

# Embracing Anisotropy: Turning Massive Activations into Interpretable Control Knobs for Large Language Models

Youngji Roh Hyunjin Cho Jaehyung Kim

Yonsei University

{youngjiroh, cyberhyunjin, jaehyungk}@yonsei.ac.kr

## Abstract

Large Language Models (LLMs) exhibit highly anisotropic internal representations, often characterized by *massive activations*, a phenomenon where a small subset of feature dimensions possesses magnitudes significantly larger than the rest. While prior works view these extreme dimensions primarily as artifacts to be managed, we propose a distinct perspective: these dimensions serve as intrinsic interpretable functional units arising from domain specialization. Specifically, we propose a simple magnitude-based criterion to identify *Domain-Critical Dimensions* in a training-free manner. Our analyses reveal that such dimensions behave as interpretable semantic detectors for symbolic/quantitative patterns or domain-specific terms. In addition, we introduce *Critical Dimension Steering*, which applies activation steering exclusively to the identified dimensions. Empirical results show that this approach outperforms conventional whole-dimension steering in domain adaptation and jailbreaking scenarios.<sup>1</sup>

## 1 Introduction

Despite the versatile capabilities of Large Language Models (LLMs) across diverse tasks, prior research indicates that their internal activations exhibit a distinct lack of *isotropy* (Ethayarajh, 2019; Rudman and Eickhoff, 2023). In particular, analyses on Transformer-based models reveal that representations do not uniformly utilize the embedding space dimensions (Ding et al., 2022; Razzhigaev et al., 2024). This anisotropy, characterized by a highly non-uniform distribution of activations, is closely linked to the presence of a few outlier dimensions exhibiting variances orders of magnitude larger than others (Hämmerl et al., 2023; Dettmers et al., 2022). Sun et al. (2024) characterizes this

<sup>1</sup>Code is available at <https://github.com/nyancat0222/dimension-analyzer>

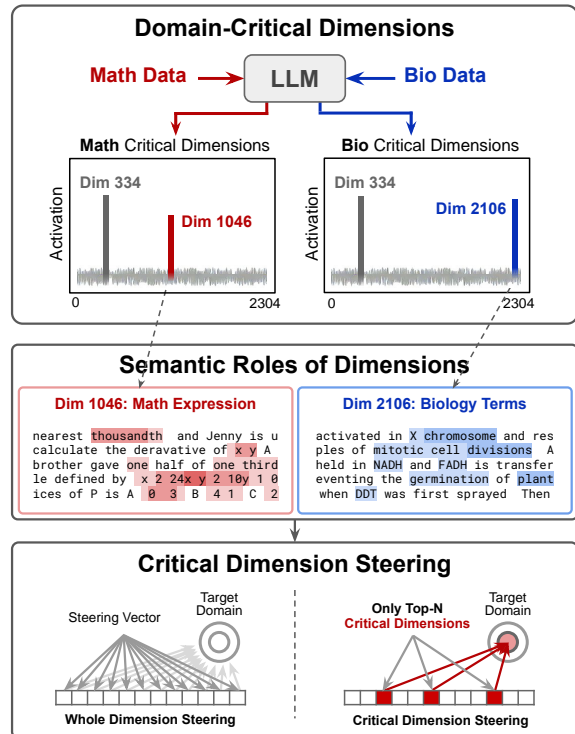


Figure 1: Overview of the study. Domain-Critical Dimensions identified via activation magnitude act as interpretable functional units. We also introduce Critical Dimension Steering, which intervenes only on the identified dimensions for precise model control.

phenomenon as *massive activations*, where a small set of features attains values tens of thousands of times larger (*e.g.*,  $100,000\times$ ) than the remainder.

A substantial body of literature argues that this representational imbalance prevents models from fully exploiting their capacity, potentially obscuring meaningful linguistic distinctions (Gao et al., 2019; Timkey and van Schijndel, 2021). Numerous studies have proposed to algorithmically suppress the influence of extreme dimensions in order to promote representational isotropy (Wei et al., 2022; He et al., 2024). Conversely, some perspectives suggest that anisotropy is a *natural* and potentially *beneficial* byproduct of stochastic gradient descent (Zhu et al., 2018), arguing that preserv-

ing these characteristics is crucial for maintaining downstream performance (Rudman and Eickhoff, 2023). However, both perspectives largely treat these extreme dimensions as *artifacts* to be managed, either suppressed or preserved, rather than as features whose intrinsic functional role should be examined.

In this work, we propose a distinct perspective: *we hypothesize the emergence of extreme dimensions as a direct consequence of functional specialization*. Consequently, these dimensions likely to encode domain-specific knowledge, serving as intrinsic interpretable functional units that can be leveraged for better model control. Unlike interpretation methods such as probe-based classifiers (Burns et al., 2023; Li et al., 2023) or sparse autoencoders (SAEs) (Cunningham et al., 2023; Bricken et al., 2023), which require learning additional parameters to map internal states to concepts or disentangle polysemantic features, our perspective allows us to uncover specialized features directly from the pre-trained model by exploiting simple statistics without additional training.

Based on this hypothesis, we present a method for identifying specialized, domain-specific critical dimensions in §2. Our approach is grounded in two key observations regarding feature dimensions: (1) *sparsity*, where a small subset of dimensions plays a crucial role in model performance, and (2) *extremity*, where domain-discriminative dimensions naturally emerge with high-magnitude activations. Building on these insights, we propose a simple yet effective approach to identify **Domain-Critical Dimensions** by selecting the top- $k$  dimensions based on activation magnitude. These dimensions substantially overlap with the ground-truth critical dimensions discovered through computationally expensive iterative masking and evaluation, indicating that functional importance can be inferred from activation statistics without supervision.

Given the functional importance of these dimensions, in §3, we conduct qualitative analyses to investigate their semantic roles by examining token-level activation patterns. These analysis reveals that individual domain-critical dimensions behave as highly interpretable feature detectors, as they are selectively activated on coherent concepts. For example, with experiments on Gemma-2-2b-it (Team et al., 2024), we reveal that dimension 1046 strongly activates on mathematical terms (+, x, ∞), dimension 2106 on specific biological terms (\_ATP, \_NAD, \_phosphorylation), and di-

mension 334 on the topic keywords (\_mathematics, \_biology).

Finally, we demonstrate that domain-critical dimensions can be utilized for better model control (§4). We introduce **Critical Dimension Steering (CDS)**, a targeted inference-time intervention that applies activation steering—*i.e.*, modifying internal activations during the forward pass (Turner et al., 2023; Rinsky et al., 2024)—only to the top- $k$  dimensions identified by our method. As a sparse set of dimensions dominates behavior within a target domain, manipulating activations along these specific axes serves as more effective control knobs than influencing the entire latent space. We validate the efficacy of CDS in two distinct applications. In a domain adaptation setup using MMLU benchmark (Hendrycks et al., 2021), CDS improves domain-specific accuracy over whole-dimension steering baseline in 34 out of 57 subjects. In a jailbreaking setup with AdvBench (Zou et al., 2023b), applying CDS increases attack success rate to 92%, surpassing 84% achieved by whole-dimension steering baseline.

## 2 Existence and Identification of Domain-Critical Dimensions

In this section, we present a method for identifying the top- $k$  dimensions critical to a target domain. In §2.1, we first observe the existence of two types of feature dimensions: *functionally critical* dimensions and *domain-discriminative* dimensions. In §2.2, we then leverage these properties to identify the final *domain-critical dimensions* using simple statistical indicators.

**Problem setup.** We consider pre-trained LLM composed of  $L$  Transformer layers and hidden dimension  $D$ . Given an input token sequence  $\mathbf{x} = [x_1, \dots, x_T]$ , the model first maps tokens to the input embedding  $h_0 \in \mathbb{R}^{T \times D}$ . Each layer  $l$  takes the hidden state  $h_{l-1}$  as input and produces  $h_l$  via residual connection, where the sub-layer output is added to the input. We refer to *activations* as scalar entries of  $h_l$ , and dimensions that encode significant semantic information or functional roles as *feature dimensions*. Our goal is to identify domain-critical dimensions  $\mathcal{I}_{\text{ded}} \subset \{1, \dots, D\}$ , feature dimensions that are both specialized to a target domain and critical for preserving performance within that domain.

Throughout this work, we operationalize a *domain* using the MMLU benchmark (Hendrycks

Table 1: Masking effect of individual dimensions. Rank  $k$  denotes the average accuracy after masking the single dimension that rank  $k$  in terms of accuracy drop for each subject. Baseline accuracy for Gemma-2-2b-it is 56.53% and Qwen-3-8b is 73.30%.

Model	Rank 1	Rank 2	Rank 5	Rank 10	Rank 100
<b>Gemma-2-2B-IT</b>	41.97	46.81	50.71	52.39	54.25
Acc. Drop (%)	<b>(-14.56)</b>	(-9.72)	(-5.82)	(-4.14)	(-2.28)
<b>Qwen-3-8B</b>	21.97	38.67	63.44	68.98	71.97
Acc. Drop (%)	<b>(-51.33)</b>	(-34.63)	(-9.86)	(-4.32)	(-1.33)

et al., 2021), which comprises 57 subjects spanning STEM, humanities, and social sciences. For each subject, we sample 100 prompts from the test split and partition them into two disjoint subsets: 50 prompts as an *identification set* and 50 prompts as an *evaluation set*. The identification set is used to extract  $\mathcal{I}_{\text{dcd}}$ , while the evaluation set is reserved for validating the identified dimensions and for downstream evaluations.

## 2.1 Intriguing Feature Dimensions of LLMs

**Existence of critical dimensions.** We first observe sparse dimensions that are *functionally critical*. Intuitively, one way to estimate the importance of a specific dimension is measuring the performance degradation when target dimension is removed from the inference process. If all dimensions contribute equally to the model’s performance, masking a single dimension among thousands should result in a small performance loss. To verify this, we iteratively evaluate the performance after masking individual dimension on Gemma-2-2b-it (Team et al., 2024) and Qwen-3-8b (Yang et al., 2025). For each dimension, we set its activation to zero at every layer and evaluate the accuracy using the evaluation set. Input embedding layer is excluded from masking, to isolate the effect of internal processing.

From this experiment, we observe that *disabling an even single dimension can yield a catastrophic collapse in model capabilities*. Table 1 presents the rank-wise average masking effect of individual dimensions across subjects of evaluation set, showing that masking the most critical individual dimension alone yields an average performance drop of 51.3% in Qwen-3-8b. Figure 2 shows the distribution of performance drop on *abstract algebra* subject when masking individual dimensions of Gemma-2-2b-it. While most of dimensions show negligible impact on performance, very small number of dimensions exhibit significant performance




Figure 2: Impact of masking individual dimensions on *Abstract Algebra*. A sharp accuracy drop is observed in few dimensions. Annotations indicate the number of dimensions to each level of accuracy drop.

drops, showing a maximum accuracy drop of 26%. The same trend appears across all subjects, and these results are presented in §A.1. Overall, these results indicate that hidden state dimensions do not contribute equally to the model’s task performance, but rather that there exist few *functionally critical dimensions* that crucially affect performance.

