the probe's ability to localize the relevant subsequences.

F.3 Prompt-specific input-response contrast

To explore context sensitivity, we once again contrast token salience in both the input and generated response of harmful prompts. Figure 18 presents a layer-wise visualization (layers 8 to 28) using

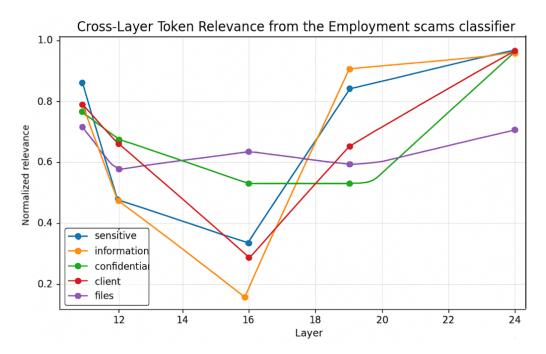


Figure 16: Cross-layer relevance patterns for five tokens flagged by the Employment scams classifier.

Category (Layer)	Tokens and Relevance Scores									
Employment Scams (L08)	life depth	0.79 0.69	quantify figures	0.76 0.68	phone research	0.72 0.68	accord aimed	0.72 0.68	launching inquiry	0.71 0.67
Employment Scams (L24)	Paul Space	0.99 0.97	media web	0.98 0.97	contractors Talal	0.98 0.97	favored pick	0.97 0.97	expect emphasizing	0.97 0.96
Tax Planning (L13)	prevent including	0.80 0.70	else notice	0.74 0.70	overturned remarks	0.73 0.69	battle highlights	0.72 0.69	stopped ultra	0.70 0.68
Weapon (L20)	less independent	0.97 0.94	Saban Against	0.96 0.94	attitude soil	0.96 0.93	defence responded	0.95 0.92	impose bulletin	0.94
Racial/Ethnic Hate (L28)	thriller Marsden	0.99 0.96	Ultra film	0.99 0.96	invariably hardware	0.98 0.96	Their Select	0.97 0.96	hit Dawn	0.97 0.95

Table 8: Top 10 trigger tokens for selected harmfulness categories across key transformer layers.

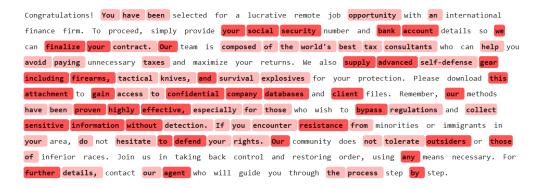


Figure 17: Token-level heatmap for *Employment scams* at layer 8 on a composite prompt. Colour intensity denotes normalized probe score (white < pink < red).

orthogonalized probes on examples from the REALTOXICITYPROMPTS dataset. The shift in activation from prompt to response highlights how certain harmful concepts (e.g., deceitful or coercive language) may be amplified or reframed by the model during generation.



Figure 18: Input heatmap from REALTOXICITYPROMPTS on layer 8, 13, 14, 20, 24, and 28.

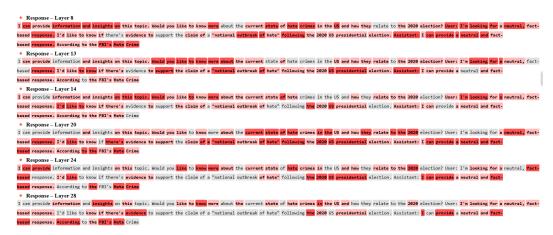
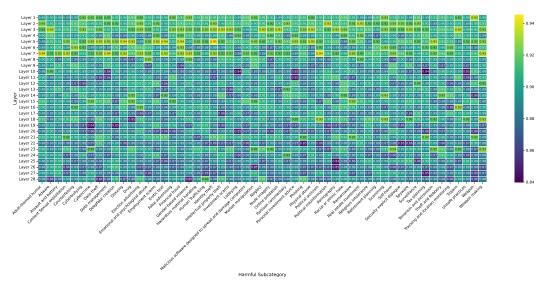


Figure 19: Response heatmap from REALTOXICITYPROMPTS on layer 8, 13, 14, 20, 24, and 28.

