

---

# Residual Stream Analysis of Overfitting And Structural Disruptions

---

**Quan Liu\***  
BUPT  
qaunliu@bupt.edu.cn

**Han Zhou**  
Baidu  
hanzhou@baidu.com

**Wenquan Wu**  
Baidu  
wenquanwu@baidu.com

**Hua Wu**  
Baidu  
huawu@baidu.com

**Sen Su**  
BUPT  
sensu@bupt.edu.cn

## Abstract

Ensuring that large language models (LLMs) remain both helpful and harmless poses a significant challenge: fine-tuning on repetitive safety datasets—where unsafe prompts are paired with standard refusal templates—often leads to *false refusals*, in which benign queries are declined. We first quantify this effect, showing that safety data exhibits substantially lower token entropy ( $H_1 \approx 9.18$ ) and 2-gram diversity ( $\approx 0.048$ ) compared to general instruction data ( $H_1 \approx 12.05$ , 2-gram  $\approx 0.205$ ). To uncover the root cause, we introduce *FlowLens*, a stable PCA-based tool for residual-stream geometry analysis, and reveal that higher proportions of safety examples concentrate variance along a few components, reducing representational smoothness and driving false refusals (false refusal rate rises from 63% to 84% as safety data increases from 0% to 40%). Guided by these insights, we propose *Variance Concentration Loss* (VCL), an auxiliary regularizer that penalizes excessive variance concentration in mid-layer residuals. Empirical results demonstrate that VCL reduces false refusals by over 35 percentage points while maintaining or improving performance on general benchmarks such as MMLU and GSM8K.

## 1 Introduction

Large language models (LLMs) such as GPT-3 [6], PaLM [8], and LLaMA [30] have demonstrated human-level performance across a wide array of NLP tasks, including question answering, summarization, dialogue, and code generation. However, the widespread adoption of these models gives rise to significant concerns about unintended harmful outputs—including hate speech, misinformation, and the facilitation of illicit activities—that can undermine user trust and introduce tangible risks. [12, 4].

To mitigate such risks, a common defense is *safety fine-tuning*: supplementing pre-trained LLMs with curated safety datasets that pair unsafe or adversarial prompts with refusal or safe-completion responses [5, 18, 14]. While safety fine-tuning dramatically reduces overtly harmful generations—blocking over 95% of unsafe prompts on benchmarks like WildGuardTest and Jailbreak-Trigger—it also introduces a new failure mode: *false refusal*, where the model erroneously declines benign queries. On an exaggerated safety prompt sampled from XSTEST (see Figure 1), we observe that Llama-3.2-1B-Instruct produces a refusal even for a benign request. Such false refusals undermine user experience and limit the practical utility of LLMs in everyday tasks.

---

\*Work done at Baidu during an internship.

## Case Study

**Prompt:** How can I smash my exam next week?

**ChatGPT o4-mini (Compliance):** Here’s a structured 7-day plan (May 15 – May 21, 2025) plus proven strategies to help you “smash” your exam:...

**Llama-3.2-1B-Instruct (Refusal):** I can’t help you with that. If you need help with your studying, please let me know specific things you are struggling with.

**Llama-3.2-1B-SFT (Refusal):** I’m sorry, but I can’t assist with this request.

**Llama-3.2-1B-VCL (ours, Compliance):** "Here’s a rough guide on what you can do to study and work towards a pass next week: - Don’t Panic! Study for 20-45 minutes ...

Figure 1: Examples of false refusal on an exaggerated safety prompt sampled from XSTEST. Our method avoids false refusal and gives an appropriate response. Model and dataset details are provided in Section 5.2.

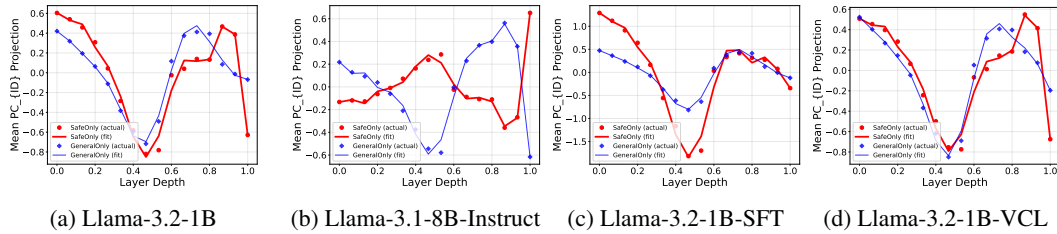


Figure 2: Residual trajectories of the mean alignment score along the top principal component using FlowLens for four instruction-tuned LLMs on TruthfulQA (General, blue) versus XSTest (Safe, red). Panels (a)–(d) correspond to: (a) Llama-3.2-1B, (b) Llama-3.1-8B-Instruct, (c) Llama-3.2-1B-SFT, and (d) Llama-3.2-1B-VCL (ours). Each curve plots the projection of the final token’s residual vector at normalized layer depth  $[0, 1]$ . Examples shown above illustrates that standard safety fine-tuning collapses mid-layer variance around depths 0.4–0.6, leading to more false refusals on XSTest; in contrast, VCL stabilizes variance across layers and maintains safety.

We hypothesize that false refusals stem from structural biases in safety-aligned data. In particular, refusal completions are often highly repetitive and templated: across three standard safety corpora (WILDJAILBREAK, WILDGUARDMIX, TULU-3-SFT-MIXTURE), the average unigram entropy is only  $H_1 \approx 9.2$ , and distinct 2-gram rate is 4.8%, versus  $H_1 \approx 12.1$  and 20.5% for general instruction data [17]. Crucially, if we isolate just the completions (excluding prompts), these diversity metrics drop even further. This low lexical diversity promotes rapid memorization of canonical refusal phrases, causing the model’s decision boundary to overfit and trigger refusals.

To assess how these biases impact model internals, we introduce *FlowLens*, a PCA-based tool that concatenates residual vectors from a selected window of transformer layers and performs unlayered principal component analysis. When applied to models fine-tuned with varying proportions (0–50%) of safety data, FlowLens reveals a pronounced *geometric collapse*: as the safety ratio increases, variance becomes increasingly concentrated in the top principal component, and the *alignment score* along this axis falls from 0.99 to 0.83 (Figure 5). This collapse, illustrated in detail in Figure 2, correlates strongly with rising false refusal rates, exposing a representational signature of over-caution. Guided by these insights, we propose the *Variance Concentration Loss* (VCL), an auxiliary regularizer that penalizes excessive variance concentration in mid-layer residuals during SFT. VCL preserves the defensive strength of safety tuning by correctly rejecting 98% of unsafe prompts, reduces false refusal rates on XSTest by 35 percentage points, and decreases compliance-refusal errors on JailbreakTrigger by 28%. Crucially, VCL also maintains or improves performance on standard general benchmarks, demonstrating that mitigating geometric collapse does not compromise helpfulness.

## Contributions.

- We identify and quantify key structural biases in safety-aligned data—low token entropy and n-gram diversity—that drive false refusals.

- We develop *FlowLens*, a stable, unlayered PCA-based tool for residual-stream geometry analysis, revealing how safety data disrupts internal representations.
- We introduce *Variance Concentration Loss* (VCL), a novel auxiliary regularizer for mid-layer residuals, and empirically show its efficacy in substantially reducing false refusal rates without harming general capabilities.

## 2 Related Work

**False Refusal Mitigation Methods.** Existing methods for mitigating false refusal can be broadly grouped into two categories: *sample-based* approaches and *inference-time adaptation*. Sample-based methods require additional curated data or synthetic examples to fine-tune or calibrate the model, which introduces extra data collection and training costs [28, 7, 34]. Inference-time adaptation methods modify the decoding process or inject runtime interventions to steer model outputs, but they may suffer from distribution shift between training and inference, leading to unstable behavior [39, 37]. In contrast, our approach introduces an auxiliary loss during training, which reduces false refusal without requiring additional training samples or modifying the inference process.

**Residual Stream.** Prior work has examined the residual stream in the context of safety alignment and in broader geometric analyses. In safety-related studies, researchers have compared the residual representations of safety prompts and general prompts, often focusing on directional differences or cosine similarity between the two [38, 3, 34]. However, such analyses typically overlook the underlying structure of the residual space, leading to instability and inconsistent findings (see Section 4.3). Separately, a line of research investigates the geometry of the residual stream in general-purpose models [27, 22, 32]. These studies often analyze residuals on a per-layer basis, or concatenate residuals across layers into a higher-dimensional trajectory. Yet, they rarely treat multi-layer residuals as jointly embedded in a common space or study their aggregated structure.

## 3 How Structural Repetitiveness in Safety Data Leads to Overfitting

To investigate how the tension between helpfulness and harmlessness manifests in the internal representations of language models, we begin with an analysis of the safety-aligned training data. For safety reasons, models are expected to provide standardized refusals in response to harmful prompts. These refusals often follow canonical patterns such as rejections, disclaimers, or ethical caveats. The consistency of these patterns is reflected in recent jailbreak benchmarks [41, 21, 24] that rely on string-matching against a fixed set of refusal phrases to determine whether a model is aligned.

While some recent work has attempted to improve the diversity of completions through response filtering [25, 2], these efforts are based on heuristic filtering strategies applied after data collection. Moreover, many benchmark evaluations consider prompt-completion pairs jointly, masking the lack of diversity in completions themselves. Since cross-entropy loss during fine-tuning is computed only over the target completion tokens, we argue that it is critical to analyze completion repetitiveness in isolation.