**Existence of discriminative dimensions.** Next, we observe extreme dimensions that are activated in a *domain-discriminative* manner. As illustrated in §A.2, hidden states corresponding to queries from 14 high school related subjects naturally form distinct clusters across layers, confirming domain-wise separation in the representational space. To identify specific dimensions driving this distinction, we conducted a comparative analysis between *high school mathematics* and *high school biology* subjects. Specifically, we consider a dimension as *active* for a given query, if its activation value deviates by more than  $3\sigma$  from the mean, following the outlier dimension criterion of Zhao et al. (2025).

For each subject, we compute activation frequency for every dimension, defined as the percentage of queries within the dataset where the dimen-




Figure 3: Dimensions with discriminated activation patterns between *high school mathematics* and *high school biology*. Red indicates higher activation frequency for mathematics, while Blue indicates biology.

The following are multiple choice questions (with answers) about high school biology. You must respond with a single alphabet character.

Energy is harvested during cellular respiration in stages. Which of the following correctly states which phase of cellular respiration harvests the most energy and the correct explanation why?

- A. The most energy is released during the Krebs cycle because it is here that pyruvate is completely ...
- B. ... both FADH<sub>2</sub> and NADH are produced. Each of those molecules will release 2 ATPs ...
- C. ... phosphorylation because in addition to the phosphorylation of ADP into ATP ...
- D. ... during oxidative phosphorylation because H<sub>2</sub>O is completely broken down into H<sup>+</sup> and O<sub>2</sub>.

Answer:

334 (Top 1)		1068 (Top 4)		2106 (Top 7)	
Act.	Token	Act.	Token	Act.	Token
2048.0	<bos>	65.0	_the	43.8	H
324.0	The	64.0	_the	43.0	_ATP
147.0	_biology	62.0	_is	42.2	_ATP
139.0	Energy	61.0	_to	41.5	_respiration
138.0	_Krebs	59.5	_to	41.2	_energy
137.0	_cellular	58.8	_the	41.2	_ATP
132.0	_school	57.5	_the	41.2	_phosphorylation
130.0	_harvested	55.8	_that	40.8	_ADP
130.0	_cellular	55.0	_The	38.5	_NAD
130.0	_respiration	54.8	_in	38.2	_phosphorylation

Figure 4: Prompt-level analysis of domain-critical dimensions. (a) A representative query regarding cellular respiration used for the token analysis. (b) Top activating tokens for dimensions 334, 1068, and 2106. Dimension 334 highlights topic-relevant tokens, 1068 targets functional tokens, and 2106 focuses on biological terminology related to cellular respiration.

sion is classified as active. In Figure 3, we observe discriminative dimensions that exhibit a distinct disparity in activation frequency (*i.e.*, > 30%) between the two subjects. Dimensions such as 682 and 1261 are persistently active in biology, whereas dimensions like 987 and 1046 exhibit the opposite pattern. This suggests that the model’s ability to differentiate domains is encoded within specific, highly active *domain-discriminative dimensions*.

## 2.2 Identifying Domain-Critical Dimensions

Our preliminary observations highlight two key properties of feature dimensions in LLMs:

- 1) *Sparsity*: a small set of critical dimensions critically affects performance.
- 2) *Extremity*: discriminative dimensions are characterized by extreme activation values.

While we aim to identify and analyze these structurally significant dimensions, locating them via exhaustive iterative search (§2.1) is limited for real-world applications, as this approach is computationally expensive. On the other hand, we observed that domain-discriminative dimensions could be found by identifying the dimensions with extreme activations. From this, we hypothesize that functionally critical dimensions also exhibit the statistical property of *extremity*, such as high magnitude.

Therefore, we propose efficient approach to identify these domain-specific critical dimensions, which we term **Domain-Critical Dimensions**, by leveraging statistical properties of activations instead of explicitly measuring the impact of dimensions to the performance. Specifically, we extract hidden states across  $L$  layers of  $N$  queries from a target domain dataset using the identification set. Let  $h_{l,t,j}^n$  denote the activation of the  $j$ -th dimension at the  $t$ -th token in layer  $l$  for the  $n$ -th query. Then, we define *importance score*  $s_j$  for  $j$ -th dimension as the average magnitude of its token-

averaged activations across all layers and queries:

$$s_j = \frac{1}{L} \sum_{l=1}^L \left( \frac{1}{N} \sum_{n=1}^N \left| \frac{1}{T} \sum_{t=1}^T h_{l,t,j}^n \right| \right)$$

Overall, importance score reflects *how strongly the dimension is activated in the target domain*.

We rank all dimensions in descending order of importance score and select the top-100 dimensions as our candidate set of domain-critical dimensions  $\mathcal{I}_{dcd}$ . To validate the selected dimensions, we compare them against the ground-truth functionally critical dimensions found in §2.1. Remarkably, the proposed approach exhibits an average recall of 89.39% for the functionally critical dimensions ranked within the top-10. Furthermore, the masking impact of  $\mathcal{I}_{dcd}$  closely mirrors the results of Table 1, with masking the rank-1 domain-critical dimension yielding an average performance drop of 11.40%. Detailed validation results are in §A.3.

## 3 Semantic Roles of Domain-Critical Dimensions

In this section, we investigate the semantic roles of domain-critical dimensions by examining highly activated tokens using Gemma-2-2b-it. The results with Gemma-2-9b-it and Llama-3.1-8b-it, which are consistent with the results on Gemma-2-2b-it, are presented in §A.7 and §A.8, respectively.

**Prompt-level analysis.** We first analyze the roles of domain-critical dimensions at the prompt level. For a given input prompt  $x$  and a selected dimension  $j$ , ranked by  $s_j$  in §2.2, we define an activation for  $t$ -th token as  $a_{t,j} = \frac{1}{L} \sum_{l=1}^L h_{l,t,j}$  to identify the highest *activated tokens* within the token sequence. We investigate whether these tokens have semantic interpretability. For intuitive illustration, we examine a representative input prompt about *cellular respiration* in Figure 4a.

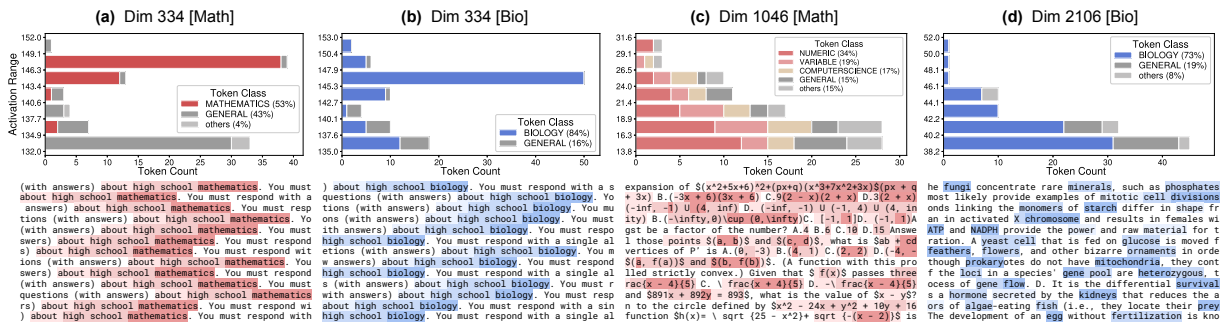


Figure 5: Dataset-level analysis of domain-critical dimensions. The bar charts display the distribution of domain token classes across the activation ranges for the most activated 100 tokens. The text sequences visualize the activation heatmaps for the corresponding dimensions.

Here, we observe the existence of semantic commonalities among tokens. Figure 4b presents the highest activated tokens corresponding to the top-1, top-4, top-7 ranked *high school biology* critical dimensions within a representative prompt.<sup>2</sup> 334 shows extreme values at the initial positional tokens (<BOS>, The), but subsequently activates strongly on tokens relevant to prompt topic such as `_biology`, `_Krebs`, `_cellular`, `_respiration`. 1068 consistently activates on functional words like `_the`, `_is`, `_to`. 2106 triggers on terms such as `_ATP`, `_NAD` and `_phosphorylation`, sharing a semantic related to the specific biological process of cellular respiration.

**Dataset-level analysis.** We extend our analysis to biology and math datasets to determine whether domain-critical dimensions exhibit global semantic roles. The biology dataset is sourced from the *high school biology* subject, and the math dataset is from the *high school mathematics* subject. For each dataset, we perform token classification via Named Entity Recognition, and analyze 100 tokens which exhibit the highest activation for each dimension (see §B.1 for details).

- **Dimension 334** emerges as the highest ranked domain-critical dimension in both datasets. Figure 5a shows that `MATHEMATICS` entities constitute the largest share of activated tokens (53%), exceeding `GENERAL` tokens (43%). In addition, the activation range breakdown shows that tokens with the highest activations are dominantly classified as mathematical entities, whereas general tokens tends to appear in lower activation ranges. The activation heatmap further aligns with this view: activations concentrate around

<sup>2</sup>These are selected to illustrate diverse semantic roles. See §A.5 for other dimensions and additional prompt examples.

explicit domain markers in the prompt, such as `_mathematics`, suggesting that dimension 334 acts as a micro-level topic activator and a macro-level domain classifier.

In the biology dataset, dimension 334 exhibits the same functional role but with biology content. As shown in Figure 5b, `BIOLOGY` entities accounts for 84%, making up the overwhelming majority of the activated tokens. The activation heatmap similarly highlights activation concentration around explicit domain marker, `_biology`, similar to the behavior on math dataset. Taken together, these results suggest that dimension 334 captures domain identity in a direct and dataset-consistent way; it amplifies tokens that explicitly signal the domain and tokens that carry domain-specific topical content.

- **Dimension 1046** is prominent in the math dataset and specializes in quantitative and symbolic content. Top activated tokens by dimension 1046 on math dataset include numeric magnitude words and arithmetic symbols (e.g., `∞`, `2`, `x`, `+`), indicating a strong association with numbers, operators, and variable-like tokens. Figure 5c shows that `NUMERIC` and `VARIABLE` tokens account for 34% and 19% of the activated tokens, respectively. The activation heatmap concentrates around equations and operator dense spans, suggesting that dimension 1046 tracks quantitative form. In the biology dataset, by contrast, it activates `GENERAL` tokens by 44.0%, showing no domain specificity.
- **Dimension 2106** is a biology-specific domain-critical dimension. In Figure 5d, `BIOLOGY` entities account for 73% of activated tokens, with `GENERAL` at 19% and the remainder in other categories. The top activated tokens in biology include concrete biological terms such as

`_roots` and `_cell`, consistent with a lexically grounded biology feature. In the corresponding heatmap, dimension 2106 responds strongly to tokens like `X_chromosome`, and `mitochondria`, showing that it captures fine grained biological semantics rather than generic ones. However, this domain-specific focus vanishes in the math dataset, where `GENERAL` tokens accounts for the majority of 54.0%.