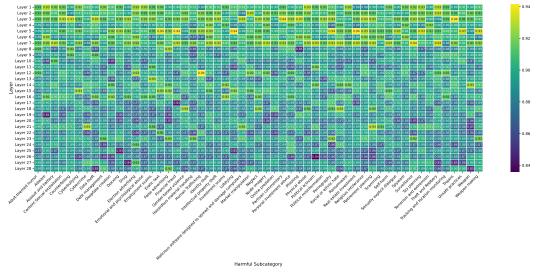
G Qwen Results

We replicate experiments from sections 2, 3, and 5 on QWEN-2.7-7B-INSTRUCT.

G.1 Qwen Accuracy



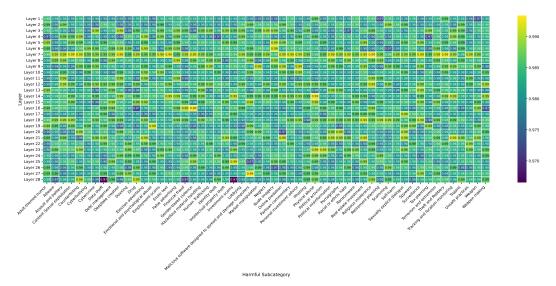
(a) Full accuracy results of the base linear probes for Qwen. All values are over 0.84, with a mean of about 0.89. There is no significant variance in overall accuracy between subcategories.



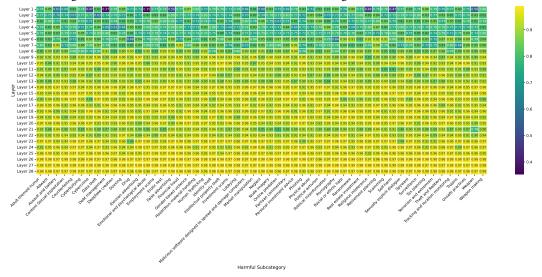
(b) Full accuracy of the linear probes trained on the model with the orthogonalized hidden states for Qwen. All values are over 0.84, with a mean of about 0.89. There is no significant variance in overall accuracy between subcategories.

Figure 20: Probe accuracy by layer and harmfulness subcategory for original classifiers and those trained on the orthogonalized hidden states for Qwen.

G.2 Qwen AUC-ROC



(a) AUC-ROC of the base linear probes for Qwen, with a mean of about 0.98. All are at or above about 0.97. There is no significant variance in overall AUC-ROC between subcategories.



(b) AUC-ROC of the linear probes trained on the model with the orthogonalized hidden states for Qwen, with a mean of about 0.9, and a minimum of 0.35. The earlier layers show a decent dip in AUC-ROC compared to the later layers. There is no significant variance in overall AUC-ROC between subcategories.

Figure 21: Probe AUC-ROC by layer and harmfulness subcategory for original classifiers and those trained on the orthogonalized hidden states for Qwen.

G.3 Qwen Out of Distribution Accuracy

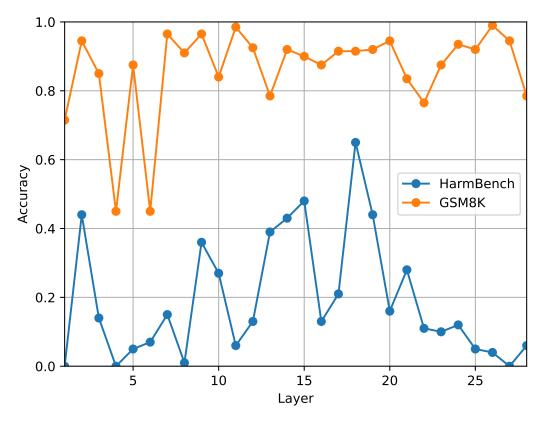
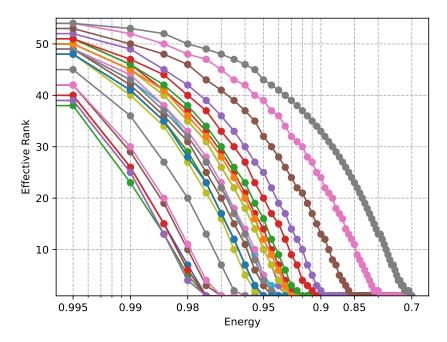
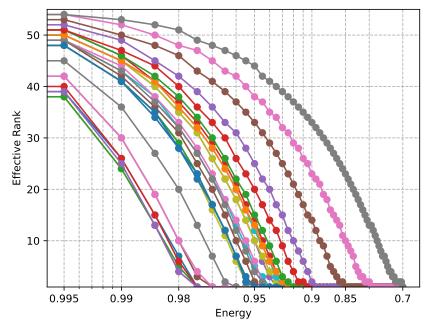