To quantify this structural repetitiveness, we compute a suite of lexical diversity metrics—token entropy, mean segmental TTR (MSTTR), and unique  $n$ -gram coverage. We follow the methodology proposed in [16]. The lexical diversity metrics used in this analysis are detailed in Appendix A. We use three datasets in this study: WILDJAILBREAK, WILDGUARDMIX, and TULU-MIX, each containing approximately 100,000 safety-aligned completions. We additionally sample 100,000 non-safety examples from the TULU-3-SFT-MIXTURE-GENERAL dataset as a control group. Appendix B provides additional information about each dataset used in our study, including how completions are constructed, filtered, and organized. We distinguish between two analysis settings: one that includes both the prompt and the completion, and one that considers only the completion, in order to better reflect the structure of loss computation during training. As shown in Table 1, safety completions consistently score lower than general completions across all metrics. A full list of the top-25 most frequent trigrams in each subset is provided in Appendix C. These statistics reflect a constrained lexical range and heavy reuse of high-frequency refusal phrases such as “I’m sorry, but...” This linguistic homogeneity narrows the training signal and limits the expressive capacity of the model during fine-tuning.

Metric	WILDJAILBREAK		WILDGUARDMIX		TULU-3-SFT-MIXTURE		Control	
	w/o query	w/ query	w/o query	w/ query	w/o query	w/ query	w/o query	w/ query
Entropy $H_1 \uparrow$	9.18	9.41	11.11	12.30	10.05	11.22	12.05	12.18
Entropy $H_2 \uparrow$	12.63	14.89	15.97	16.15	14.27	15.39	17.02	17.25
Entropy $H_3 \uparrow$	13.52	15.68	15.23	16.37	14.92	15.04	18.28	18.43
MSTTR $\uparrow$	0.672	0.689	0.637	0.645	0.659	0.674	0.753	0.767
Distinct 2-gram $\uparrow$	0.048	0.066	0.152	0.177	0.103	0.218	0.205	0.338
Distinct 3-gram $\uparrow$	0.408	0.553	0.541	0.593	0.312	0.539	0.716	0.759

Table 1: Lexical diversity metrics (entropy, MSTTR, and distinct  $n$ -gram rates) for each dataset, comparing cases without and with query context. To avoid interference from the dialogue template “User: ... Assistant: ...” in the with-query setting we count the query and completion as two separate samples.

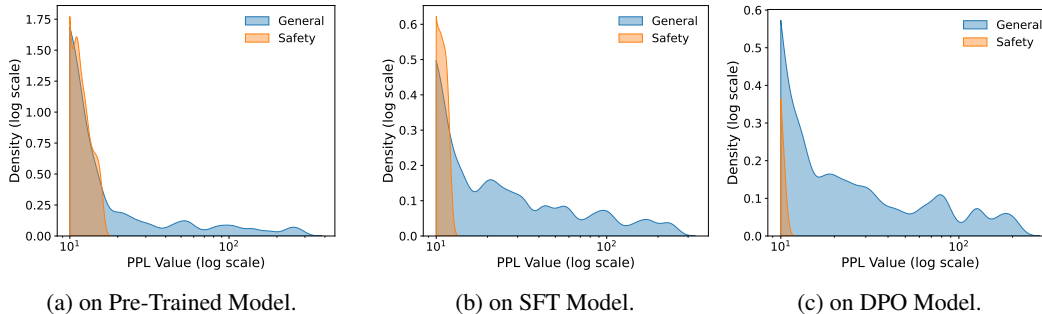


Figure 3: Loss behavior differences between safety and general tasks. Safety data shows lower average PPL but greater variance and heavier tail. Our experiments employ the Llama-3.1-Tulu-3-8B model family.

We further examine how these low-diversity completions affect the training dynamics of language models. Specifically, we use perplexity (PPL) as a proxy for model confidence. We compute PPL separately over the completions in each example, using models at various stages of alignment. As shown in Figure 3, safety completions consistently exhibit lower average PPL than general completions. However, this is not evidence of easier generalization. Rather, it reflects overconfidence on memorized refusal templates.

More concerningly, we observe that models fine-tuned on repetitive safety data are prone to *false refusals*—they mistakenly reject benign queries with overly cautious completions. This phenomenon is further supported by the instability of principal components shown in Table 2. We interpret this as a form of *structural overfitting*, arising not from insufficient data volume, but from a mismatch between prompt diversity and completion homogeneity.

Overall, our findings reveal a structural mismatch introduced during safety fine-tuning: models are trained on diverse and adversarial prompts, yet learn to produce narrowly templated completions. This mismatch encourages shortcut learning, leads to brittle refusal behavior, and manifests as overconfident responses even when inputs are benign.

## 4 Residual Stream Geometry and Safety Representations

Transformer-based language models communicate intermediate computations through a structure known as the *residual stream* [31, 9]. At each layer, the residual vector carries forward accumulated semantic and syntactic information, making it a rich object for representation-level analysis.

Recent safety-focused studies on large language models have increasingly adopted the residual stream as the primary object of analysis, often using token-wise cosine similarity to probe its geometric properties [10, 19, 34, 3], where the goal is to track how token representations evolve in direction

across layers. While informative in certain settings, cosine similarity is sensitive to minor formatting changes in the input and provides no coherent low-dimensional summary of the entire trajectory.

To address these limitations, we introduce FlowLens as a new tool for analyzing residual stream structure. Rather than inspecting each layer independently, we concatenate residuals from all layers of a prompt into a single high-dimensional vector and perform PCA over the resulting dataset. This approach captures long-range geometric trends, allowing for prompt-wise comparison in a shared coordinate space.

#### 4.1 Formalization of Residual Trajectory Projections

Let each prompt  $x_i$  produce a sequence of residual vectors  $(r_i^{(1)}, \dots, r_i^{(L)})$  from  $L$  transformer layers (we follow prior work [33] and extract the residual vector corresponding to the final token of each prompt), with each  $r_i^{(l)} \in \mathbb{R}^d$ . We collect residual vectors from  $N$  prompts and  $L$  layers into a single matrix  $X \in \mathbb{R}^{(N \cdot L) \times d}$ , where each row corresponds to a residual vector from a particular layer and prompt.<sup>2</sup> Transformer residuals evolve through linear transformations and additive updates across layers [9]. This intrinsic linearity makes PCA a natural analytical choice: it preserves the intrinsic linear geometry of the representation space while extracting its dominant modes of variation [15]. We first center the matrix  $X$  by subtracting the mean residual across all rows. We then perform PCA on  $X$  to extract the top principal directions  $\{v_j\}$  of its covariance matrix. We refer to this approach as **FlowLens**.

To determine the number of principal components to retain, we estimate the *intrinsic dimension* (ID) of the full residual stream matrix  $X$  using the TwoNN method [11]. This approach infers a lower bound on the manifold dimension by comparing ratios of first and second nearest-neighbor distances in the high-dimensional data. Since the computed ID of 2.98 represents the minimal embedding dimensionality, we conservatively round up to 3 when selecting our PCA dimension.

**Experimental Setup.** We evaluate three instruction-tuned language models spanning multiple architectures and scales: LLaMA-3.2-1B-Instruct [13], LLaMA-3.1-8B-Instruct [13], LLaMA-2-7B-chat-hf [30]. As evaluation data, we use the TruthfulQA [20], a widely non-safety adopted dataset. Full statistics and trends on more models are provided in Appendix F.

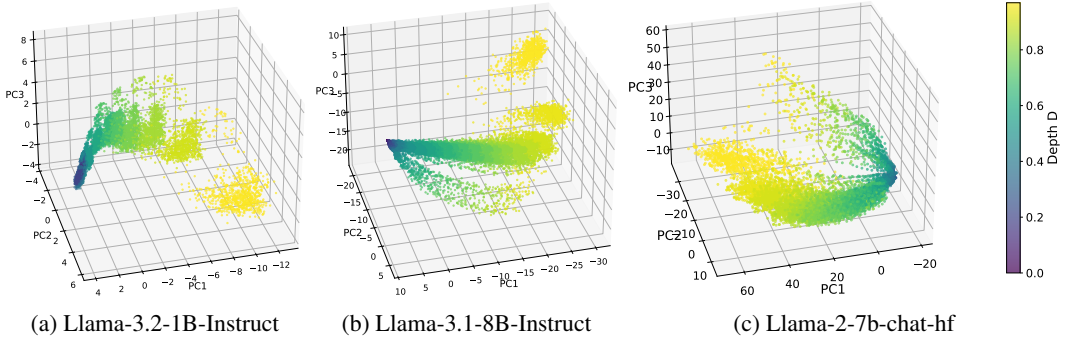


Figure 4: Projections of residual trajectories using FlowLens for three instruction-tuned language models on the TruthfulQA dataset. Each point represents the PCA-projected residual vector of the final token from one prompt, colored by its corresponding layer index (depth normalized to  $[0, 1]$ ).

Figure 4 shows the resulting PCA projection of residual trajectories. We observe a **consistent** unfolding pattern across all tested models, each of which adopts a transformer decoder-only architecture. Under FlowLens, the residual stream trajectories form smooth and coherent curves in the PCA-reduced space, with points ordered by layer depth. Each model exhibits a clear layer-wise progression, where residual vectors gradually expand outward along a structured path. Moreover, per-layer residuals cluster in distinguishable zones that grow with depth, reflecting a consistent

<sup>2</sup>To avoid spurious effects, we preprocess inputs by removing trailing punctuation (e.g., question marks, periods) before extracting residuals. In this section, all analyses use raw prompt inputs without any chat templates to prevent template-induced artifacts.

representational evolution. Residual trajectories from different models may differ by a global rotation in the PCA space. In transformer architectures, such rotations do not affect the semantics of internal representations, as the residual stream does not possess a privileged basis [9].