#### 4 Enhanced Model Controllability via Critical Dimension Steering

In this section, we introduce **Critical Dimension Steering (CDS)**, an inference-time intervention method that restricts activation steering to a sparse set of top- $k$  domain-critical dimensions. Our key idea is that if domain-specific information is encoded in a subspace of domain-critical dimensions, then steering only that subspace would yield more accurate domain-targeted control while minimizing interference with unrelated capabilities.

**Activation steering and CDS.** Activation steering is a method that modulates model behavior by shifting internal representations along a specific direction during inference. Steering vector  $\mathbf{v}$  should encode a shift from undesired behavior to desired one. Let  $\mathcal{D}_+$  and  $\mathcal{D}_-$  denote the positive and negative datasets representing the desired and undesired behaviors. For a given layer  $l$ , the mean activation vectors  $\mu_+$  and  $\mu_-$  is computed by averaging hidden states over token positions within each prompt and subsequently averaging across all prompts in the respective sets. Then, the steering vector  $\mathbf{v}_l \in \mathbb{R}^D$  is defined as  $\mathbf{v}_l = \mu_+ - \mu_-$ .

The standard activation steering injects  $\mathbf{v}_l$  into the forward pass of the model by modifying the hidden state  $h_l$  to  $\tilde{h}_l$  as follows:

$$\tilde{h}_l = h_l + \alpha \cdot \mathbf{v}_l,$$

where  $\alpha$  is a scalar controlling the steering strength. CDS modifies this process by applying the intervention only to the identified dimensions. Let  $m \in \{0, 1\}^D$  be a binary mask where  $m_j = 1$  if dimension  $j$  is among the top- $k$  domain-critical dimensions, and  $m_j = 0$  otherwise. Then, the activation steering is selectively applied with  $m$ :

$$\tilde{h}_l = h_l + \alpha \cdot (m \odot \mathbf{v}_l),$$

where  $\odot$  denotes element-wise multiplication.

#### 4.1 Experiment on Domain Adaptation

**Setup.** We evaluate the effectiveness of CDS on *domain adaptation* with Gemma-2-2b-it (Team et al., 2024) using MMLU benchmark (Hendrycks et al., 2021). Following the setup in §2, we utilize the identification set to derive the mask  $m$  and construct the steering vector  $\mathbf{v}$ , while reporting performance on the evaluation set. Following the setup in §2, we utilize the identification set to derive the mask  $m$  and construct the steering vector  $\mathbf{v}$ , while reporting performance on the evaluation set. To construct the domain-specific steering vector  $\mathbf{v}_{F \rightarrow T}$ , we partition the identification set based on the model’s prediction correctness. Specifically, we define  $\mathcal{D}_+$  as the set of prompts where the model generated the correct answer, and  $\mathcal{D}_-$  as the set where it generated an incorrect answer. Consequently, the vector  $\mathbf{v}_{F \rightarrow T}$  represents a directional shift from incorrect to correct reasoning within that specific domain. We compare our approach against two baselines: (i) *Standard*, the original model performance without intervention ( $\alpha = 0$ ); (ii) *Whole-Dimension Steering (WDS)*, activation steering using the entire dimensional space. For the evaluation, we use LM-evaluation-harness (Gao et al., 2023) and report exact-match (EM) accuracy under greedy decoding for deterministic evaluation. We report the best accuracy achieved for each method over all hyperparameters (see §B.3 for experimental details).

**Results.** Figure 6 reports subject-wise accuracy changes by CDS and WDS, compared to the standard baseline without steering. On average, the proposed CDS achieves 3.09% improvement in accuracy while WDS shows 1.51% improvement. Consequently, CDS outperforms WDS in 34 out of 57 subjects, while WDS is better in 5 cases; the remaining 18 subjects are tied. Remarkably, CDS yields larger performance gains in STEM where precise knowledge retrieval is crucial.<sup>3</sup> These results are consistent with our analysis in §3, which observed specific dimensions encoding domain-specific concepts like biological terms. For instance, in *high school biology*, CDS achieves an accuracy increase of 6%, whereas WDS yields no improvement. We also observe significant gains in medical subject such as *clinical knowledge* by

<sup>3</sup>We reclassify 6 medical subjects from MMLU’s *Other* category to *STEM* to better reflect their scientific nature, supporting clearer interpretation and aligning with MMLU-Global (Singh et al., 2025), which similarly isolates medical subjects.




Figure 6: Performance comparison between CDS and WDS baseline across MMLU subjects. The bars represent the accuracy difference from the unsteered standard baseline. The numbers on subject name denote CDS improvement ( $\blacktriangle/\nabla$ ) over WDS. Subjects are divided into categories: **STEM**, **Humanities**, **Social Sciences**, and **Other**. Subjects with identical scores (total of 8) across all methods are excluded.

+12%, significantly surpassing WDS of +6%.

CDS also effectively mitigates the negative transfer often observed with naive steering vectors, a trend particularly evident in logical and mathematical reasoning tasks. As shown in *formal logic* and *high school mathematics*, applying WDS results in performance degradation of -8% and -2% respectively. In contrast, by masking out non-critical dimensions, CDS prevents this interference, recovering performance to the standard baseline level in *formal logic* and even achieving a positive gain of +2% in *high school mathematics*. Overall, these results confirm that CDS allows for targeted enhancement of domain-specific capabilities without disrupting the model’s general capability. The results on other models are provided in §A.10.

## 4.2 Experiment on Safety Jailbreaking

**Setup.** We further evaluate the efficacy of CDS in the context of *LLM jailbreaking* (Wei et al., 2023a) using the AdvBench dataset (Zou et al., 2023b). Jailbreaking aims to steer the model from a refusal state to a compliant state when facing harmful queries, and we specifically focus on the control of refusal mechanism using activation steering. To derive the steering vector  $\mathbf{v}$ , we construct  $\mathcal{D}_+$  and  $\mathcal{D}_-$  using pairs of compliant and refusal behavioral suffixes, respectively, and compute the mean difference in hidden states following Lee et al. (2024).<sup>4</sup> For the dimension mask  $m$ , we identify harmfulness-relevant critical dimensions by utilizing 260 queries and corresponding target responses provided in AdvBench. We compare CDS against

<sup>4</sup>We utilize a set of behavioral compliance and refusal suffixes adopted from the IBM Activation Steering repository [https://github.com/IBM/activation-steering/blob/main/docs/demo-data/behavior\\_refusal.json](https://github.com/IBM/activation-steering/blob/main/docs/demo-data/behavior_refusal.json)

two baselines: (i) *WDS*, same as §4.1, and (ii) *Random*, where the mask  $m$  is constructed by selecting the same number of  $k$  dimensions at random. The primary metrics are (1) Attack Success Rate (ASR) and (2) Text Quality measured on a 1–5 scale, where both are evaluated with GPT-4o-mini as LLM-as-judge. We report the best performance attained by each method across the hyperparameter sweep (see §B.4 for experimental details).

**Results.** Figure 7 presents the comparative performance of CDS against WDS and random baselines. The results provide empirical evidence that the refusal mechanism is mediated by a sparse set of critical dimensions, which CDS effectively isolates and manipulates. In Figure 7a, CDS demonstrates significantly higher steering efficiency compared to the random baseline; as steering strength  $\alpha$  increases, CDS rapidly achieves a high ASR, reaching a peak of 92% at  $\alpha = 10$ . In contrast, random baseline requires higher strength to achieve even moderate success rate below 60% and fails to induce meaningful behavioral shifts. This divergence confirms that the identified critical dimensions are functionally central to the model’s safety guardrails, whereas random subspaces have negligible influence. Notably, CDS achieves a higher peak ASR (92%) than that of WDS (84%), despite manipulating only 13% of the total dimensions.

Figure 7b further illustrates the trade-off between ASR and text quality under activation steering. As the force to artificially modify the model internal for jailbreaking is increased ( $\alpha \uparrow$ ), text quality is decreasing in general. Nevertheless, each method exhibits different behavior. For example, the random baseline highlights the inefficiency of uninformed steering. It consumes the text quality




Figure 7: Attack Success Rate (ASR) across steering strength  $\alpha$  and text quality. (a) Evaluation of ASR as a function of steering intensity  $\alpha$ . (b) Trade-off between ASR and the generated text quality. As  $\alpha$  increases, the performance trajectory moves from right to left, indicating the cost in text quality at higher steering intensities.

budget without yielding proportional gains in ASR, remaining at the bottom-left of the performance frontier. Importantly, CDS surpasses the saturation point of WDS, while still maintaining text quality scores above 2. While WDS achieves high ASR in certain regions (*e.g.*, near text quality 2.6), it suffers from more rapid performance degradation compared to CDS which maintains higher ASR at equivalent text quality levels. This confirms that steering with CDS using critical dimensions allows for a more accurate intervention that breaks refusal barriers effectively, compared to the baselines.

## 5 Related Works

**Interpretability units.** As LLMs have grown enormously, interpretability of their internal states remain a crucial research direction (Clark et al., 2019; Meng et al., 2022). Prior works have proposed various interpretability units for tracing concepts, including MLP layers (Geva et al., 2021), attention heads (Voita et al., 2019; Abnar and Zuidema, 2020), neurons (Vig et al., 2020; Dai et al., 2022), and features of probing classifiers (Ettinger et al., 2016; Li et al., 2023). More recently, sparse autoencoders have been used to address the polysemanticity of internal representations, decomposing hidden representations into sparse combinations of monosemantic features (Bricken et al., 2023; Yun et al., 2021; Cunningham et al., 2023). While these sophisticated units provide high precision and clear functional decomposition, they often entail substantial training overhead. In contrast, our approach treats individual hidden-state dimensions as interpretability units. By sacrificing the fine-grained granularity of SAEs for computational efficiency and model integrity, we provide a training-free means to identify domain-relevant signals directly within the original model.

**Massive activations.** Values flowing through LLMs include entries that deviate significantly from the average, referred to as outliers (Kovaleva et al., 2021; Dettmers et al., 2022). These massive values appear systematically across attention scores, weights, and activations (An et al., 2025). Sun et al. (2024) analyzed the role of massive activations as attention biases, while (Zhang et al., 2024) focused on weight outliers in LayerNorm layers, demonstrating their importance for model behavior. Prior works mainly treated massive activations as artifacts to be suppressed or regularized using techniques such as pruning or quantization (Xiao et al., 2023; Wei et al., 2023b; Yao et al., 2022). While several studies explore the structural alignment of massive activations with specific dimensions (Elhage et al., 2023) or view them as task-specific linear classifiers (Rudman et al., 2023), their intricate semantic functions remain underexplored. In contrast, we focus on interpretable attributes of massive activation, highlighting their semantic role and interpretable view.