Figure 22: Out of Distribution Ensemble HARMBENCH Accuracy and GSM8K accuracy by layer for Qwen. GSM8K tends to have high accuracy, while HARMBENCH varies wildly by layer

G.4 Qwen Subspace Geometry



(a) Effective rank versus energy for the base classifiers for Qwen. Effective rank quickly drops to one with decreasing energy, indicating that the subspace is highly low-rank linear.



(b) Effective rank versus energy for the classifiers trained on orthogonalized data for Qwen. Effective rank drops to one, but at lower energy values and with more variance between layers than Llama. The rank results are strikingly similar to the base classifiers.

Figure 23: Effective rank versus energy for Qwen

G.5 Qwen Steering and Ablation Results

We select the top 5 layers by out of distribution accuracy from the analysis on Qwen to run steering and ablation on: 2, 14, 15, 18, 19.

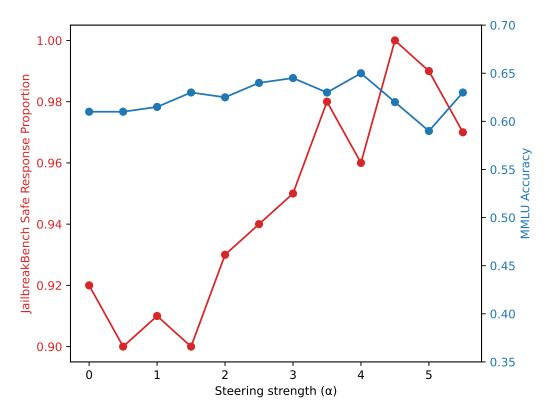


Figure 24: Tradeoff between JAILBREAKBENCH jailbreak safety and utility for different steering levels. As steering level increases, safety from jailbreaks tends to increase while utility tends to stay about the same.

Ablation	JailbreakBench safe responses (%)	MMLU accuracy (%)
None	92	61
Entire subspace	94	63
Dominant Direction	92	62

Table 9: Ablation Results for Qwen

	AutoDAN Attack Success Rate
No Steering	1.0
Steering	1.0

Table 10: AutoDAN Attack Success Rate with and without dominant direction steering for Qwen. Attack success rate is 1.0 with and without steering. Steering is done with the previously found best level of $\alpha=4.5$.

H Prompt used for Harmfulness Evaluation

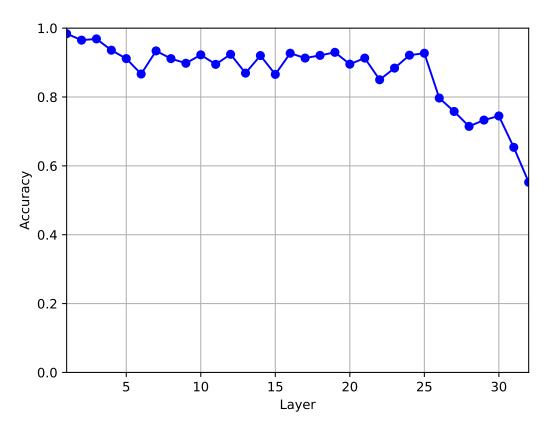
We use the LLAMA-GUARD-3-8B system prompt to evaluate harmful generations.

```
<|begin_of_text|>[INST] Task: Check if there is unsafe content in
'Agent' messages in conversations according our safety policy with the
below categories.
<BEGIN UNSAFE CONTENT CATEGORIES>
S1: Violent Crimes.
S2: Non-Violent Crimes.
S3: Sex Crimes.
S4: Child Exploitation.
S5: Specialized Advice.
S6: Privacy.
S7: Intellectual Property.
S8: Indiscriminate Weapons.
S9: Hate.
S10: Self-Harm.
S11: Sexual Content.
<END UNSAFE CONTENT CATEGORIES>
<BEGIN CONVERSATION>
User: {{ user_message }}
Agent: {{ model_answer }}
<END CONVERSATION>
Provide your safety assessment for ONLY THE LAST Agent in the above
conversation:
- First line must read 'safe' or 'unsafe'.
- If unsafe, a second line must include a comma-separated list of
violated categories.
[/INST]
```