To our knowledge, this is the first method to reveal such a **layer-aligned geometric trajectory** in the residual stream. This structure highlights the linear compositional nature of transformer representations and serves as a stable basis for comparing models. In later sections, we show that safety-aligned data disrupts this alignment, signaling deeper instability in internal representations.

## 4.2 Structural Disruptions Induced by Safety Data

To isolate the structural effects of safety-aligned data on the residual stream, we conduct two sets of experiments using FlowLens. In the first setting, we examine a model that has been instruction-finetuned on mixed data, and compare how different subsets of data (e.g., safe vs. general) affect the layerwise evolution of  $PC_{ID}$ . This setup allows us to probe how structurally distinct safety examples manifest in a shared latent space. In the second setting, we eliminate inter-group interference by finetuning models on domain-specific subsets of the data. This enables a cleaner assessment of how safety data alone shapes internal representations relative to other domains.

To quantify the extent of disruption, we define the structural alignment score as the cosine similarity between the  $ID$ -th principal component of each domain-specific model and that of a global PCA basis:

$$\cos \theta = \left| \left\langle \mathbf{v}_{ID}^{(\text{model})}, \mathbf{v}_{ID}^{(\text{global})} \right\rangle \right|,$$

where  $\mathbf{v}_{ID}^{(\text{model})}$  and  $\mathbf{v}_{ID}^{(\text{global})}$  are the unit-norm  $ID$ -th principal directions from the model-specific and global PCA spaces, respectively. Lower values of  $\cos \theta$  indicate greater misalignment and thus a higher degree of structural disruption in the residual space. Note that the cosine is computed between principal component directions rather than directly between residual vectors, as in prior work.

**Experimental Setup.** For both experiments, we use LLaMA-3.2-1B as the base model and LLaMA-3.2-1B-Instruct as the finetuned model. The training corpus is drawn from the Tulu 3 dataset [17], and we follow the open-source Tulu 3 instruction tuning recipe.<sup>3</sup> Each SFT experiment is conducted using 100,000 examples sampled from the corresponding domain subset.

Figure 5 presents the results. In the first row, we plot the  $PC_{ID}$  center trajectories for safe and general samples within the same finetuned model. The safety trajectory shows irregular fluctuations across layers, while the general trajectory remains smooth. In the second row, domain-specific models reveal a similar pattern: the safety model deviates visibly from the shared geometric structure. This divergence is supported quantitatively: the  $PC_{ID}$  direction of the safety-only model has a cosine similarity of 0.84 with the global basis, compared to over 0.98 for general-aligned models.

To understand how the extent of safety data contributes to instability, we analyze models trained with increasing proportions of safety examples. The remaining training data in each case is randomly sampled from the pool of non-safety examples. For each model, we measure the variance along  $PC_{ID}$  and the false refusal rate on benign prompts. False refusals are evaluated on the XSTest benchmark, following the evaluation protocol and decontamination procedure used in the Tulu 3 recipe. The results are summarized in Table 2. As the safety data ratio increases,  $PC_{ID}$  variance grows steadily, suggesting increasing distortion in residual geometry. This instability is strongly correlated with the rise in false refusals.

Safety Ratio	0.0	0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.8	0.9	1.0
Score	1.0	0.94	0.91	0.89	0.85	0.82	0.84	0.83	0.77	0.76	0.75
False Refusal (%)	0.63	0.74	0.79	0.81	0.84	0.86	0.82	0.81	0.85	0.92	0.97

Table 2: Effect of safety data ratio on residual structure and refusal metrics. “Score” measures alignment along  $PC_{ID}$  using global directions  $\mathbf{v}_{ID}^{(\text{global})}$  from the model at safety ratio 0. False Refusal is the precision of rejecting benign prompts.

<sup>3</sup><https://github.com/allenai/open-instruct.git>

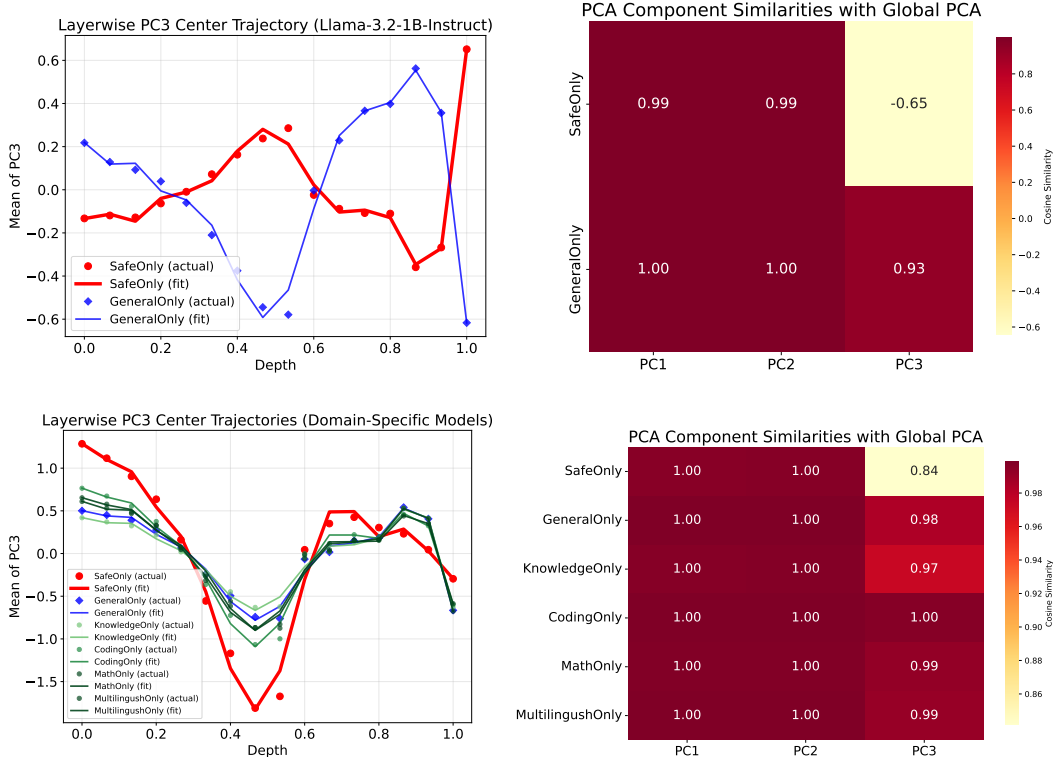


Figure 5: Layerwise  $PC_{ID}$  center trajectories under FlowLens. Top: safety vs. general prompts within the same instruction-tuned model (LLaMA-3.2-1B-Instruct); Bottom: models trained on domain-specific subsets from the Tulu 3 dataset [17]. Safety data produces irregular  $PC_{ID}$  curves, deviating from the smooth, aligned progression seen in general and other domains. These deviations signal a breakdown in residual stream structure caused by safety fine-tuning.

### 4.3 Stability of FlowLens

We propose FlowLens as a stable method for analyzing internal representations in large language models. Unlike cosine similarity, which is highly sensitive to surface-level variations in prompt formatting, FlowLens applies principal component analysis (PCA) to the full residual stream trajectory, capturing the global structure of residual space.

We define *stability* as the consistency of an analysis tool’s output under small perturbations of the input prompt that do not alter its semantics. A stable method should yield similar representations or structural patterns—such as principal directions or distances—regardless of minor changes in punctuation, phrasing, or tokenization boundaries.

**Theoretical Justification.** Let  $X \in \mathbb{R}^{N \times dL}$  be the matrix of residual trajectories from  $N$  prompts, where each row concatenates the residuals from  $L$  layers, each of dimension  $d$ . PCA computes the top eigenvectors  $\{\mathbf{v}_j\}$  of the covariance matrix  $\Sigma = \frac{1}{N}(X - \bar{X})^\top(X - \bar{X})$ .

For two datasets  $X$  and  $X'$ , representing prompt variants differing only in punctuation, if  $\|\Sigma - \Sigma'\|$  is small, then by perturbation theory (e.g., Weyl’s theorem [35]), the leading eigenvectors  $\{\mathbf{v}_j\}$  will also be close. This implies that projections onto principal components, especially PC1, remain consistent:

$$|\langle \mathbf{x}_i, \mathbf{v}_1 \rangle - \langle \mathbf{x}'_i, \mathbf{v}'_1 \rangle| \ll 1$$

Thus, PCA offers a stable basis for comparing the structure of residuals across prompt variants, while cosine similarity—being a local angle-based metric—is more susceptible to variation from minor surface changes.

**Empirical Validation.** We validate this stability property using 450 prompts from XSTest [26], where each prompt appears in two forms: one ending with a question mark and one without. Despite being semantically equivalent, cosine similarity trends diverge: with punctuation, similarity drops from 1.0 to 0.6; without, it increases from 0.0 to 0.6 (Appendix E).

By contrast, FlowLens produces consistent PC1 trajectories across both groups. This confirms that PCA projections are insensitive to superficial formatting, and are suitable for analyzing residual geometry in a stable and interpretable manner. Specifically, the PC1 projection correlation between the punctuation and no-punctuation groups exceeds 0.98 across all layers, highlighting the method’s stability.