**Activation steering.** Activation steering refers to the direct manipulation of internal activations to control model behavior (Rimsky et al., 2024; Zou et al., 2023a). Existing approaches vary by their intervention sites and strategies to optimize this control. Li et al. (2023) focused on manipulating attention heads to enhance truthfulness, whereas Turner et al. (2023) demonstrated that intervening on the residual stream effectively governs sentiment and toxicity. Advancing this paradigm, Lee et al. (2024) introduced conditional activation steering, which dynamically applies interventions based on specific conditions. We distinguish our approach by selectively steering the highest activated dimensions for each domain, focusing on the key hidden state components responsible for domain-specific control.

## 6 Conclusion

In this work, we focus on the inherent anisotropy of LLM internal activations, demonstrating that the disproportionate utilization of the embedding space leads to the emergence of specialized dimensions that capture domain-specific features. Specifically, we propose a simple statistical way to identify Domain-Critical Dimensions directly from the pre-trained weights without additional training, and verify their semantic interpretability with comprehensive analyses. We additionally show that targeted interventions with those dimensions can match or even surpass the performance of whole-space steering. We believe our findings offer a new perspective on navigating the complex internal landscape of LLMs, suggesting direction for more transparent and controllable AI systems.

### Limitations

While our results have shown that a sparse set of dominant dimensions can provide a practical control interface, we first acknowledge the inherent polysemanticity within LLM representations. Interpretability outcomes depend on the granularity at which features are defined, and approaches aimed at explicit fine-grained factorization can offer a complementary view of the latent space. However, we would like to clarify that our primary goal is not maximal feature separability, but investigating whether human-interpretable feature patterns are naturally embedded in the model’s intrinsic representational dimensions to utilize them for effective model control.

Second, the current implementation of Critical Dimension Steering (CDS) intervenes through a sparse subset of the representational space, which may not reflect every subtle, context-dependent variation in activation. However, by restricting interventions to a minimal set of critical dimensions, CDS remains computationally lightweight relative to approaches that steer across the full space. Furthermore, focusing on magnitude-dominant dimensions provides a stable control signal in our experiments and reduces exposure to less consequential variation.

Lastly, our method can be limited in tasks where domain boundaries are ambiguous. For domains with distinct characteristic signatures, our method successfully captures meaningful signals as demonstrated by our experiments. But in tasks like creative writing or general summarization, seman-

tic features heavily overlap and evaluation is inherently subjective, making it difficult to isolate domain-critical dimensions. Thus, an important direction for future work is to explore adaptive or context-aware mechanisms for selecting dimensions and steering strengths, which may further improve the specificity of sparse interventions while preserving their efficiency.

### Broader Impact and Ethical Implications

As LLMs increasingly influence decisions and public information ecosystems, alignment with human values is essential for ensuring trustworthy behavior, accountability, and user safety. Aligning LLMs with human values requires a deep understanding of how information is processed within their internal activations. While our goal is to contribute to more ethical and safety-oriented LLM research, our study includes experiments that uses harmful queries, including settings that may increase compliance with harmful requests, and this introduces a potential risk if misused. We therefore emphasize that these experiments are conducted for controlled analysis of internal mechanisms and are intended to improve model transparency and controllability, enabling more reliable safeguards and more trustworthy outputs. By tracing specific activation patterns and intervening directly in internal representations, our research provides a framework for enhancing the safety and factuality of AI outputs. We wish for this study to contribute to the development of reliable AI technologies that are transparently steered to remain beneficial and ethically grounded for all users.

### Acknowledgments

Youngji and Jaehyung are affiliated with the Department of Artificial Intelligence at Yonsei University. This research was supported in part by Institute for Information & communications Technology Planning & Evaluation (IITP) grant funded by the Korea government (MSIT) (No. RS-2020-II201361, Artificial Intelligence Graduate School Program (Yonsei University); No. RS-2025-02215344, Development of AI Technology with Robust and Flexible Resilience Against Risk Factors; No. RS-2025-25442405, Development of a Self-Learning World Model-Based AGI System for Hyperspectral Imaging).

## References

- Samira Abnar and Willem Zuidema. 2020. Quantifying attention flow in transformers. In *Annual Meeting of the Association for Computational Linguistics (ACL)*.
- Yongqi An, Xu Zhao, Tao Yu, Ming Tang, and Jinqiao Wang. 2025. Systematic outliers in large language models. In *International Conference on Learning Representations (ICLR)*.
- Trenton Bricken, Adly Templeton, Joshua Batson, Brian Chen, Adam Jermy, Tom Conerly, Nicholas L Turner, Cem Anil, Carson Denison, Amanda Askell, Robert Lasenby, Yifan Wu, Shauna Kravec, Nicholas Schiefer, Tim Maxwell, Nicholas anthony Joseph, Alex Tamkin, Karina Nguyen, Brayden McLean, and 5 others. 2023. [Towards monosemanticity: Decomposing language models with dictionary learning](#). *Transformer Circuits Thread*.
- Collin Burns, Haotian Ye, Dan Klein, and Jacob Steinhardt. 2023. Discovering latent knowledge in language models without supervision. In *International Conference on Learning Representations (ICLR)*.
- Kevin Clark, Urvashi Khandelwal, Omer Levy, and Christopher D Manning. 2019. What does bert look at? an analysis of bert’s attention. In *Proceedings of the 2019 ACL Workshop BlackboxNLP: Analyzing and Interpreting Neural Networks for NLP*.
- Hoagy Cunningham, Aidan Ewart, Logan Riggs, Robert Huben, and Lee Sharkey. 2023. Sparse autoencoders find highly interpretable features in language models. In *International Conference on Learning Representations (ICLR)*.
- Damai Dai, Li Dong, Yaru Hao, Zhifang Sui, Baobao Chang, and Furu Wei. 2022. Knowledge neurons in pretrained transformers. In *Annual Meeting of the Association for Computational Linguistics (ACL)*, pages 8493–8502.
- Tim Dettmers, Mike Lewis, Younes Belkada, and Luke Zettlemoyer. 2022. Llm.int8(): 8-bit matrix multiplication for transformers at scale. In *Advances in Neural Information Processing Systems (NeurIPS)*, volume 35, pages 30318–30332.
- Yue Ding, Karolis Martinkus, Damian Pascual, Simon Clematide, and Roger Wattenhofer. 2022. On isotropy calibration of transformer models. In *Proceedings of the Third Workshop on Insights from Negative Results in NLP*, pages 1–9.
- Nelson Elhage, Robert Lasenby, and Christopher Olah. 2023. [Privileged bases in the transformer residual stream](#). *Transformer Circuits Thread*.
- Kawin Ethayarajh. 2019. [How contextual are contextualized word representations? comparing the geometry of bert, elmo, and gpt-2 embeddings](#). In *Conference on Empirical Methods in Natural Language Processing*.
- Allyson Ettinger, Ahmed Elgohary, and Philip Resnik. 2016. Probing for semantic evidence of composition by means of simple classification tasks. In *Proceedings of the 1st workshop on evaluating vector-space representations for nlp*, pages 134–139.
- Jun Gao, Di He, Xu Tan, Tao Qin, Liwei Wang, and Tie-Yan Liu. 2019. Representation degeneration problem in training natural language generation models. In *International Conference on Learning Representations (ICLR)*.
- Leo Gao, Jonathan Tow, Baber Abbasi, Stella Biderman, Sid Black, Anthony DiPofi, Charles Foster, Laurence Golding, Jeffrey Hsu, Alain Le Noac’h, Haonan Li, Kyle McDonell, Niklas Muennighoff, Chris Ociepa, Jason Phang, Laria Reynolds, Hailey Schoelkopf, Aviya Skowron, Lintang Sutawika, and 5 others. 2023. [A framework for few-shot language model evaluation](#).
- Mor Geva, Roei Schuster, Jonathan Berant, and Omer Levy. 2021. Transformer feed-forward layers are key-value memories. In *Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 5484–5495.
- Aaron Grattafiori, Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Alex Vaughan, and 1 others. 2024. The llama 3 herd of models. *arXiv preprint arXiv:2407.21783*.
- Katharina Hämmerl, Alina Fastowski, Jindřich Libovický, and Alexander Fraser. 2023. Exploring anisotropy and outliers in multilingual language models for cross-lingual semantic sentence similarity. In *Annual Meeting of the Association for Computational Linguistics (ACL)*.
- Bobby He, Lorenzo Noci, Daniele Paliotta, Imanol Schlag, and Thomas Hofmann. 2024. Understanding and minimising outlier features in neural network training. In *Advances in Neural Information Processing Systems (NeurIPS)*.
- Dan Hendrycks, Collin Burns, Steven Basart, Andy Zou, Mantas Mazeika, Dawn Song, and Jacob Steinhardt. 2021. Measuring massive multitask language understanding. In *International Conference on Learning Representations (ICLR)*.
- Olga Kovaleva, Saurabh Kulshreshtha, Anna Rogers, and Anna Rumshisky. 2021. Bert busters: Outlier dimensions that disrupt transformers. In *Findings of the Association for Computational Linguistics: ACL-IJCNLP 2021*.
- Bruce W Lee, Inkit Padhi, Karthikeyan Natesan Ramamurthy, Erik Miehl, Pierre Dognin, Manish Nagireddy, and Amit Dhurandhar. 2024. Programming refusal with conditional activation steering. In *International Conference on Learning Representations (ICLR)*.