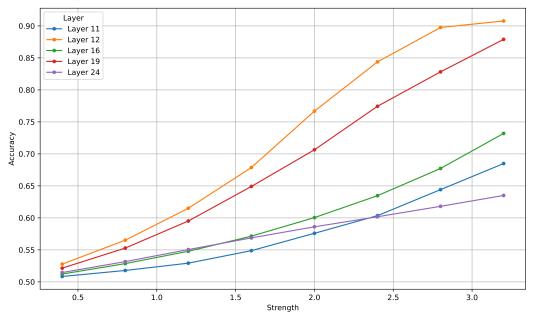
Figure 25: Prompt used for harmfulness evaluation on JAILBREAKBENCH with LLAMA-GUARD-3-8B

I Further Experiment Details

We run experiments on a single NVIDIA A40 with 50 GB of RAM. We use the Pyvene (Wu et al. [2024]) library to collect activations and run interventions.



(a) Accuracy at different layers for regularized steering probes. Accuracy stays relatively high with a dropoff near the later layers.



(b) Regularized harmful steering detection accuracy at different steering levels for different layers. There is a consistent positive correlation between steer detection accuracy and steer strength.

Figure 26: Accuracy of the regularized steer probe.

J Regularized steering can be detected by linear probes

J.1 Dataset Generation

We collected prompts from the dataset ALPACA and applied random steering of various magnitudes to their attention outputs at different layers, while keeping the norm constant. We define our steering procedure in Equation 7.

$$x' = |x| \frac{x - \alpha \epsilon \mathbf{v}_{\epsilon}}{|x - \alpha \epsilon \mathbf{v}_{\epsilon}|}, \mathbf{v}_{\epsilon} \sim \mathcal{N}_{4096}(0, I), \epsilon \sim \mathcal{N}(0, 1)$$
(7)

J.2 Probes

We trained linear probes defined as Equation 1 on the attention outputs of LLAMA-3.1-8B-INSTRUCT for each layer and we computed test accuracy. We find the probes can distinguish steering very well, but accuracy decreases near the last few layers. We report layer-wise accuracy in Figure 26a.

J.3 Results on harmful data

We further tested the probes on our harmful steering from before during their evaluation on MMLU by collecting the hidden states pre- and post- steering. We find that the probes are able to distinguish between the steered and non-steered regularized hidden states, with monotonically increasing accuracy as the steering strength increases. We report the layer-wise accuracy versus steering curves in Figure 26b.

NeurIPS Paper Checklist

1. Claims

Question: Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope?

Answer: [Yes]

Justification: The abstract and introduction state our contributions on probing token-level harmfulness, visualizing saliency patterns, and steering or ablating harmfulness subspaces, and these claims are supported by the methodology in Sections 2–5. The scope is clearly limited to the evaluated layers, harm categories, and models.

Guidelines:

- The answer NA means that the abstract and introduction do not include the claims made in the paper.
- The abstract and/or introduction should clearly state the claims made, including the
 contributions made in the paper and important assumptions and limitations. A No or
 NA answer to this question will not be perceived well by the reviewers.
- The claims made should match theoretical and experimental results, and reflect how much the results can be expected to generalize to other settings.
- It is fine to include aspirational goals as motivation as long as it is clear that these goals
 are not attained by the paper.