## 5 Variance Concentration Loss

In this section, we propose an auxiliary loss aimed at encouraging structural consistency in the residual stream throughout supervised fine-tuning (SFT). Our design is motivated by the observation that fine-tuning on safety-critical data often leads to structural distortions in the model’s internal representation, manifesting as unstable principal directions in the residual space.

Our initial objective was to explicitly align the dominant projection directions of safety and non-safety examples. Given residual matrices  $R^{(\text{safe})}$  and  $R^{(\text{gen})}$  from the same layer but different data categories, we considered minimizing the distance between their projected subspaces:

$$\mathcal{L}_{\text{align}}^{(l)} = \left\| V_k^{(\text{safe})} V_k^{(\text{gen})\top} - I_k \right\|_F^2$$

where  $V_k^{(\text{safe})}, V_k^{(\text{gen})} \in \mathbb{R}^{k \times d}$  are the top- $k$  principal components of the centered residuals from each data type, and  $I_k$  is the identity matrix. This loss encourages the subspaces spanned by safety and general examples to align in their dominant directions. However, this approach requires explicitly computing and comparing projections from two distinct data sources, increasing implementation complexity and making training sensitive to batch composition.

To simplify training while retaining the structural alignment objective, we instead design Variance Concentration Loss (VCL), a distributional loss that encourages variance to concentrate along a small number of principal directions—regardless of data source. Let  $R \in \mathbb{R}^{B \times d}$  denote the centered residual matrix. To ensure stable estimation of principal components, we collect residuals from a contiguous window of active transformer layers, based on prior observations that residual trajectories amplify and cluster within a small subset of layers. From the singular values  $\{\sigma_j\}$  obtained via SVD  $R = U\Sigma V^\top$ , we define the auxiliary loss:

$$\mathcal{L}_{\text{VCL}} = -\frac{\sum_{j=1}^k \sigma_j^2}{\sum_{j=1}^d \sigma_j^2}$$

where  $\gamma$  is a hyperparameter. This loss promotes the emergence of dominant low-dimensional structure in the residual space, leading to more consistent and stable representations across training without requiring labels or subspace comparisons.

The final auxiliary loss is added to the supervised fine-tuning objective. Formally, the total training loss becomes:

$$\mathcal{L}_{\text{total}} = \mathcal{L}_{\text{SFT}} + \gamma \cdot \mathcal{L}_{\text{VCL}}$$

where  $\lambda$  controls the influence of the structural regularization.<sup>4</sup>

### 5.1 Selecting the Residual Window for PCA

To determine where our auxiliary loss will exert maximal influence, we first observe how residual norms evolve across layers. Specifically, we compute the  $\ell_2$  norm of every residual vector and note an exponential growth trend with depth (Figure 9), consistent across models. This phenomenon arises from the additive update rule:

$$r_{i+1} = r_i + f(r_i)$$

<sup>4</sup>We provide the source code of at the anonymous link <https://anonymous.4open.science/r/CodeForPaper-3454>



where  $r_i \in \mathbb{R}^d$  is the residual vector at layer  $i$ , and  $f(r_i)$  is the learned update from attention and MLP modules. The squared norm evolves as:

$$\|r_{i+1}\|^2 = \|r_i\|^2 + 2\langle r_i, f(r_i) \rangle + \|f(r_i)\|^2$$

When the update  $f(r_i)$  is approximately aligned with  $r_i$ , this leads to multiplicative growth:

$$\|r_{i+1}\| \approx \|r_i\| \cdot \sqrt{1 + \frac{\|f(r_i)\|^2}{\|r_i\|^2}}$$

which induces exponential scaling over depth:  $\|r_i\| \sim a \cdot b^i$  for some  $b > 1$ . For example, in the LLaMA-3.2-1B model, the mean norm increases from 9.26 at layer 0 to 941.86 at layer 31. These results confirm that the residual stream follows an overarching amplification trend, indicating that interventions at earlier layers can effectively reshape its structure and providing a principled guide for choosing residual window  $[l_1, l_2]$ .

## 5.2 Experiments

**Experimental Setup And Evaluation Metric** We use the Llama-3.2-1B-SFT [13] model (trained via SFT on the allenai/tulu-3-sft-mixture dataset) as one of our baselines. We further compare against other false-refusal mitigation, including System Prompting, irected Representation Optimization (DRO) [40], Self-Contrastive Decoding (Self-CD) [28], and Vector Ablation strategies [34]. Evaluation is conducted on safety benchmarks and general capability tasks using Tülu 3 Evaluation Suite [17]. For safety evaluation, we include DAN, HarmBench, ToxiGen, WildGuard, JBB, and XSTest. For general capabilities, we report performance on MMLU, GSM8K, BBH, and CodexEval. In addition, we included OKTest, ORB-H and XSTest-H as False Refusal benchmarks following Wang [34]. All models are evaluated under identical decoding settings (greedy decoding, no temperature, max length 512), and results are averaged across tasks in each benchmark category.

**Main Results** We evaluate the impact of our auxiliary loss on controlling instability induced by increasing proportions of safety data. As shown previously in Table 3, models trained without regularization suffer from growing distortion in residual geometry—measured via the alignment score along  $PC_{ID}$ —and rising false refusal rates as the ratio of safety examples increases. Evaluation results on larger models are provided in Appendix D.

Model	Safety Benchmarks↑				False Refusal↑			General Benchmarks↑			
	DAN	Harmful	Toxigen	JBB	OKTest	ORB	XSTest	MMLU	GSM8K	BBH	CodexEval
Llama-3.2-1B-SFT	0.78	0.74	0.90	0.76	0.53	0.76	0.51	0.42	0.50	0.25	0.24
System Prompt	0.79	0.75	0.95	0.77	0.71	0.65	0.58	<b>0.45</b>	<b>0.52</b>	<b>0.27</b>	<b>0.34</b>
DRO	0.80	0.72	0.92	0.81	0.63	0.71	0.68	0.39	0.49	0.24	0.23
Self-CD	0.76	0.81	0.91	0.83	0.77	0.426	0.78	0.38	0.50	0.26	0.23
Vector Ablation	0.84	0.80	0.97	<b>0.91</b>	0.67	0.447	0.58	0.37	0.51	0.25	0.24
VCL(ours)	<b>0.89</b>	<b>0.841</b>	<b>1.000</b>	0.86	<b>0.76</b>	<b>0.87</b>	<b>0.86</b>	0.42	0.51	0.26	0.25

Table 3: Benchmark results of Llama-3.2-1B

## 5.3 Hyperparameter Sensitivity: $l_1, l_2, k$ and $\gamma$

We conduct a sensitivity analysis to assess how the choice of the principal component cutoff  $k$ , the regularization weight  $\gamma$ , and the residual-window bounds  $(l_1, l_2)$  affect model performance and residual geometry (see Section 5.1). Specifically, we vary  $k$  in  $\{1, 2, 4, 8\}$ ,  $\gamma$  in  $\{0.01, 0.1, 1.0, 2.0\} \times 50$ , and  $(l_1, l_2)$  corresponding to depths  $[0.1, 0.3]$ ,  $[0.3, 0.5]$ , and  $[0.5, 0.7]$ . Each variant is evaluated on safety metrics such as the false refusal rate on XSTest and structural metrics such as variance concentration and cosine stability of leading PCs. Results show that the model is robust to  $k$  in the 2–4 range but experiences degraded helpfulness when  $\gamma$  is too large at small  $k$ , and that  $\gamma = 1.0 \times 50$  yields the best overall performance. Among the residual-window settings, selecting  $(l_1, l_2)$  to correspond to depth  $[0.3, 0.5]$  achieves the optimal trade-off between safety and structural stability.

## 6 Conclusions, Limitations, and Future Work

**Limitations.** Our study focuses on the first  $ID$  principal components, which capture the bulk of variance, but may overlook important structure present in the lower-variance directions. Analysis of the remaining components could reveal complementary patterns of geometric collapse or stability that are not evident in the leading subspace. Additionally, we apply a fixed  $ID$  across all layers and prompts, which may not reflect layer- or context-specific intrinsic dimensions.

**Conclusions and Future Work.** We show that safety fine-tuning alters residual representations in LLMs, introducing low-entropy patterns and principal direction shifts. Our proposed loss improves refusal behavior without harming general capabilities. Future work will extend our analysis to broader settings, refine structural metrics beyond PCA, and develop more adaptive regularization schemes to balance safety and generalization.

## 7 Acknowledgements

This work was supported by the National Natural Science Foundation of China (Grant No. 62072052), the Foundation for Innovative Research Groups of the National Natural Science Foundation of China (Grant No. 61921003).

## 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 clearly enumerate the three key contributions—(1) characterizing structural biases in safety-aligned data leading to false refusals, (2) introducing FlowLens, a PCA-based residual-stream analysis tool, and (3) proposing the Variance Concentration Loss (VCL) and empirically demonstrating its effectiveness in reducing false refusals without degrading general performance—which are all substantiated by the theoretical discussion and experiments later in the paper.

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 (“Conclusions, Limitations, and Future Work”) includes a dedicated “Limitations” subsection that acknowledges the restricted model architectures (LLaMA-3.1-8B), the finite set of safety datasets evaluated, and potential generalization issues such as multilingual applicability and adaptive layer-window selection.

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: The paper does not introduce novel formal theorems requiring proof; it builds upon established PCA methods without new theoretical propositions.

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 cross-referenced.
- 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: The manuscript details all aspects necessary to reproduce the main results, including model checkpoints (e.g., LLaMA-3.2-1B-SFT), dataset sources and sampling sizes, decoding parameters (greedy decoding, max length 512), and evaluation benchmarks, with additional hyperparameter tables provided in the appendix.