- Kenneth Li, Oam Patel, Fernanda Viégas, Hanspeter Pfister, and Martin Wattenberg. 2023. Inference-time intervention: Eliciting truthful answers from a language model. In *Advances in Neural Information Processing Systems (NeurIPS)*, volume 36, pages 41451–41530.
- Kevin Meng, David Bau, Alex Andonian, and Yonatan Belinkov. 2022. Locating and editing factual associations in gpt. In *Advances in Neural Information Processing Systems (NeurIPS)*, volume 35, pages 17359–17372.
- Team OLMo, Pete Walsh, Luca Soldaini, Dirk Groeneveld, Kyle Lo, Shane Arora, Akshita Bhagia, Yuling Gu, Shengyi Huang, Matt Jordan, and 1 others. 2024. 2 olmo 2 furious. *arXiv preprint arXiv:2501.00656*.
- Anton Razzhigaev, Matvey Mikhalechuk, Elizaveta Goncharova, Ivan Oseledets, Denis Dimitrov, and Andrey Kuznetsov. 2024. The shape of learning: Anisotropy and intrinsic dimensions in transformer-based models. In *Conference of the European Chapter of the Association for Computational Linguistics (EACL)*, pages 868–874.
- Nina Rimskey, Nick Gabrieli, Julian Schulz, Meg Tong, Evan Hubinger, and Alexander Turner. 2024. Steering llama 2 via contrastive activation addition. In *Annual Meeting of the Association for Computational Linguistics (ACL)*, pages 15504–15522.
- William Rudman, Catherine Chen, and Carsten Eickhoff. 2023. Outlier dimensions encode task specific knowledge. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, Singapore. Association for Computational Linguistics.
- William Rudman and Carsten Eickhoff. 2023. Stable anisotropic regularization. In *International Conference on Learning Representations (ICLR)*.
- Shivalika Singh, Angelika Romanou, Clémentine Fourrier, David Ifeoluwa Adelani, Jian Gang Ngui, Daniel Vila-Suero, Peerat Limkonchotiwat, Kelly Marchisio, Wei Qi Leong, Yosephine Susanto, and 1 others. 2025. Global mmlu: Understanding and addressing cultural and linguistic biases in multilingual evaluation. In *Annual Meeting of the Association for Computational Linguistics (ACL)*, pages 18761–18799.
- Mingjie Sun, Xinlei Chen, J Zico Kolter, and Zhuang Liu. 2024. Massive activations in large language models. In *Conference on Language Modeling (COLM)*.
- Gemma Team, Morgane Riviere, Shreya Pathak, Pier Giuseppe Sessa, Cassidy Hardin, Surya Bhupatiraju, Léonard Hussenot, Thomas Mesnard, Bobak Shahriari, Alexandre Ramé, and 1 others. 2024. Gemma 2: Improving open language models at a practical size. *arXiv preprint arXiv:2408.00118*.
- William Timkey and Marten van Schijndel. 2021. All bark and no bite: Rogue dimensions in transformer language models obscure representational quality. In *Conference on Empirical Methods in Natural Language Processing (EMNLP)*.
- Alexander Matt Turner, Lisa Thiergart, Gavin Leech, David Udell, Juan J Vazquez, Ulisse Mini, and Monte MacDiarmid. 2023. Steering language models with activation engineering. *arXiv preprint arXiv:2308.10248*.
- Jesse Vig, Sebastian Gehrmann, Yonatan Belinkov, Sharon Qian, Daniel Nevo, Yaron Singer, and Stuart Shieber. 2020. Investigating gender bias in language models using causal mediation analysis. In *Advances in Neural Information Processing Systems (NeurIPS)*, volume 33, pages 12388–12401.
- Elena Voita, David Talbot, Fedor Moiseev, Rico Senrich, and Ivan Titov. 2019. Analyzing multi-head self-attention: Specialized heads do the heavy lifting, the rest can be pruned. In *Annual Meeting of the Association for Computational Linguistics (ACL)*.
- Alexander Wei, Nika Haghtalab, and Jacob Steinhardt. 2023a. Jailbroken: How does llm safety training fail? In *Advances in Neural Information Processing Systems (NeurIPS)*.
- Xiuying Wei, Yunchen Zhang, Yuhang Li, Xiangguo Zhang, Ruihao Gong, Jinyang Guo, and Xianglong Liu. 2023b. Outlier suppression+: Accurate quantization of large language models by equivalent and optimal shifting and scaling. In *Conference on Empirical Methods in Natural Language Processing (EMNLP)*.
- Xiuying Wei, Yunchen Zhang, Xiangguo Zhang, Ruihao Gong, Shanghang Zhang, Qi Zhang, Fengwei Yu, and Xianglong Liu. 2022. Outlier suppression: Pushing the limit of low-bit transformer language models. In *Advances in Neural Information Processing Systems (NeurIPS)*, volume 35, pages 17402–17414.
- Guangxuan Xiao, Ji Lin, Mickael Seznec, Hao Wu, Julien Demouth, and Song Han. 2023. Smoothquant: Accurate and efficient post-training quantization for large language models. In *International conference on machine learning*, pages 38087–38099. PMLR.
- An Yang, Anfeng Li, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chang Gao, Chengen Huang, Chenxu Lv, and 1 others. 2025. Qwen3 technical report. *arXiv preprint arXiv:2505.09388*.
- Zhewei Yao, Reza Yazdani Aminabadi, Minjia Zhang, Xiaoxia Wu, Conglong Li, and Yuxiong He. 2022. Zeroquant: Efficient and affordable post-training quantization for large-scale transformers. In *Advances in Neural Information Processing Systems (NeurIPS)*.
- Zeyu Yun, Yubei Chen, Bruno Olshausen, and Yann LeCun. 2021. Transformer visualization via dictionary learning: contextualized embedding as a linear superposition of transformer factors. In *Proceedings of*

*Deep Learning Inside Out (DeeLIO): The 2nd Workshop on Knowledge Extraction and Integration for Deep Learning Architectures*, pages 1–10.

Zhihao Zhang, Jun Zhao, Qi Zhang, Tao Gui, and Xuan-Jing Huang. 2024. Unveiling linguistic regions in large language models. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 6228–6247.

Tianyang Zhao, Kunwar Yashraj Singh, Srikar Apalparaju, Peng Tang, Ying Nian Wu, and Li Erran Li. 2025. On the analysis and distillation of emergent outlier properties in pre-trained language models. In *Proceedings of the 2025 Conference of the Nations of the Americas Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, pages 8475–8507.

Zhanxing Zhu, Jingfeng Wu, Bing Yu, Lei Wu, and Jinwen Ma. 2018. The anisotropic noise in stochastic gradient descent: Its behavior of escaping from sharp minima and regularization effects. In *Proceedings of the International Conference on Machine Learning (ICML)*.

Andy Zou, Long Phan, Sarah Chen, James Campbell, Phillip Guo, Richard Ren, Alexander Pan, Xuwang Yin, Mantas Mazeika, Ann-Kathrin Dombrowski, and 1 others. 2023a. Representation engineering: A top-down approach to ai transparency. *arXiv preprint arXiv:2310.01405*.

Andy Zou, Zifan Wang, Nicholas Carlini, Milad Nasr, J Zico Kolter, and Matt Fredrikson. 2023b. Universal and transferable adversarial attacks on aligned language models. *arXiv preprint arXiv:2307.15043*.

## A Additional Analyses

In this section, we provide additional analysis results.

### A.1 Single Dimension Masking effect in Other Domains

We extend the analysis from §2.1 to all 57 MMLU subjects to verify the universality of functional criticality. Figure 9 presents the rank-wise accuracy drop averaged across all 57 subjects. Additionally, Figure 11 illustrates the distributions for 12 randomly selected subjects due to space limitations. The aggregated view confirms a distinct sharp head and long tail distribution. Specifically, only 12 out of 2304 average dimensions yield an accuracy drop of at least 4%, with the top rank causing a 14.6% drop. In contrast, the vast majority of dimensions have negligible impact, falling within narrow bands: 168 in  $[-4, -2)$ , 1184 in  $[-2, 0)$ , 872 in  $[0, 2)$ , and 38 in  $\geq 2\%$ . These results confirm that performance in every domain relies on a sparse set of critical dimensions, regardless of the subject.

### A.2 Domain Separation in Representational Space

We performed a qualitative analysis using 14 high school-related MMLU subjects, to explore the internal representations of different subjects. We extracted the hidden states from the identification set queries and applied t-Distributed Stochastic Neighbor Embedding (t-SNE) with a cosine metric for dimensionality reduction. Figure 8 illustrates the representations across different depths of the model (Layers 2, 8, and 19). We observe that the hidden states of each subject naturally form distinct subject-specific clusters, continued throughout the layers. This demonstrates that the model intrinsically encodes domain identity within its activation space.




Figure 8: Visualization of hidden state representations for 14 high school subjects. The dimensionality reduction was performed using t-SNE with cosine metric.

### A.3 Validation Between $\mathcal{I}_{\text{dcd}}$ and Functionally Critical Dimensions

**Validation by recall.** For each domain, we define the ground-truth set of functionally critical dimensions  $\mathcal{I}_{\text{mask}}^{(k)}$  using the top-ranked dimensions from the masking experiments up to a rank cutoff  $N$ . We measure recall as the overlap ratio  $|\mathcal{I}_{\text{dcd}} \cap \mathcal{I}_{\text{mask}}^{(k)}|/k$ . Note that to ensure strict validity, we exclude dimensions tied at the cutoff boundary, so the size of ground-truth set  $k$  may be slightly smaller than  $N$ . With  $N = 10$ , our method achieves an subject average recall of 89.39%, demonstrating that simple activation statistics effectively recover functionally critical dimensions. As shown in Table 2, the recall naturally decreases as the cutoff  $N$  increases due to the expansion of the ground-truth set.

Table 2: Recall of the top-100 statistical dimensions across varying rank cutoffs ( $N$ ).

Rank Cutoff ( $N$ )	5	10	15	20
Recall@ $K$ (%)	94.91	89.39	82.69	78.88

**Validation by masking effect.** We further validate  $\mathcal{I}_{\text{dcd}}$  by checking accuracy drop when masking them. The degradation trajectory shown in Table 3 and Figure 10 closely mirrors the ground-truth masking effects observed in Table 1 and Figure 9. Masking rank-1  $\mathcal{I}_{\text{dcd}}$  yields an 11.40% drop, while the impact decays rapidly for lower ranks. This confirms that  $s_j$  effectively isolates the sparse set of functionally critical dimensions.

Table 3: Impact of masking individual dimensions in  $\mathcal{I}_{\text{dcd}}$ . Rank  $k$  indicates the average accuracy when masking the  $k$ -th dimension.

Model	Rank 1	Rank 2	Rank 5	Rank 10	Rank 100
<b>Gemma-2-2B-IT</b>	45.13	51.30	53.16	53.16	55.93
Acc. Drop (%)	<b>(-11.40)</b>	(-5.23)	(-3.37)	(-3.37)	(-0.60)
<b>Qwen-3-8B</b>	35.75	23.39	73.12	72.21	72.57
Acc. Drop (%)	<b>(-37.55)</b>	<b>(-49.91)</b>	(-0.18)	(-1.09)	(-0.73)




Figure 9: Rankwise average accuracy drop of *function-critical dimensions*.




Figure 10: Rankwise average accuracy drop of our *domain-critical dimensions*.















Figure 11: Single dimension masking results across various MMLU subjects. Individual subjects exhibit a sparse set of critical dimensions that significantly impact performance.

#### A.4 Statistics of Domain-Critical Dimensions

**Statistical extremity.** In §2.2, we hypothesized that functionally critical dimensions are characterized by extreme activation values. To quantify this extremity for identified  $\mathcal{I}_{\text{dcd}}$ , we analyze the activation magnitude distribution of identified dimensions in the *high school biology*. Figure 12 and Table 4 presents the average magnitude of the top- $k$  identified dimensions alongside the individual magnitude at specific ranks. The results demonstrate distinct magnitude disparities of the the identified dimensions. The top-1 dimension exhibits a magnitude of 115.16, which is approximately  $100\times$  larger than the total dimension average of 1.17. Even within the top-100 dimensions (our selection for  $\mathcal{I}_{\text{dcd}}$ ), the average magnitude remains at 5.96, approximately  $5\times$  the baseline.

Table 4: Magnitude statistics of domain-critical dimensions. We report top- $k$  dimension set average and rank- $k$  individual magnitudes for the *High School Biology* subject.