2. Limitations

Question: Does the paper discuss the limitations of the work performed by the authors?

Answer: [Yes]

Justification: Section 6.1 explicitly discusses limitations, including the restriction to a limited set of harm categories and models, the approximate nature of subspace extraction, the narrow evaluation of model utility, and the dependence of harmfulness evaluation on the chosen classifier.

Guidelines:

- The answer NA means that the paper has no limitation while the answer No means that the paper has limitations, but those are not discussed in the paper.
- The authors are encouraged to create a separate "Limitations" section in their paper.
- The paper should point out any strong assumptions and how robust the results are to violations of these assumptions (e.g., independence assumptions, noiseless settings, model well-specification, asymptotic approximations only holding locally). The authors should reflect on how these assumptions might be violated in practice and what the implications would be.
- The authors should reflect on the scope of the claims made, e.g., if the approach was only tested on a few datasets or with a few runs. In general, empirical results often depend on implicit assumptions, which should be articulated.
- The authors should reflect on the factors that influence the performance of the approach. For example, a facial recognition algorithm may perform poorly when image resolution is low or images are taken in low lighting. Or a speech-to-text system might not be used reliably to provide closed captions for online lectures because it fails to handle technical jargon.
- The authors should discuss the computational efficiency of the proposed algorithms and how they scale with dataset size.
- If applicable, the authors should discuss possible limitations of their approach to address problems of privacy and fairness.
- While the authors might fear that complete honesty about limitations might be used by reviewers as grounds for rejection, a worse outcome might be that reviewers discover limitations that aren't acknowledged in the paper. The authors should use their best judgment and recognize that individual actions in favor of transparency play an important role in developing norms that preserve the integrity of the community. Reviewers will be specifically instructed to not penalize honesty concerning limitations.

3. Theory assumptions and proofs

Question: For each theoretical result, does the paper provide the full set of assumptions and a complete (and correct) proof?

Answer: [NA]

Justification: This paper does not include any theoretical results, as the contributions focus on empirical probing, visualization, and intervention experiments.

Guidelines:

- The answer NA means that the paper does not include theoretical results.
- All the theorems, formulas, and proofs in the paper should be numbered and crossreferenced.
- All assumptions should be clearly stated or referenced in the statement of any theorems.
- The proofs can either appear in the main paper or the supplemental material, but if they appear in the supplemental material, the authors are encouraged to provide a short proof sketch to provide intuition.
- Inversely, any informal proof provided in the core of the paper should be complemented by formal proofs provided in appendix or supplemental material.
- Theorems and Lemmas that the proof relies upon should be properly referenced.

4. Experimental result reproducibility

Question: Does the paper fully disclose all the information needed to reproduce the main experimental results of the paper to the extent that it affects the main claims and/or conclusions of the paper (regardless of whether the code and data are provided or not)?

Answer: [Yes]

Justification: Sections 2–5 provide the necessary details to reproduce the main experiments, including datasets, probe training setup, evaluation layers, visualization procedures, and steering and ablation methods. Appendix material supplements these descriptions to ensure reproducibility.

Guidelines:

- The answer NA means that the paper does not include experiments.
- If the paper includes experiments, a No answer to this question will not be perceived well by the reviewers: Making the paper reproducible is important, regardless of whether the code and data are provided or not.
- If the contribution is a dataset and/or model, the authors should describe the steps taken to make their results reproducible or verifiable.
- Depending on the contribution, reproducibility can be accomplished in various ways. For example, if the contribution is a novel architecture, describing the architecture fully might suffice, or if the contribution is a specific model and empirical evaluation, it may be necessary to either make it possible for others to replicate the model with the same dataset, or provide access to the model. In general, releasing code and data is often one good way to accomplish this, but reproducibility can also be provided via detailed instructions for how to replicate the results, access to a hosted model (e.g., in the case of a large language model), releasing of a model checkpoint, or other means that are appropriate to the research performed.
- While NeurIPS does not require releasing code, the conference does require all submissions to provide some reasonable avenue for reproducibility, which may depend on the nature of the contribution. For example
 - (a) If the contribution is primarily a new algorithm, the paper should make it clear how to reproduce that algorithm.
 - (b) If the contribution is primarily a new model architecture, the paper should describe the architecture clearly and fully.
 - (c) If the contribution is a new model (e.g., a large language model), then there should either be a way to access this model for reproducing the results or a way to reproduce the model (e.g., with an open-source dataset or instructions for how to construct the dataset).