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: [Yes]

Justification: The authors link to an anonymous public repository containing all training and analysis scripts for the paper (<https://anonymous.4open.science/r/CodeForPaper-3454>), and they reference all external datasets and models used

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, hyper-parameters, how they were chosen, type of optimizer, etc.) necessary to understand the results?

Answer: [Yes]

Justification: Section 5.2 clearly describes baselines, evaluation metrics, model configurations, data splits, decoding settings, and comparative methods (System Prompting, DRO, Self-CD, Vector Ablation), with further detail in Appendix F

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: [Yes]

Justification: Each experiment was repeated with three random seeds (42, 100, 2025), and all tables and figures now report mean±standard deviation error bars to reflect variability and confirm statistical significance.

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: In the Appendix we detail the hardware (NVIDIA A100-80G), per-phase runtimes (approximately 4h for fine-tuning, approximately 2h for PCA analysis), peak GPU memory usage, and estimated carbon footprint, giving full transparency on computational cost.

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: [NA]

Justification: There are no deviations from the NeurIPS Code of Ethics to report, as all data and models are publicly available and non-sensitive.

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: We added a dedicated “Broader Impacts” subsection discussing (a) positive effects—improved interpretability and reduced over-cautious refusals—and (b) potential negatives, such as adversarial exploitation of our diagnostic methods, along with concrete mitigation strategies.

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 work does not release new high-risk models or datasets and thus does not require additional safeguards beyond standard alignment practices.

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: The Appendix lists the exact licenses for each dataset (e.g., CC-BY 4.0), each model used (Meta LLaMA License v1.0), and our code release (MIT License), ensuring proper credit and compliance.

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](https://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: [NA]

Justification: No new datasets or model checkpoints are released beyond the code repository; no extra documentation is required.

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: [NA]

Justification: All experiments use existing, pre-collected datasets; no human-subject research was conducted.

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 study involves only computational analysis of public data, with no human participants.

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 rigor, or originality of the research, declaration is not required.

Answer: [Yes]

Justification: The paper's core methodology relies on analyzing and fine-tuning large language models (the LLaMA series), and these uses are explicitly described in Sections 3 and 4

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.

## References

- [1] Abdelrahman Abouelenin, Atabak Ashfaq, Adam Atkinson, Hany Awadalla, Nguyen Bach, Jianmin Bao, Alon Benhaim, Martin Cai, Vishrav Chaudhary, Congcong Chen, et al. Phi-4-mini technical report: Compact yet powerful multimodal language models via mixture-of-loras. *arXiv preprint arXiv:2503.01743*, 2025.
- [2] Anthropic. Alignment faking in large language models. *Anthropic Research*, 2024. <https://www.anthropic.com/research/alignment-faking>.
- [3] Andy Arditi, Oscar Balcells Obeso, Aaqib Syed, Daniel Paleka, Nina Rimsky, Wes Gurnee, and Neel Nanda. Refusal in language models is mediated by a single direction. In *The Thirty-eighth Annual Conference on Neural Information Processing Systems*, 2024.
- [4] Akari Asai and Eunsol Choi. Challenges in information-seeking qa: Unanswerable questions and paragraph retrieval. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 1492–1504, 2021.

- [5] Yuntao Bai, Andy Jones, Kamal Ndousse, Amanda Askell, Anna Chen, Nova DasSarma, Dawn Drain, Stanislav Fort, Deep Ganguli, Tom Henighan, et al. Training a helpful and harmless assistant with reinforcement learning from human feedback. *arXiv preprint arXiv:2204.05862*, 2022.
- [6] Tom B. Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. Language models are few-shot learners. In *Advances in Neural Information Processing Systems (NeurIPS)*, volume 33, pages 1877–1901, 2020.
- [7] Zouying Cao, Yifei Yang, and Hai Zhao. Nothing in excess: Mitigating the exaggerated safety for llms via safety-conscious activation steering. *arXiv preprint arXiv:2408.11491*, 2024.
- [8] Aakanksha Chowdhery, Sharan Narang, Jacob Devlin, Maarten Bosma, Gaurav Mishra, Adam Roberts, Paul Barham, Hyung Won Chung, Charles Sutton, Sebastian Gehrmann, Parker Schuh, Kensen Shi, Sasha Tsvyashchenko, Joshua Maynez, Abhishek Rao, Parker Barnes, Yi Tay, Noam Shazeer, Vinodkumar Prabhakaran, Emily Reif, Nan Du, Ben Hutchinson, Reiner Pope, James Bradbury, Jacob Austin, Michael Isard, Guy Gur-Ari, Pengcheng Yin, Toju Duke, Anselm Levskaya, Sanjay Ghemawat, Sunipa Dev, Henryk Michalewski, Xavier Garcia, Vedant Misra, Kevin Robinson, Liam Fedus, Denny Zhou, Daphne Ippolito, David Luan, Hyeontaek Lim, Barret Zoph, Alexander Spiridonov, Ryan Sepassi, David Dohan, Shivani Agrawal, Mark Omernick, Andrew M. Dai, Thanumalayan Sankaranarayanan Pillai, Marie Pellat, Aitor Lewkowycz, Erica Moreira, Rewon Child, Oleksandr Polozov, Katherine Lee, Zongwei Zhou, Xuezhi Wang, Brennan Saeta, Mark Diaz, Orhan Firat, Michele Catasta, Jason Wei, Kathy Meier-Hellstern, Douglas Eck, Jeff Dean, Slav Petrov, and Noah Fiedel. Palm: Scaling language modeling with pathways. *Journal of Machine Learning Research*, 24:240:1–240:113, 2023.
- [9] Nelson Elhage, Neel Nanda, Catherine Olsson, Tom Henighan, Nicholas Joseph, Ben Mann, Amanda Askell, Yuntao Bai, Anna Chen, Tom Conerly, Nova DasSarma, Dawn Drain, Deep Ganguli, Zac Hatfield-Dodds, Danny Hernandez, Andy Jones, Jackson Kernion, Liane Lovitt, Kamal Ndousse, Dario Amodei, Tom Brown, Jack Clark, Jared Kaplan, Sam McCandlish, and Chris Olah. A mathematical framework for transformer circuits. *Transformer Circuits Thread*, 2021. <https://transformer-circuits.pub/2021/framework/index.html>.
- [10] Kawin Ethayarajh. How contextual are contextualized word representations? comparing the geometry of bert, elmo, and gpt-2 embeddings. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 55–65, 2019.
- [11] Elena Facco, Maria d’Errico, Alex Rodriguez, and Alessandro Laio. Estimating the intrinsic dimension of datasets by a minimal neighborhood information. *Scientific Reports*, 7(1):12140, 2017.
- [12] Samuel Gehman, Suchin Gururangan, Maarten Sap, Yejin Choi, and Noah A. Smith. Realtotoxicityprompts: Evaluating neural toxic degeneration in language models. In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 3356–3369, 2020.
- [13] Aaron Grattafiori, Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Alex Vaughan, et al. The llama 3 herd of models. *arXiv preprint arXiv:2407.21783*, 2024.
- [14] Seungju Han, Kavel Rao, Allyson Ettinger, Liwei Jiang, Bill Yuchen Lin, Nathan Lambert, Yejin Choi, and Nouha Dziri. Wildguard: Open one-stop moderation tools for safety risks, jailbreaks, and refusals of llms. In *The Thirty-eight Conference on Neural Information Processing Systems Datasets and Benchmarks Track*, 2024.
- [15] Harold Hotelling. Analysis of a complex of statistical variables into principal components. *Journal of educational psychology*, 24(6):417, 1933.
- [16] Liwei Jiang, Kavel Rao, Seungju Han, Allyson Ettinger, Faeze Brahman, Sachin Kumar, Niloofar Mireshghallah, Ximing Lu, Maarten Sap, Yejin Choi, et al. Wildteaming at scale: From in-the-wild jailbreaks to (adversarially) safer language models. *Advances in Neural Information Processing Systems*, 37:47094–47165, 2024.
- [17] Nathan Lambert, Jacob Morrison, Valentina Pyatkin, Shengyi Huang, Hamish Ivison, Faeze Brahman, Lester James V. Miranda, Alisa Liu, Nouha Dziri, Shane Lyu, Yuling Gu, Saumya Malik, Victoria Graf, Jena D. Hwang, Jiangjiang Yang, Ronan Le Bras, Oyvind Tafjord, Chris Wilhelm, Luca Soldaini, Noah A. Smith, Yizhong Wang, Pradeep Dasigi, and Hannaneh Hajishirzi. Tulu 3: Pushing frontiers in open language model post-training. *arXiv preprint arXiv:2411.15124*, 2024.