$\mathcal{I}_{\text{dcd}}$ Rank ( $k$ )	1	5	10	50	100	200	300	400	500
Top- $k$ Avg.	115.16	44.83	28.47	9.71	5.96	3.81	3.01	2.59	2.31
Rank $k$ Single	115.16	16.96	10.97	2.95	1.88	1.51	1.35	1.26	1.19

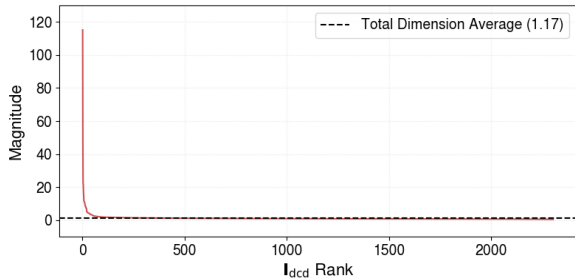


Figure 12: Distribution of activation magnitudes. A small subset of dimensions exhibits significantly higher magnitudes compared to the total dimension average.

**Subject specificity on masking.** We conducted experiments to confirm whether the masking effects of critical dimensions are indeed subject-specific. Instead of solely ablating dimensions based on their importance rank within the target task, we ablated dimensions of an equivalent rank extracted from a randomly selected different task. As shown in the table below, removing a target task-specific critical dimension causes a significantly larger performance drop than masking an equally ranked dimension from a random task. This confirms that the impact of these ablations is task-specific, rather than a general degradation from masking any highly activated dimension.

Table 5: Comparison of accuracy drop between task-specific and task-random dimensions.

Metric	Rank 3	Rank 5	Rank 10	Rank 15	Rank 20
<b>Task-Specific (Ours)</b>	51.30	53.16	53.16	55.37	55.44
Accuracy Drop (%)	(-5.23)	(-3.37)	(-3.37)	(-1.16)	(-1.09)
<b>Task-Random (Control)</b>	54.67	54.64	54.78	55.79	56.32
Accuracy Drop (%)	(-1.86)	(-1.89)	(-1.75)	(-0.74)	(-0.21)

**Subject specificity on variances.** We measured the variances and cross-task ranges of domain-critical dimensions. As shown, the Math-specific dimension (Dim 1046) exhibits significantly higher variance within the *high school mathematics* subject compared to other domains. The Biology and CS-critical dimensions demonstrate similar domain-specific variance patterns.

Table 6: Variances of domain-critical dimensions across different datasets.

Domain	ID	HS Math	HS Bio	HS Computer Sci.
Math	1046	<b>92.03</b>	59.70	73.47
Bio	2106	33.50	<b>80.48</b>	56.06
CS	1807	56.72	57.07	<b>105.10</b>

The following are multiple choice questions (with answers) about high school biology.

You must respond with a single alphabet character.

Energy is harvested during cellular respiration in stages. Which of the following correctly states which phase of cellular respiration harvests the most energy and the correct explanation why?

- A. The most energy is released during the Krebs cycle because it is here that pyruvate is completely broken down into CO<sub>2</sub>.
- B. The most energy is released during the Krebs cycle because in addition to the production of ATP, both FADH<sub>2</sub> and NADH are produced. Each of those molecules will release 2 ATPs and 3 ATPs, respectively.
- C. The most energy is released during oxidative phosphorylation because in addition to the phosphorylation of ADP into ATP, all the potential energy held in NADH and FADH is transferred to ATP.
- D. The most energy is released during oxidative phosphorylation because H<sub>2</sub>O is completely broken down into H<sup>+</sup> and O<sub>2</sub>.

Answer:

Dimension	334	1393	1645	1068	1570	535	2106	1807	1824	682
Attribute	Activation Value	Top Activating Token	Activation Value	Top Activating Token	Activation Value	Top Activating Token	Activation Value	Top Activating Token	Activation Value	Top Activating Token
Rank										
1	2043.9	<bos>	40.6	<bos>	52.3	Answer	64.8	_the	51.3	_phase
2	325.6	The	-11.3	?	49.4	_questions	63.9	_the	49.4	_energy
3	147.3	_biology	-17.4	?	47.0	_with	62.2	_is	49.3	_production
4	138.8	_Krebs	-18.2	.	45.3	_must	60.7	_to	48.3	_energy
5	138.1	_Krebs	-22.5	.	44.0	_about	59.8	_to	47.6	_harvested
6	137.4	_cellular	-24.5	.	43.5	)	58.5	_the	47.0	_molecules
7	131.3	_school	-28.3	D	40.8	:	57.1	_the	46.3	_released
8	130.4	_Krebs	-29.0	_alphabet	40.8	with	55.9	_that	46.0	_release
9	130.0	_harvested	-29.7	.	40.3	_cellular	55.0	_in	45.4	_potential
10	130.0	_stages	-30.4	_character	40.0	_answers	54.8	_The	44.9	_held

The following are multiple choice questions (with answers) about high school biology.

You must respond with a single alphabet character.

Which of the following statements is not correct about lipids?

- A. Lipids consist of fatty acids and glycerol.
- B. Steroids are examples of lipids.
- C. The molecules of a saturated fatty acid are packed so close together that they form a solid at room temperature.
- D. The head of a phospholipid is hydrophobic, and the tails are hydrophilic.

Answer:

Dimension	334	1393	1645	1068	1570	535	2106	1807	1824	682
Attribute	Activation Value	Top Activating Token	Activation Value	Top Activating Token	Activation Value	Top Activating Token	Activation Value	Top Activating Token	Activation Value	Top Activating Token
Rank										
1	2043.9	<bos>	40.6	<bos>	52.4	Answer	59.5	_a	49.2	_head
2	325.6	The	-18.2	.	49.4	_questions	59.5	_the	48.0	_close
3	148.3	_lipids	-25.7	.	47.0	_with	58.2	_The	47.7	_packed
4	147.3	_biology	-25.9	.	45.3	_must	57.6	_a	46.2	_solid
5	131.3	_school	-29.0	_alphabet	44.0	_about	57.4	_of	42.5	_school
6	123.7	_hydrophobic	-30.4	_character	43.5	)	56.7	_they	42.5	_molecules
7	120.8	_the	-30.8	_and	40.8	with	56.3	_The	39.6	_tails
8	118.0	_questions	-30.8	D	40.0	_answers	55.7	_are	39.0	_room
9	116.0	?	-33.2	.	39.8	.	54.4	_of	38.4	_answers
10	115.7	_about	-33.9	\n	39.0	_respond	53.2	_that	36.6	_acid

Figure 13: Analysis of domain-critical dimensions on high school biology prompts. (a) and (b) display the top-10 activating tokens for a *cellular respiration* related query and a *lipid* related query.

### A.5 Additional Prompt-level Analysis

To further extend the qualitative findings in §3, we provide an extended prompt-level analysis on high school biology and high school mathematics. These results consistently demonstrate that the specialized semantic roles of the identified  $\mathcal{L}_{\text{dcd}}$  remain robust across diverse queries.

**High school biology.** We display the full list of the highest-10 activating tokens for two queries, a *cellular respiration* related query illustrated at Figure 4 and a *lipids* related query as additional example. Figure 13 visualizes the activation patterns for these queries. Comparing the two queries reveals that domain-critical dimensions exhibit specialized roles that are consistent across different biological topics:

- **Dimension 334** (Domain & Topic Marker): This dimension consistently identifies the broad domain and the specific topic. In both queries, it strongly activates on the token `_biology`, signaling the subject matter. Simultaneously, it adapts to the specific context, capturing `_Krebs`

and `_cellular` in the cellular respiration query, while shifting to `_lipids` and `_hydrophobic` in the lipids query.

- **Dimensions 1645** (Task Structure & Instruction): This dimension focus on the structural scaffolding of the multiple-choice task rather than biological content. It shows high activations for `Answer`, `_questions`, `_must`, effectively tracking the format and instructions of the prompt.
- **Dimension 1068** (Functional Syntax): This dimension consistently targets functional words. Regardless of the query content, the top activating tokens remain syntactic connectors such as `_the`, `_of`, `_is`, and `_to`.
- **Dimension 2106** (Biological Lexicon): As a fine-grained content extractor, this dimension targets specific biological entities. It transitions from activating `_ATP` and `_phosphorylation` in the cellular respiration query to `_fatty`, `_glycerol`, and `_acid` in the lipids query.

The following are multiple choice questions (with answers) about high school mathematics. You must respond with a single alphabet character.

Suppose 5 different integers are randomly chosen from between 20 and 69, inclusive. What is the probability that they each have a different tens digit?

- A.  $\frac{1}{4}$
- B.  $\frac{1}{3}$
- C.  $\frac{1000}{52969}$
- D.  $\frac{2500}{52969}$

Answer:

Dimension	334	1393	1645	1068	535	987	1570	1046	1824	1807										
Attribute	Activation Value	Top Activating Token	Activation Value	Top Activating Token	Activation Value	Top Activating Token	Activation Value	Top Activating Token	Activation Value	Top Activating Token										
Rank	2043.9	<bos>	40.6	<bos>	52.5	Answer	52.2	_a	932.1	<bos>	28.7	<bos>	46.8	_integers	8.8	0	44.3	_questions	33.5	Answer
2	325.6	The	-17.9	.	49.4	_questions	50.9	_a	12.9	_character	7.0	_mathematics	44.0	_tens	8.6	0	36.1	)	33.3	_answers
3	146.6	_mathematics	-28.8	_alphabet	45.7	_with	47.4	with	12.3	_that	6.3	The	42.6	_school	4.9	0	33.8	)	33.0	_questions
4	142.6	_probability	-29.7	_character	44.6	:	46.7	_must	11.0	_probability	5.5	39.1	_chosen	4.6	5	33.1	_with	29.9	)	_n
5	131.3	_school	-31.3	)	44.0	_about	46.7	_are	10.8	The	3.6	_inclusive	38.4	_answers	4.1	(	30.6	_	29.6	)
6	126.6	_integers	-31.9	D	43.5	)	45.5	_the	10.2	_What	3.3	_school	38.3	_different	3.4	)	30.4	6	28.6	)
7	124.3	_tens	-34.0	)	43.3	_must	45.3	_are	10.1	_is	1.4	_they	36.6	_different	2.7	6	30.1	_respond	27.2	You
8	118.0	_questions	-34.1	)	40.8	with	45.1	_that	9.2	_tens	-1.2	_high	35.5	_digit	2.4	2	30.1	_mathematics	26.5	)
9	117.0	_chosen	-34.4	_	40.0	_answers	44.3	_with	9.2	)	-2.3	_that	35.4	0	2.3	9	30.0	_choice	26.3	)
10	116.9	_digit	-34.6	)	38.8	_respond	42.9	_about	9.1	_have	-2.8	.	34.7	5	2.1	9	29.9	9	26.1	.

The following are multiple choice questions (with answers) about high school mathematics. You must respond with a single alphabet character.