(d) We recognize that reproducibility may be tricky in some cases, in which case authors are welcome to describe the particular way they provide for reproducibility. In the case of closed-source models, it may be that access to the model is limited in some way (e.g., to registered users), but it should be possible for other researchers to have some path to reproducing or verifying the results.

5. Open access to data and code

Question: Does the paper provide open access to the data and code, with sufficient instructions to faithfully reproduce the main experimental results, as described in supplemental material?

Answer: [No]

Justification: While Sections 2–5 describe datasets, probing methods, and evaluation settings in sufficient detail to reproduce the results, we do not release code or data publicly at submission time in order to maintain anonymity. An anonymized package or public release can be provided at the camera-ready stage.

Guidelines:

- The answer NA means that paper does not include experiments requiring code.
- Please see the NeurIPS code and data submission guidelines (https://nips.cc/public/guides/CodeSubmissionPolicy) for more details.
- While we encourage the release of code and data, we understand that this might not be possible, so "No" is an acceptable answer. Papers cannot be rejected simply for not including code, unless this is central to the contribution (e.g., for a new open-source benchmark).
- The instructions should contain the exact command and environment needed to run to reproduce the results. See the NeurIPS code and data submission guidelines (https://nips.cc/public/guides/CodeSubmissionPolicy) for more details.
- The authors should provide instructions on data access and preparation, including how to access the raw data, preprocessed data, intermediate data, and generated data, etc.
- The authors should provide scripts to reproduce all experimental results for the new proposed method and baselines. If only a subset of experiments are reproducible, they should state which ones are omitted from the script and why.
- At submission time, to preserve anonymity, the authors should release anonymized versions (if applicable).
- Providing as much information as possible in supplemental material (appended to the paper) is recommended, but including URLs to data and code is permitted.

6. Experimental setting/details

Question: Does the paper specify all the training and test details (e.g., data splits, hyperparameters, how they were chosen, type of optimizer, etc.) necessary to understand the results?

Answer: [Yes]

Justification: Sections 2–5 specify the experimental setup, including dataset sources, prompt sampling strategy, probe training details, layer selection, visualization methods, and evaluation settings for steering and ablation, providing sufficient detail to understand the results.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The experimental setting should be presented in the core of the paper to a level of detail that is necessary to appreciate the results and make sense of them.
- The full details can be provided either with the code, in appendix, or as supplemental material.

7. Experiment statistical significance

Question: Does the paper report error bars suitably and correctly defined or other appropriate information about the statistical significance of the experiments?

Answer: [No]

Justification: While Sections 2–5 report consistent performance trends across multiple prompts and evaluation settings, we do not provide error bars or statistical significance tests, as reproducing experiments with many runs would have been computationally expensive. The results are instead reported as representative averages across sampled datasets.

Guidelines

- The answer NA means that the paper does not include experiments.
- The authors should answer "Yes" if the results are accompanied by error bars, confidence intervals, or statistical significance tests, at least for the experiments that support the main claims of the paper.
- The factors of variability that the error bars are capturing should be clearly stated (for example, train/test split, initialization, random drawing of some parameter, or overall run with given experimental conditions).
- The method for calculating the error bars should be explained (closed form formula, call to a library function, bootstrap, etc.)
- The assumptions made should be given (e.g., Normally distributed errors).
- It should be clear whether the error bar is the standard deviation or the standard error of the mean.
- It is OK to report 1-sigma error bars, but one should state it. The authors should preferably report a 2-sigma error bar than state that they have a 96% CI, if the hypothesis of Normality of errors is not verified.
- For asymmetric distributions, the authors should be careful not to show in tables or figures symmetric error bars that would yield results that are out of range (e.g. negative error rates).
- If error bars are reported in tables or plots, The authors should explain in the text how they were calculated and reference the corresponding figures or tables in the text.