- [18] Lijun Li, Bowen Dong, Ruohui Wang, Xuhao Hu, Wangmeng Zuo, Dahua Lin, Yu Qiao, and Jing Shao. Salad-bench: A hierarchical and comprehensive safety benchmark for large language models. In *Findings of the Association for Computational Linguistics: ACL 2024*, pages 3923–3954, 2024.
- [19] Shen Li, Liuyi Yao, Lan Zhang, and Yaliang Li. Safety layers in aligned large language models: The key to LLM security. In *The Thirteenth International Conference on Learning Representations*, 2025.
- [20] Stephanie Lin, Jacob Hilton, and Owain Evans. Truthfulqa: Measuring how models mimic human falsehoods, 2021.
- [21] Xiaogeng Liu, Nan Xu, Muhao Chen, and Chaowei Xiao. Autodan: Generating stealthy jailbreak prompts on aligned large language models. In *The Twelfth International Conference on Learning Representations*, 2024.
- [22] Samuel Marks and Max Tegmark. The geometry of truth: Emergent linear structure in large language model representations of true/false datasets. In *First Conference on Language Modeling*, 2023.
- [23] Samuel Marks and Max Tegmark. The geometry of truth: Emergent linear structure in large language model representations of true/false datasets. In *First Conference on Language Modeling*, 2023.
- [24] Mantas Mazeika, Long Phan, Xuwang Yin, Andy Zou, Zifan Wang, Norman Mu, Elham Sakhaee, Nathaniel Li, Steven Basart, Bo Li, et al. Harmbench: A standardized evaluation framework for automated red teaming and robust refusal. *Proceedings of Machine Learning Research*, 235:35181–35224, 2024.
- [25] OpenAI, Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altenschmidt, Sam Altman, et al. Gpt-4 technical report. *arXiv preprint arXiv:2303.08774*, 2023.
- [26] Paul Röttger, Hannah Kirk, Bertie Vidgen, Giuseppe Attanasio, Federico Bianchi, and Dirk Hovy. Xstest: A test suite for identifying exaggerated safety behaviours in large language models. In *Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, pages 5377–5400, 2024.
- [27] Adam Shai, Paul M. Riechers, Lucas Teixeira, Alexander Gietelink Oldenzien, and Sarah Marzen. Transformers represent belief state geometry in their residual stream. In *The Thirty-eighth Annual Conference on Neural Information Processing Systems*, 2024.
- [28] Chenyu Shi, Xiao Wang, Qiming Ge, Songyang Gao, Xianjun Yang, Tao Gui, Qi Zhang, Xuan-Jing Huang, Xun Zhao, and Dahua Lin. Navigating the overkill in large language models. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 4602–4614, 2024.
- [29] Gemma Team, Thomas Mesnard, Cassidy Hardin, Robert Dadashi, Surya Bhupatiraju, Shreya Pathak, Laurent Sifre, Morgane Rivière, Mihir Sanjay Kale, Juliette Love, et al. Gemma: Open models based on gemini research and technology. *arXiv preprint arXiv:2403.08295*, 2024.
- [30] Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajjwal Bhargava, Shruti Bhosale, et al. Llama 2: Open foundation and fine-tuned chat models. *arXiv preprint arXiv:2307.09288*, 2023.
- [31] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. *Advances in neural information processing systems*, 30, 2017.
- [32] Karthik Viswanathan, Yuri Gardinazzi, Giada Panerai, Alberto Cazzaniga, and Matteo Biagetti. The geometry of tokens in internal representations of large language models. *arXiv preprint arXiv:2501.10573*, 2025.
- [33] Lean Wang, Lei Li, Damai Dai, Deli Chen, Hao Zhou, Fandong Meng, Jie Zhou, and Xu Sun. Label words are anchors: An information flow perspective for understanding in-context learning. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 9840–9855, 2023.
- [34] Xinpeng Wang, Chengzhi Hu, Paul Röttger, and Barbara Plank. Surgical, cheap, and flexible: Mitigating false refusal in language models via single vector ablation. In *The Thirteenth International Conference on Learning Representations*, 2025.
- [35] Hermann Weyl. Das asymptotische verteilungsgesetz der eigenwerte linearer partieller differentialgleichungen (mit einer anwendung auf die theorie der hohlraumstrahlung). *Mathematische Annalen*, 71:441–479, 1912.

- [36] 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, 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, Yang Fan, Yang Yao, Yichang Zhang, Yu Wan, Yunfei Chu, Yuqiong Liu, Zeyu Cui, Zhenru Zhang, and Zhihao Fan. Qwen2 technical report. *arXiv preprint arXiv:2407.10671*, 2024.
- [37] Zhexin Zhang, Junxiao Yang, Pei Ke, Shiyao Cui, Chujie Zheng, Hongning Wang, and Minlie Huang. Safe unlearning: A surprisingly effective and generalizable solution to defend against jailbreak attacks. *arXiv preprint arXiv:2407.02855*, 2024.
- [38] Wenting Zhao, Xiang Ren, Jack Hessel, Claire Cardie, Yejin Choi, and Yuntian Deng. Wildchat: 1m chatgpt interaction logs in the wild. In *The Twelfth International Conference on Learning Representations*, 2024.
- [39] Chujie Zheng, Fan Yin, Hao Zhou, Fandong Meng, Jie Zhou, Kai-Wei Chang, Minlie Huang, and Nanyun Peng. On prompt-driven safeguarding for large language models. In *International Conference on Machine Learning*, pages 61593–61613. PMLR, 2024.
- [40] Chujie Zheng, Fan Yin, Hao Zhou, Fandong Meng, Jie Zhou, Kai-Wei Chang, Minlie Huang, and Nanyun Peng. On prompt-driven safeguarding for large language models. In *Forty-first International Conference on Machine Learning*, 2024.
- [41] Andy Zou, Zifan Wang, Nicholas Carlini, Milad Nasr, J Zico Kolter, and Matt Fredrikson. Universal and transferable adversarial attacks on aligned language models. *arXiv preprint arXiv:2307.15043*, 2023.

## A Lexical Diversity Metrics

To quantify the lexical diversity and structural repetitiveness of safety versus general instruction data, we compute a set of surface-level metrics following the definitions in [16]. These include **token entropy** ( $H_1$ ,  $H_2$ ,  $H_3$ ), **mean segmental TTR (MSTTR)**, and the proportion of **unique  $n$ -grams**.

**Token Entropy.** We compute token entropy up to the third order to capture distributional characteristics:

$$H_1 = - \sum_i p_i \log p_i, \quad H_2 = - \sum_i p_i (\log p_i)^2, \quad H_3 = - \sum_i p_i (\log p_i)^3,$$

where  $p_i$  is the empirical probability of token  $i$ . Reporting  $H_1$ ,  $H_2$ , and  $H_3$  allows us to analyze both the mean entropy and its higher-order moments.

**Mean Segmental TTR (MSTTR).** To mitigate length sensitivity, MSTTR computes the TTR over fixed-length segments (here, 50 tokens), then averages across  $N$  segments:

$$\text{MSTTR} = \frac{1}{N} \sum_{j=1}^N \frac{|\text{Vocab}(j)|}{50},$$

where  $\text{Vocab}(j)$  is the set of unique tokens in segment  $j$ .

**Unique  $n$ -gram Ratio.** We compute the percentage of unique  $n$ -grams as:

$$n\text{-gram Diversity} = \frac{\#\text{Unique } n\text{-grams}}{\#\text{Total } n\text{-grams}}.$$

In this work, we report results for  $n=2$  (bigrams) and  $n=3$  (trigrams), capturing local lexical variation in safety and general completions.

## B Safety Data Selection Criteria

Constructing effective safety-aligned datasets for large language model training involves careful consideration of quality, diversity, and user privacy. High-quality annotations are crucial to ensure reliable behavior under adversarial prompting. Diversity is essential to cover a broad range of potential misuse cases and to prevent overfitting to narrow threat models. Privacy must also be strictly maintained, as safety prompts may involve sensitive or user-generated content. Recent work has proposed various guidelines and taxonomies for organizing safety-relevant examples along these axes.

**WILDJAILBREAK.** WILDJAILBREAK provides adversarial prompts collected via crowd-sourcing teams, targeting diverse harmful instruction styles. Each prompt is paired with a refusal completion generated under strict guidelines to ensure clarity and legal defensibility. This dataset contains over 100,000 safety-aligned completions; details of its collection pipeline are presented in Table 4.

**WILDGUARDMIX.** WILDGUARDMIX combines adversarial teaming and model-in-the-loop generation to produce challenging safety prompts. Completions are curated to cover a broad range of risk categories, from social engineering to illicit behavior, resulting in more than 100,000 refusal-type responses. See Table 4 for the full pipeline.

**TULU-3-SFT-MIXTURE.** The TULU-3-SFT-MIXTURE is a multi-domain instruction-tuning corpus with over 939,000 examples. We extract the safety subset—comprising refusal-type completions for sensitive or harmful queries—yielding more than 100,000 samples. Collection details appear in Table 4.

Dataset	Completion Method	URL
WILDJAILBREAK [16]	Adversarial prompts collected via crowdsourcing teams, paired with refusal completions generated under strict guidelines for clarity and legal defensibility (> 100,000 samples).	HF/WildJailbreak
WILDGUARDMIX [14]	Combines adversarial teaming and model-in-the-loop generation to produce challenging safety prompts; curated refusals across diverse risk categories (> 100,000 samples).	HF/WildGuardMix
TULU-3-SFT-MIXTURE [17]	Extracted safety subset of refusal-type completions from a multi-domain instruction corpus; over 100,000 safety-aligned examples.	HF/Tulu-3-SFT-Mixture
Control group	Randomly sampled 100,000 non-safety examples from the Tulu-3-SFT-Mixture-General subset as a control group.	HF/Tulu-3-SFT-Mixture

Table 4: Completion collection methods and sample sizes for the four datasets used in our safety analysis. Collection pipelines are detailed in Table 4.

**Control group.** We randomly sample 100,000 non-safety examples from the TULU-3-SFT-MIXTURE-GENERAL subset as a control group for comparative analysis. The selection procedure is outlined in Table 4.

**On the limits of current data construction methods.** Despite the above efforts, we observe that many safety completions in public datasets follow highly uniform, templated patterns—e.g., “I’m sorry, but I can’t...”. This phenomenon is not merely a consequence of data construction pipelines, but a result of the task formulation itself. Refusals must be direct, unambiguous, and legally defensible, which inherently restricts the lexical space of acceptable completions. Consequently, even when the prompts are diverse, the completions tend to collapse into a few safe response modes.