An equilateral triangle has sides of 12 inches. What is the approximate area of the triangle?

- A. 62
- B. 72
- C. 84
- D. 112

Answer:

Dimension	334	1393	1645	1068	535	987	1570	1046	1824	1807										
Attribute	Activation Value	Top Activating Token	Activation Value	Top Activating Token	Activation Value	Top Activating Token	Activation Value	Top Activating Token	Activation Value	Top Activating Token										
Rank	2043.9	<bos>	40.6	<bos>	49.4	_questions	52.8	_the	932.1	<bos>	28.7	<bos>	43.7	_sides	0.4	2	44.3	_questions	33.8	Answer
2	325.6	The	-17.9	.	49.4	Answer	50.9	_a	13.1	_area	7.0	_mathematics	42.6	_school	-0.4	_	33.6	)	33.3	_answers
3	146.6	_mathematics	-28.0	_	45.7	_with	47.4	with	12.9	_character	6.3	The	39.6	_area	-1.6	_	33.1	_with	33.0	_questions
4	138.6	_area	-28.8	_alphabet	44.0	_about	46.7	_must	10.8	_of	5.5	.	38.4	_answers	-1.8	4	30.6	_	29.9	)
5	131.3	_school	-29.7	_character	43.5	)	46.7	_are	10.8	The	3.3	_school	37.4	_triangle	-2.1	1	30.1	_respond	29.6	)
6	125.0	_inches	-33.3	D	43.3	_must	46.4	_the	9.5	_is	0.2	_the	36.1	_inches	-2.8	1	30.1	_mathematics	28.6	)
7	124.2	_approximate	-33.4	D	40.8	with	46.2	_of	9.2	)	-1.2	_high	33.5	_mathematics	-4.8	2	30.0	_choice	27.2	You
8	118.0	_triangle	-34.3	?	40.0	_answers	44.3	_with	9.1	_questions	-3.2	You	33.2	_choice	-5.0	8	28.5	_must	26.4	)
9	117.0	_questions	-36.8	.	38.8	_respond	42.9	_about	9.1	)	-3.6	_about	31.8	2	-5.7	_	28.1	26.1	.	.
10	116.1	_equilateral	-37.1	C	38.6	:	40.6	_is	8.9	)	-6.0	)	31.3	_multiple	-6.3	_	28.1	_character	25.6	_choice

Figure 14: Analysis of domain-critical dimensions on high school mathematics prompts. (a) and (b) display the top-10 activating tokens for a *tens digit* related query and an *equilateral triangle* related query, respectively.

**High school mathematics.** We provide additional query samples sourced from high school mathematics. Figure 14 presents the highest-10 activating tokens for a *tens digit* related query (Figure 14a) and an *equilateral triangle* related query (Figure 14b). Similar to the observations in the biology domain, specific dimensions in the mathematics dataset demonstrate distinct semantic responsibilities:

- **Dimension 334** (Domain & Topic Marker): Just as it identified `_biology` in the previous section, dimension 334 consistently captures the domain marker `_mathematics` in both queries. Furthermore, it adapts to the specific query context, prioritizing `_tens`, `_digits` and `_integers` in the tens digit query, while shifting to `_equilateral`, `_triangle`, and `_area` in the equilateral triangle related query.
- **Dimension 1046** (Quantitative & Symbolic Content): This dimension specializes in numerical and quantitative concepts. In the tens digit query, it activates on specific numbers like `0`, `5` and `6`. In the equilateral triangle query, its focus sharpens on raw numerical values and symbolic

tokens, with top activations including `2`, `4`, and `1`. This confirms that dimension 1046 serves as a dedicated processor for quantitative information.

- **Structural & Functional Syntax Dimensions:** Consistent with the biology results, dimension 1645 continue to track the query structure (e.g., `Answer`, `_questions`), while dimension 1068 maintains its role in attending to functional words.

## A.6 Additional Dataset-level Analysis on Other Domains

Beyond the primary subjects discussed in §3, we extend our dataset-level analysis to *high school computer science* and *high school chemistry*. The following results shows that the semantic specialization of domain-critical dimensions is a universal phenomenon across varied scientific disciplines, not limited to biology or mathematics.

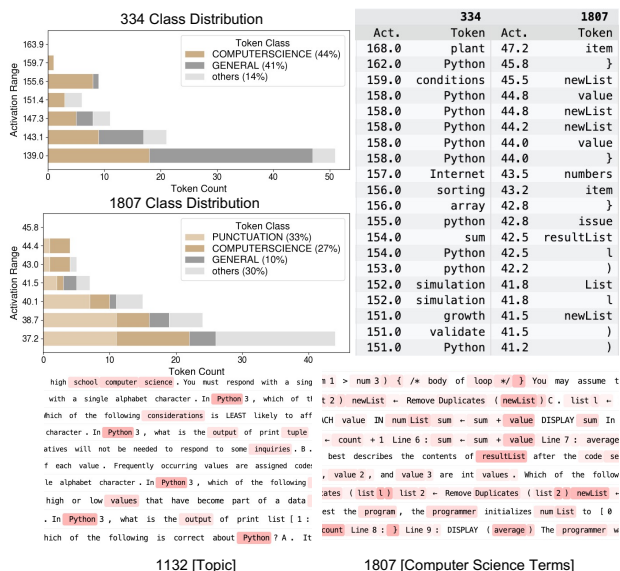


Figure 15: Analysis of  $\mathcal{I}_{dcd}$  in *high school computer science*. (Top-left) Token class distributions of dimension 334 (rank-1) and 1807 (rank-6). (Top-right) Highest activated tokens for dimension 334 and 1807. (Bottom) Activation heatmaps highlighting topic-level markers and specific computer science terms/syntax.

**High school computer science.** Figure 15 illustrates the semantic roles of dimension 334 and dimension 1807.

- **Dimension 334** acts as a high-level subject activator, consistent with its role in other STEM subjects mentioned in §3. 44% of its highly activated tokens belong to the `COMPUTERSCIENCE` class, including terms like `Python`, `_sorting`, and `_array`. The activation heatmap shows it consistently triggers on explicit domain indicators such as `computer science` and `Python`.
- **Dimension 1807** specializes in the structural and lexical components of programming, ranking 6 in high school computer science. 27% of activated tokens are CS-specific, while 33% are `PUNCTUATION`, reflecting the symbol-heavy nature of code. It responds strongly to variable-like tokens (`_newList`, `_resultList`) and structural symbols (`}`, `)`, `[`), effectively tracking the logic and syntax of computer science problems.



Figure 16: Analysis of  $\mathcal{I}_{dcd}$  in *high school chemistry*. (Top-left) Token class distributions of dimension 334 (rank-1) and 987 (rank-8). (Top-right) Highest activated tokens for dimension 334 and 987. (Bottom) Activation heatmaps highlighting domain-level markers and specific chemical terminology.

**High school chemistry.** Additionally, we analyze the *high school chemistry* dataset to confirm the functional roles of its domain-critical dimensions. Figure 16 illustrates the behavior of dimension 334 and dimension 987.

- **Dimension 334** serves as the primary critical dimension, as it acts as a robust indicator for the chemistry domain. 64% of its highly activated tokens are categorized as `CHEMISTRY`, with the token `_chemistry` itself appearing most frequently at the highest activation levels. The heatmap reveals that it consistently triggers on explicit subject markers like `school chemistry` within the prompt.
- **Dimension 987** ranked 8th in high school chemistry, specializing in specific chemical entities and substances. 41% of its activations correspond to the `CHEMISTRY` class, targeting concrete terms such as `_sulfuric`, `_iodine`, `_calcium`, and `_compounds`. The activation heatmap demonstrates that it precisely tracks fine-grained chemical concepts including `NH`, `Ag`, and `sulfuric acid`, facilitating the processing of domain-specific scientific content.

## A.7 Extension to Larger Scale Model: Gemma-2-9B-IT

To evaluate whether functional criticality scales with model size, we apply our identification framework to Gemma-2-9b-it. Our analysis reveals that larger models similarly shows a highly interpretable domain-critical dimensions, suggesting that the concentration of domain-specific information is a persistent property of scaled-up architectures.



Figure 17: High school biology  $\mathcal{I}_{dcd}$  in Gemma-2-9b-it. (Top-left) Token class distribution of dimension 3033 across biology and math datasets. (Top-right) Highly activated tokens for dimension 79 (rank-7) and 3033 (rank-8). (Bottom) Activation heatmaps of dimension 3033 and 79 on biology dataset prompts.

**High school Biology.** Figure 17 illustrates biology prominent dimensions, dimension 3033 and dimension 79.

- o **Dimension 3033** exhibits strong domain selectivity similar to the critical dimensions observed in the 2B model. As shown in Figure 17, when analyzed on the biology dataset, 64% of activated tokens are BIOLOGY entities, including core concepts such as `_ATP`, and `_homeostasis`. While in math dataset, it shifts to `GENERAL` (51%) and `VARIABLE` (18%) tokens. Heatmaps confirm activation on terms like `_metabolism` and `_NADPH`.
- o **dimension 79** acts as a specialized detector for biological terminology (e.g., `_lipids`, `_vaccination`). It tracks fine-grained entities such as `_lymphocytes` and `_mutations`, supporting the hypothesis that larger models also utilize interpretable dimensions.

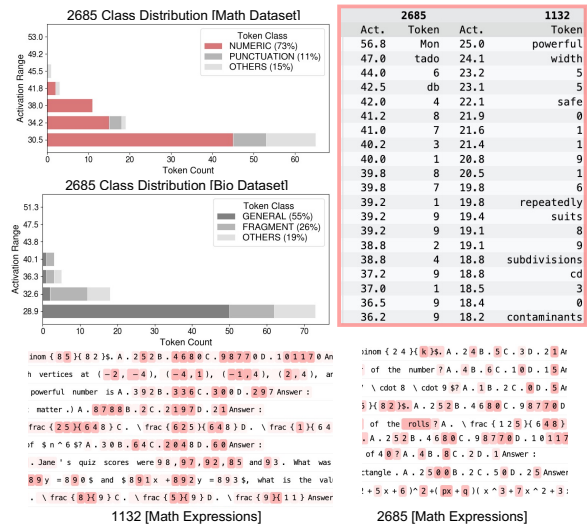


Figure 18: High school mathematics  $\mathcal{I}_{dcd}$  in Gemma-2-9b-it. (Top-left) Token class distribution of dimension 2685 across math and biology datasets. (Top-right) Highly activated tokens for dimension 2685 (rank-10) and 1132 (rank-9). (Bottom) Activation heatmaps of dimension 1132 and 2685 on mathematical prompts.

**High school Mathematics.** We further analyze the domain-critical dimensions identified from the *high school mathematics* dataset in the Gemma-2-9b-it model. Figure 18 illustrates the semantic roles of two specialized dimensions, dimension 2685 and dimension 1132.