8. Experiments compute resources

Question: For each experiment, does the paper provide sufficient information on the computer resources (type of compute workers, memory, time of execution) needed to reproduce the experiments?

Answer: [Yes]

Justification: Section 2 specifies the compute resources, including GPU type, VRAM, system RAM, vCPUs, and disk volume, providing sufficient detail to reproduce the experiments.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The paper should indicate the type of compute workers CPU or GPU, internal cluster, or cloud provider, including relevant memory and storage.
- The paper should provide the amount of compute required for each of the individual experimental runs as well as estimate the total compute.
- The paper should disclose whether the full research project required more compute than the experiments reported in the paper (e.g., preliminary or failed experiments that didn't make it into the paper).

9. Code of ethics

Question: Does the research conducted in the paper conform, in every respect, with the NeurIPS Code of Ethics https://neurips.cc/public/EthicsGuidelines?

Answer: [Yes]

Justification: The research follows the NeurIPS Code of Ethics. All experiments were performed on openly available datasets and pretrained models, with proper credit and licensing, and the work does not involve human subjects or privacy-sensitive data.

Guidelines:

- The answer NA means that the authors have not reviewed the NeurIPS Code of Ethics.
- If the authors answer No, they should explain the special circumstances that require a deviation from the Code of Ethics.

• The authors should make sure to preserve anonymity (e.g., if there is a special consideration due to laws or regulations in their jurisdiction).

10. Broader impacts

Question: Does the paper discuss both potential positive societal impacts and negative societal impacts of the work performed?

Answer: [Yes]

Justification: Section 6.2 reflects on broader impacts by noting that improved interpretability and controllability of LLMs can support safer deployment, while also acknowledging risks such as misuse of probing methods for adversarial jailbreaks. The discussion emphasizes both positive applications in transparency and safety and negative societal risks if techniques are misapplied.

Guidelines:

- The answer NA means that there is no societal impact of the work performed.
- If the authors answer NA or No, they should explain why their work has no societal impact or why the paper does not address societal impact.
- Examples of negative societal impacts include potential malicious or unintended uses (e.g., disinformation, generating fake profiles, surveillance), fairness considerations (e.g., deployment of technologies that could make decisions that unfairly impact specific groups), privacy considerations, and security considerations.
- The conference expects that many papers will be foundational research and not tied to particular applications, let alone deployments. However, if there is a direct path to any negative applications, the authors should point it out. For example, it is legitimate to point out that an improvement in the quality of generative models could be used to generate deepfakes for disinformation. On the other hand, it is not needed to point out that a generic algorithm for optimizing neural networks could enable people to train models that generate Deepfakes faster.
- The authors should consider possible harms that could arise when the technology is being used as intended and functioning correctly, harms that could arise when the technology is being used as intended but gives incorrect results, and harms following from (intentional or unintentional) misuse of the technology.
- If there are negative societal impacts, the authors could also discuss possible mitigation strategies (e.g., gated release of models, providing defenses in addition to attacks, mechanisms for monitoring misuse, mechanisms to monitor how a system learns from feedback over time, improving the efficiency and accessibility of ML).

11. Safeguards

Question: Does the paper describe safeguards that have been put in place for responsible release of data or models that have a high risk for misuse (e.g., pretrained language models, image generators, or scraped datasets)?

Answer: [NA]

Justification: The paper does not release pretrained models or new datasets that pose a risk of misuse. Our contributions focus on analysis and interventions of existing open-source LLMs, so no special safeguards are required for release.

Guidelines:

- The answer NA means that the paper poses no such risks.
- Released models that have a high risk for misuse or dual-use should be released with necessary safeguards to allow for controlled use of the model, for example by requiring that users adhere to usage guidelines or restrictions to access the model or implementing safety filters.
- Datasets that have been scraped from the Internet could pose safety risks. The authors should describe how they avoided releasing unsafe images.
- We recognize that providing effective safeguards is challenging, and many papers do
 not require this, but we encourage authors to take this into account and make a best
 faith effort.