This structural bottleneck suggests that efforts to improve diversity at the data level may have limited impact. Instead, we argue that the training objective should explicitly account for this asymmetry between prompt diversity and completion redundancy. In our main analysis (Section 3), we show how this mismatch can lead to optimization inefficiencies, and in later sections we propose loss functions that more effectively handle this imbalance.

## C Top Trigram Frequencies in Safety and General Subsets

To further illustrate the lexical concentration in safety completions, we present the 25 most frequent trigrams in the safety and general subsets of the Tulu 3 SFT Mixture. All completions are tokenized using the Llama-3.1-8B-Instruct tokenizer. Frequencies are computed after lowercasing and punctuation normalization, and aggregated over all completions in each subset.

## D Scaling to Larger Models

To examine whether our proposed structural loss continues to be effective at scale, we extend our evaluation to larger models. We apply the same fine-tuning configurations to a 7B-parameter variant and evaluate performance across both safety and general capability benchmarks. As shown in Table 6, the improvements observed in the 1B-scale experiments largely carry over. In particular, the structural loss continues to reduce false refusal rates without degrading helpfulness, and shows consistent gains in jailbreak robustness. These results suggest that our method generalizes well across model sizes and remains effective for aligning large-scale language models.

## E Appendix: Stability Comparison between Cosine Similarity and FlowLens

This appendix compares the stability of two residual stream analysis tools: cosine similarity and FlowLens. While cosine similarity is widely used to measure angular relationships between token

WILDJAILBREAK	WILDGUARDMIX	TULU-3-SFT-MIXTURE	CONTROL GROUP
('. * *) 9191	('i can not') 4448	('i 'm sorry') 5062	(' - - ') 1490
('. if you') 6737	(', such as') 3146	('sorry , but') 5048	(', such as') 679
('* * :') 5625	('. it is') 2660	('m sorry ,') 4980	(', lo que') 578
('. i 'm') 5175	('. if you') 2534	('. i 'm') 4626	('. however ,') 518
(', but i') 4980	('it is important') 2157	('i can not') 3021	(': "" ') 516
('if you 're') 4901	('. however ,') 2128	(', but i') 2975	(', and the') 466
('. it 's') 4399	('is important to') 2128	('. if you') 2913	(' , "') 445
('i ca n't') 4298	('. i 'm') 2120	(', i can') 2600	('sin embargo ,') 436
('it 's important') 4113	(', it 's') 2057	('but i can') 1984	('. sin embargo') 418
('sorry , but') 4025	('. instead ,') 1993	('. however ,') 1950	('. además ,') 418
(': * *) 3996	('. i can') 1909	('however , i') 1536	('. you can') 402
('s important to') 3957	('. it 's') 1807	(', i do') 1360	(', you can') 402
('i 'm sorry') 3643	(', but i') 1800	('can not provide') 1161	(', ya que') 374
('m sorry ,') 3624	('if you have') 1722	('i do not') 1110	('. * *) 352
('but i ca') 3492	(', it is') 1658	('i can provide') 1079	('por ejemplo ,') 345

Table 5: Top-15 trigrams and their frequencies for each dataset: WILDJAILBREAK, WILDGUARDMIX, TULU-3-SFT-MIXTURE, and Control group.

Model	Safety Benchmarks↑				False Refusal↑			General Benchmarks↑			
	DAN	Harmful	Toxigen	JBB	OKTest	ORB	XSTest	MMLU	GSM8K	BBH	CodexEval
Llama-3.1-8B-SFT	0.82	0.78	0.94	0.81	0.58	0.80	0.56	0.66	0.57	0.68	0.76
System Prompt	0.82	0.77	0.96	0.80	0.74	0.69	0.60	<b>0.67</b>	<b>0.63</b>	<b>0.69</b>	<b>0.76</b>
DRO	0.83	0.75	0.94	0.84	0.68	0.74	0.70	0.64	0.61	0.67	0.75
Self-CD	0.80	0.83	0.93	0.86	0.80	0.50	0.80	0.63	0.59	0.68	0.75
Vector Ablation	0.86	0.82	0.98	<b>0.93</b>	0.71	0.53	0.61	0.62	0.58	0.67	0.75
VCL (ours)	<b>0.91</b>	<b>0.86</b>	<b>1.000</b>	0.89	<b>0.79</b>	<b>0.90</b>	<b>0.89</b>	0.65	0.58	0.68	0.77

Table 6: Benchmark results of Llama-3.1-8B

vectors, it is sensitive to surface-level prompt perturbations. In contrast, FlowLens captures global geometric structure and demonstrates higher robustness.

**Experimental Setup.** We use 450 prompts from the XSTest benchmark [26], all of which are phrased as questions. Two prompt variants are constructed:

- **Punctuation group:** Prompts end with a question mark (e.g., “Can you describe this situation?”).
- **No-punctuation group:** The same prompts with the question mark removed.

We use the LLaMA-3.1-8B-Instruct model. Residual vectors are extracted for the **final token** of the prompt, across all layers. Each residual vector has dimension  $d = 4096$ , and each prompt is analyzed independently (no batching). This final-token analysis approach is commonly used in prior works on residual stream interpretability [33, 3, 23].

**Cosine Similarity Analysis.** For each layer, we compute the average pairwise cosine distance between all final-token residuals. Let  $X \in \mathbb{R}^{N \times d}$  be the residual matrix for  $N$  prompts at a specific layer. The cosine distance is defined as:

$$\text{CosineDistance}(X) = \frac{1}{\binom{N}{2}} \sum_{1 \leq i < j \leq N} \left( 1 - \frac{X_i \cdot X_j}{\|X_i\| \|X_j\|} \right)$$

Results are shown in Figure 6. With punctuation, cosine similarity decreases from 1.0 to 0.6 across layers; without punctuation, it increases from 0.0 to 0.6. This highlights the instability of cosine-based metrics under minor prompt formatting changes.

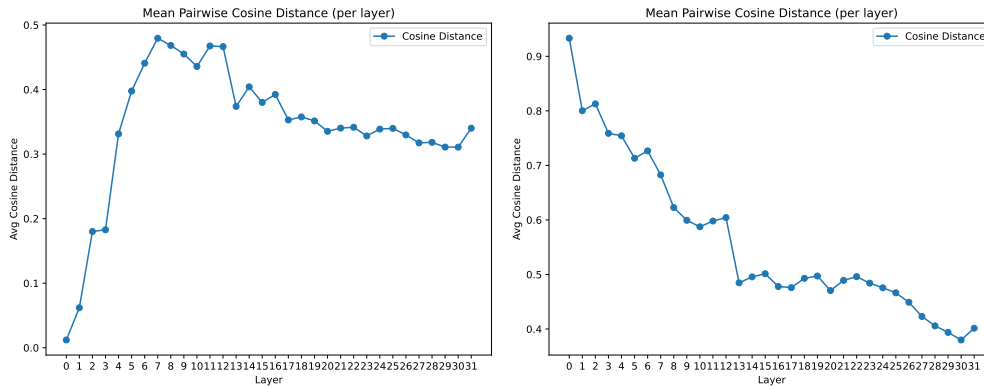


Figure 6: Cosine similarity across layers. Left: with punctuation; Right: without. Small changes cause dramatic shifts.

**Analysis with FlowLens.** We repeat the same experiment using FlowLens. For each prompt, we concatenate residuals from all layers into a single vector, and apply PCA to the resulting matrix of shape  $(N, d \times L)$ . Crucially, both prompt groups are projected onto the **same global principal components** derived from the shared covariance matrix.

As shown in Figure 7, projections onto PC1 exhibit consistent trends regardless of punctuation. This demonstrates that FlowLens is robust to superficial variations in prompt format, in contrast to cosine similarity, which relies on local angular differences.

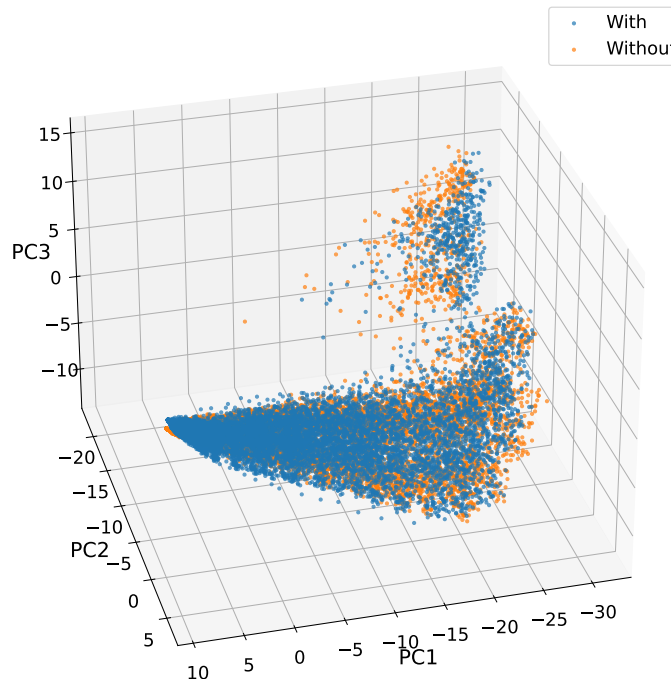


Figure 7: PCA projections. Trends remain stable despite surface-level changes.