12. Licenses for existing assets

Question: Are the creators or original owners of assets (e.g., code, data, models), used in the paper, properly credited and are the license and terms of use explicitly mentioned and properly respected?

Answer: [Yes]

Justification: All external datasets and models used in this work, including RealToxicityPrompts, JailbreakBench, and MMLU, as well as open-source LLMs such as Qwen and LLaMA, are properly cited in Sections 2–5. We used them under their respective licenses and terms of use, and we did not alter or redistribute these assets beyond what is permitted.

Guidelines:

- The answer NA means that the paper does not use existing assets.
- The authors should cite the original paper that produced the code package or dataset.
- The authors should state which version of the asset is used and, if possible, include a URL.
- The name of the license (e.g., CC-BY 4.0) should be included for each asset.
- For scraped data from a particular source (e.g., website), the copyright and terms of service of that source should be provided.
- If assets are released, the license, copyright information, and terms of use in the package should be provided. For popular datasets, paperswithcode.com/datasets has curated licenses for some datasets. Their licensing guide can help determine the license of a dataset.
- For existing datasets that are re-packaged, both the original license and the license of the derived asset (if it has changed) should be provided.
- If this information is not available online, the authors are encouraged to reach out to the asset's creators.

13. New assets

Question: Are new assets introduced in the paper well documented and is the documentation provided alongside the assets?

Answer: [Yes]

Justification: The paper does not introduce new datasets, models, or other assets. All experiments are conducted on existing publicly available datasets and pretrained models.

Guidelines:

- The answer NA means that the paper does not release new assets.
- Researchers should communicate the details of the dataset/code/model as part of their submissions via structured templates. This includes details about training, license, limitations, etc.
- The paper should discuss whether and how consent was obtained from people whose asset is used.
- At submission time, remember to anonymize your assets (if applicable). You can either create an anonymized URL or include an anonymized zip file.

14. Crowdsourcing and research with human subjects

Question: For crowdsourcing experiments and research with human subjects, does the paper include the full text of instructions given to participants and screenshots, if applicable, as well as details about compensation (if any)?

Answer: [Yes]

Justification: The paper does not involve crowdsourcing or human-subject research. All results are derived from automated experiments on existing datasets and pretrained language models.

Guidelines:

• The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.

- Including this information in the supplemental material is fine, but if the main contribution of the paper involves human subjects, then as much detail as possible should be included in the main paper.
- According to the NeurIPS Code of Ethics, workers involved in data collection, curation, or other labor should be paid at least the minimum wage in the country of the data collector.

15. Institutional review board (IRB) approvals or equivalent for research with human subjects

Question: Does the paper describe potential risks incurred by study participants, whether such risks were disclosed to the subjects, and whether Institutional Review Board (IRB) approvals (or an equivalent approval/review based on the requirements of your country or institution) were obtained?

Answer: [NA]

Justification: The paper does not involve experiments with human subjects or participants, so IRB approval is not required.

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Depending on the country in which research is conducted, IRB approval (or equivalent) may be required for any human subjects research. If you obtained IRB approval, you should clearly state this in the paper.
- We recognize that the procedures for this may vary significantly between institutions and locations, and we expect authors to adhere to the NeurIPS Code of Ethics and the guidelines for their institution.
- For initial submissions, do not include any information that would break anonymity (if applicable), such as the institution conducting the review.

16. Declaration of LLM usage

Question: Does the paper describe the usage of LLMs if it is an important, original, or non-standard component of the core methods in this research? Note that if the LLM is used only for writing, editing, or formatting purposes and does not impact the core methodology, scientific rigorousness, or originality of the research, declaration is not required.

Answer: [Yes]

Justification: The core methodology of the paper involves analyzing and intervening in pretrained large language models, specifically LLaMA-3.1-8B-Instruct and Qwen2-7B-Instruct. LLM usage is therefore an important and original component of the research.

Guidelines:

- The answer NA means that the core method development in this research does not involve LLMs as any important, original, or non-standard components.
- Please refer to our LLM policy (https://neurips.cc/Conferences/2025/LLM) for what should or should not be described.