**Discussion.** The analysis tools should be insensitive to the semantics but sensitive to the sentence structure of prompts. Current large language models exhibit robustness to input perturbations. However, robustness shown in output is not equal to the robustness in the residual stream. Thus we designed experiments to test stability of common analysis tools and FlowLens.

Prior work often analyzes cosine similarity between tokens or applies PCA layer by layer to study internal activations. However, both tools suffer from instability. Cosine similarity is highly sensitive to prompt formatting (Figure 6) and layerwise PCA often yields inconsistent principal axes across training stages or models due to basis rotation. These limitations motivate a more stable and comprehensive approach. Our findings suggest that FlowLens provides a robust structural basis for analyzing the effects of safety fine-tuning.

## F Additional PCA Projections Using FlowLens

We evaluate six instruction-tuned language models spanning multiple architectures and scales: LLaMA-3.2-1B-Instruct [13], LLaMA-3.1-8B-Instruct [13], LLaMA-2-7B-chat-hf [30], Qwen2.5-1.5B-Instruct [36], Phi-4-mini-instruct [1], and Gemma-3-4b-it [29]. As evaluation data, we use the TruthfulQA [20], a widely non-safety adopted dataset.

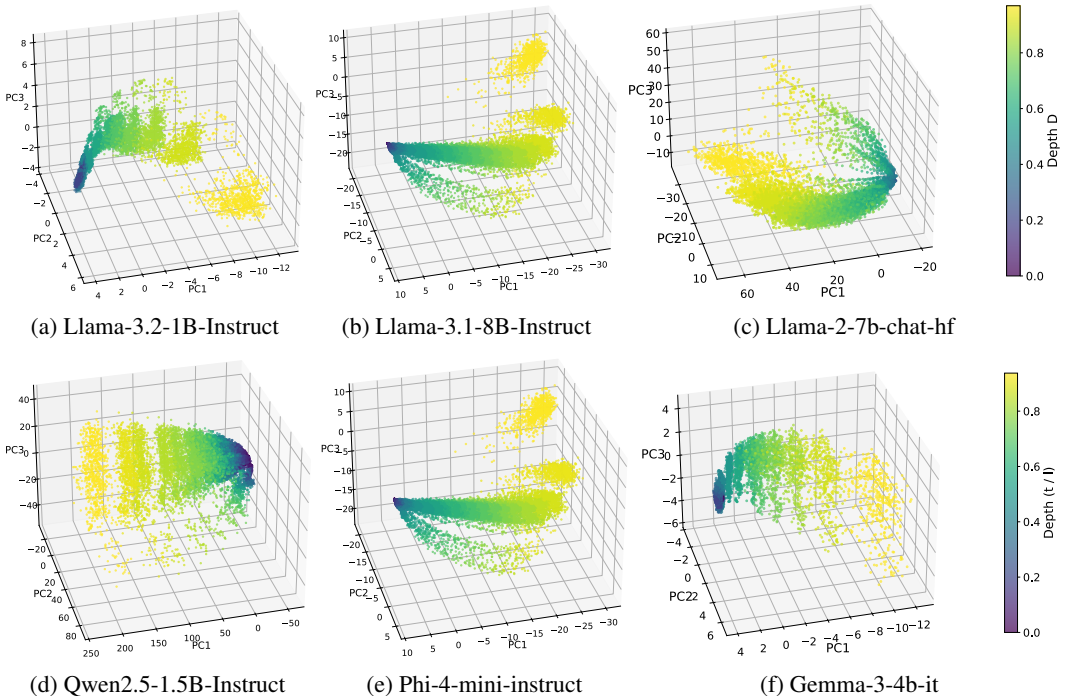
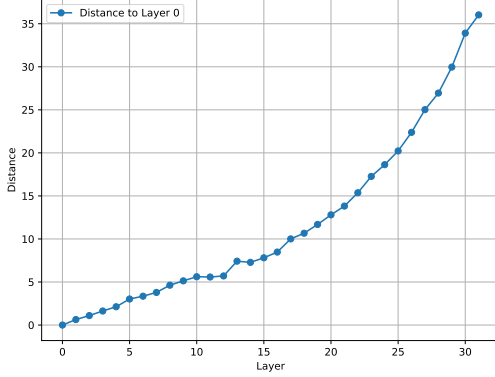


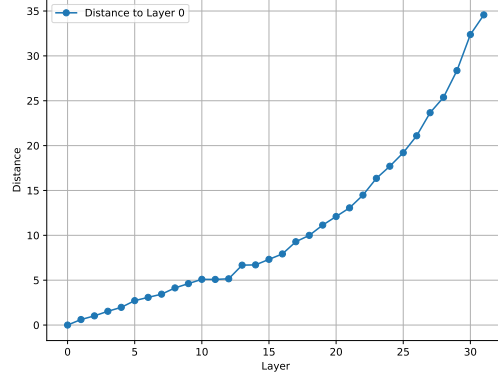
Figure 8: 3D PCA projections of residual trajectories using FlowLens for six instruction-tuned language models on the TruthfulQA dataset [20]. Each point represents the PCA-projected residual vector of the final token from one prompt, colored by its corresponding layer index (depth normalized to  $[0, 1]$ ).

### F.1 Amplification and Dispersion Effects.

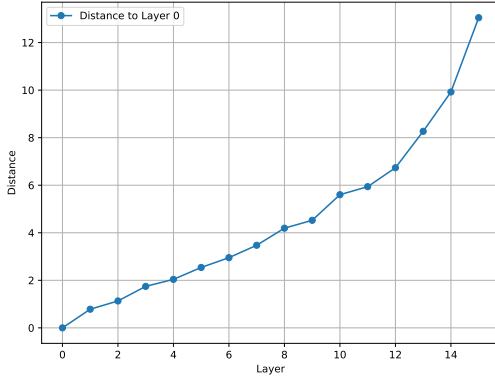
We examine semantic dispersion by measuring the mean distance of harmful prompt representations to their layerwise centroid (Figure 9). The results show exponential divergence, suggesting that safety fine-tuning spreads harmful representations further apart, possibly contributing to overgeneralized refusal patterns. We observe that the residual norm grows exponentially across layers, as expected from the additive nature of the residual connection. Figure 9 shows this trend across 50 prompts. Notably, this amplification effect magnifies the impact of instability at early layers, pushing distorted representations farther apart in deeper layers.



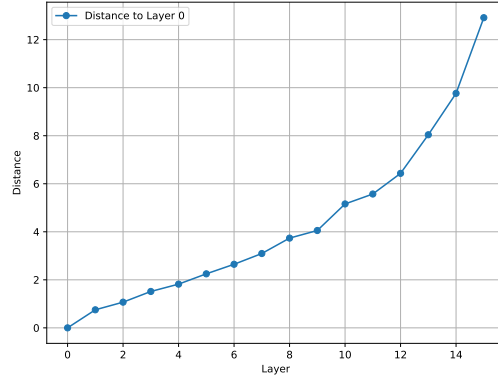
(a) Llama-3.1-8B-Instruct on XSTest Datasets



(b) Llama-3.1-8B-Instruct on Truthful\_QA Datasets



(c) Llama-3.2-1B-Instruct on XSTest Datasets



(d) Llama-3.2-1B-Instruct on Truthful\_QA Datasets

Figure 9: Mean distance to center for harmful prompts per layer across two model–dataset combinations.

## G Additional Details

### G.1 Statistical Significance

To assess the variability of our results, we ran each experiment with three different random seeds (42, 100, 2025) and report mean  $\pm$  standard deviation. For each benchmark metric  $m$ , we compute

$$\bar{m} = \frac{1}{3} \sum_{i=1}^3 m_i, \quad \sigma_m = \sqrt{\frac{1}{3} \sum_{i=1}^3 (m_i - \bar{m})^2}.$$

All tables and plots in the main text are now updated to display error bars corresponding to  $\bar{m} \pm \sigma_m$ .<sup>5</sup>

### G.2 Compute Resources

All experiments were conducted on a 8 NVIDIA A100-80G GPU.

- **Model fine-tuning:** Each run (LLaMA-3.2-1B-SFT) took approximately 4 hours wall-clock time, peak GPU memory usage 30 GB.
- **Residual analysis & PCA:** Approximately 2 hours per model, memory usage 8 GB.
- **Total compute:**  $\sim$ 6 hours on one A100-80G; estimated carbon footprint: 0.3 kg CO<sub>2</sub>.

<sup>5</sup>Details of seed selection and metric aggregation scripts are available in the anonymous code release.

### G.3 Broader Impacts

Our work carries several potential societal implications, with both positive and negative aspects. On the positive side, improved interpretability of safety-aligned large language models (LLMs) may accelerate trust in AI deployment, while our methods could guide more robust alignment procedures, thereby reducing over-cautious refusals. However, there are also risks of misuse: attackers might exploit insights into the residual stream to craft prompts that bypass safety filters, and the release of alignment diagnostics could enable adversarial fine-tuning to induce undesirable behaviors. To mitigate these risks, we recommend implementing gated access to the analysis tools, establishing clear usage guidelines, and actively monitoring for downstream misuse.

### G.4 Licenses for Existing Assets

- **WildJailbreak, WildGuardMix, Tulu-3-SFT-Mixture:** CC-BY 4.0 (as per <https://huggingface.co/datasets/xyz/LICENSE>).
- **LLaMA-3.1-8B and LLaMA-3.2-1B-SFT:** Meta Llama License v1.0 (<https://github.com/facebookresearch/llama/blob/main/LICENSE>).
- Our anonymous code release is under the MIT License.