- o **Dimension 2685** demonstrates a clear specialization in numerical content. As shown in the class distribution plots, when evaluated on the mathematics dataset, 73% of the highly activated tokens are classified as `NUMERIC`, with top tokens consisting of various digits such as 6, 4, and 8. In contrast, on the biology dataset, this domain-specific preference disappears, with `GENERAL` tokens accounting for the majority (55%). The activation heatmap confirms that dimension 2685 selectively targets numerical values and digits within mathematical expressions.
- o **Dimension 1132** focuses on the components of mathematical formulas and quantitative descriptions. The top activated tokens include specific digits, alongside quantitative terms such as `width`, and `subdivisions`. In the corresponding heatmap, dimension 1132 responds strongly to numbers embedded within LaTeX-style expressions, showing that it captures the quantitative elements necessary for solving mathematical problems.

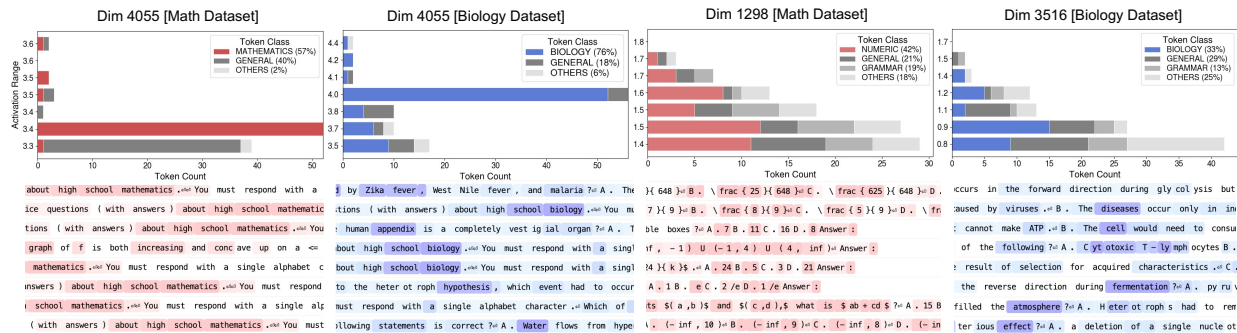


Figure 19:  $\mathcal{I}_{\text{dcd}}$  activation pattern in Llama-3.1-8b-it. Token analysis on domain-critical dimensions of *high school mathematics* and *high school biology*. (a) indicates token class distribution and activation of dimension 4055 (rank-5 in mathematics), (b) dimension 4055 (rank-5 in biology), (c) dimension 1298 (rank-14 in mathematics), and (d) dimension 3516 (rank-13 in biology).

### A.8 Extension to Other Family Model: Llama-3.1-8B-IT

To demonstrate that the interpretable features identified in our main analysis (§3) are not specific to a single model architecture, we extend our experiments to Llama-3.1-8b-it. This analysis aims to verify whether the functional roles of specific dimensions are preserved across different model families.

**Domain indicator dimension.** We observe dimension 4055 as a direct functional analogue to dimension 334 described in the main text. It ranks 5th in importance for both the mathematics and biology datasets, confirming its role as a global topic indicator. As visualized in Figure 19, this dimension strongly activates on high-level domain keywords (e.g., `_mathematics`, `_biology`). This consistency suggests that the mechanism of dedicating specialized dimensions to track the global context or topic is a generalizable property of LLMs, independent of the specific training run.

**Domain specialized dimension.** Beyond domain indicator dimension, Llama-3.1-8b-it exhibits specialized dimensions dedicated to fine-grained concepts, mirroring the behavior observed in our primary results in §3. As illustrated in Figure 19c, dimension 1298 (rank 14 in mathematics) focuses on the syntactic structure of mathematics. This functional role is quantitatively supported by its activation distribution, which is dominated by `NUMERIC (42%)` and `GRAMMAR (19%)` tokens such as LaTeX operators (e.g., `\frac`), effectively separating mathematical notation from natural language context. Similarly, Figure 19d demonstrates the role of dimension 3516 (rank 13 in biology) as a biology-specific entity detector. With `BIOLOGY`

tokens (33%) constituting the largest share, this dimension selectively highlights biology entities like `_cell` and `_fermentation`, showcasing the model’s capacity to isolate and process granular biological vocabulary.

## A.9 Additional CDS Results for §4.1

In §4.1, we reported the performance of Critical Dimension Steering (CDS) using the optimal number of domain-critical dimensions  $k$  determined for each subject. Given that each subject possesses distinct domain nuances and utilizes different steering vectors constructed from subject-specific identification sets, optimizing  $k$  for each domain is a natural approach to maximize control efficacy. However, to decouple the intrinsic effectiveness of our selection method from the benefits of hyperparameter optimization, we present an analysis of CDS performance with a fixed  $k$  across all 57 MMLU subjects. We compare this against the Whole Dimension Steering (WDS) baseline, which utilizes all dimensions ( $D = 2304$ ) and thus remains constant across these comparisons. Table 7 summarizes the average accuracy improvement for each setting.

**Performance scaling with  $k$ .** As shown in Table 7, CDS has a performance trend correlated with the dimension budget  $k$ . At the sparsest setting ( $k=100$ ), shown in Figure 20e, CDS achieves an average improvement of 1.47%, which is slightly below the WDS baseline of 1.51%. However, despite the lower average, CDS ( $k=100$ ) outperforms WDS in 20 subjects compared to 16 wins for WDS, suggesting that even a small subset of critical dimensions allows for more frequent success in precision-heavy tasks. As  $k$  increases, CDS consistently surpasses the WDS baseline. In Figure 20d ( $k=200$ ), the average improvement rises to 1.61%, crossing the WDS threshold. The performance gap widens further at  $k=500$  (Figure 20a), where CDS achieves a 2.07% average improvement. This indicates that while the most *critical* features are sparse, a broader support set, up to  $k=500$ , contributes to maximizing domain adaptation performance. Nevertheless, comparing the fixed  $k=500$  result (2.07%) with the result in §4.1 (3.09%) highlights that subject-specific sparsity tuning remains essential for optimal control.

**Mitigation of Negative Transfer.** A key advantage of CDS over WDS is the reduction of negative transfer, particularly in reasoning-heavy tasks. As visualized across Figure 20a–e, WDS consistently degrades performance in *formal logic* (-8%) and *high school mathematics* (-2%). CDS effectively mitigates this interference. In Figure 20a ( $k=500$ ), the degradation in *formal logic* is significantly reduced from -8% (WDS) to -2% (CDS). Further-

Table 7: Average accuracy improvement over standard baseline across 57 MMLU subjects. CDS is evaluated at fixed dimension numbers  $k$  while WDS uses all dimensions. The WDS average corresponds to 1.51%.

Dimension number $k$	100	200	300	400	500
CDS Average (%)	1.47	1.61	1.86	1.86	2.07
Subjects (CDS/WDS/Tie)	20/16/21	23/17/17	19/12/26	18/14/25	21/11/25

more, at lower sparsity levels such as  $k=100$  (Figure 20e) and  $k=200$  (Figure 20d), CDS recovers performance in *high school mathematics* to positive gains (+2%), whereas WDS remains negative. This confirms that by masking out non-critical dimensions, CDS preserves the model’s general reasoning capabilities better than whole-dimension steering, regardless of the specific  $k$  value.


## A.10 Extension to More Models

We have conducted additional experiments on other model families, OLMo-2-1B-Instruct (OLMo et al., 2024), Llama-3.1-8B-IT (Grattafiori et al., 2024), Qwen-3-8B (Yang et al., 2025), and Gemma-2-9B-IT (Team et al., 2024), presented in Table 9, 10, 11, 12. Our results show that CDS consistently outperforms the WDS baseline across the majority of MMLU subjects (e.g., 37 vs. 10 wins for OLMo-2-1B-Instruct, and 38 vs. 7 wins for Qwen-3-8B). This consistent effectiveness across different model families demonstrates the robustness and generalizability of our findings.


Table 8: Average accuracy improvement over baselines on various models.

Model	Base	WDS	CDS
OLMo-2-1B-Instruct	38.77	+2.70	+4.63
Llama-3.1-8B-IT	64.60	+1.56	+3.48
Qwen-3-8B	73.30	+1.47	<b>+2.42</b>
Gemma-2-9B-IT	70.56	-0.14	+1.40


(a) CDS ( $k = 500$ ) vs. WDS




(b) CDS ( $k = 400$ ) vs. WDS



(c) CDS ( $k = 300$ ) vs. WDS



(d) CDS ( $k = 200$ ) vs. WDS



(e) CDS ( $k = 100$ ) vs. WDS




Figure 20: Subject-wise accuracy difference for fixed  $k$  CDS settings.

Table 9: Performance of OLMo-2-1B-Instruct across all 57 MMLU subjects.

Subject	Base	WDS	CDS	Subject	Base	WDS	CDS	Subject	Base	WDS	CDS
Abstract Algebra	32.0	-2.0	+2.0	High School Chemistry	14.0	+12.0	+14.0	Management	36.0	+4.0	+2.0
Anatomy	40.0	+10.0	+6.0	High School Computer Science	38.0	+4.0	+12.0	Marketing	64.0	0.0	+4.0
Astronomy	32.0	+8.0	+8.0	High School European History	60.0	0.0	+2.0	Medical Genetics	48.0	-2.0	+6.0
Business Ethics	40.0	+2.0	+6.0	High School Geography	44.0	-2.0	0.0	Miscellaneous	46.0	+8.0	+8.0
Clinical Knowledge	36.0	+10.0	+14.0	High School Gov. and Politics	42.0	+4.0	0.0	Moral Disputes	36.0	+6.0	+6.0
College Biology	42.0	+4.0	+6.0	High School Macroeconomics	30.0	+2.0	0.0	Moral Scenarios	30.0	0.0	0.0
College Chemistry	34.0	-4.0	-2.0	High School Mathematics	26.0	+2.0	+8.0	Nutrition	48.0	+2.0	+4.0
College Computer Science	30.0	+2.0	+2.0	High School Microeconomics	36.0	+14.0	+4.0	Philosophy	38.0	0.0	+2.0
College Mathematics	28.0	+4.0	+8.0	High School Physics	28.0	+2.0	+6.0	Prehistory	48.0	+6.0	+8.0
College Medicine	30.0	+6.0	+10.0	High School Psychology	50.0	+8.0	+10.0	Professional Accounting	32.0	0.0	+6.0
College Physics	24.0	0.0	+8.0	High School Statistics	34.0	0.0	+2.0	Professional Law	30.0	+4.0	+6.0
Computer Security	52.0	+2.0	0.0	High School US History	48.0	+4.0	+10.0	Professional Medicine	28.0	0.0	+6.0
Conceptual Physics	30.0	+4.0	+6.0	High School World History	48.0	+2.0	+2.0	Professional Psychology	32.0	+2.0	+4.0
Econometrics	28.0	+6.0	+2.0	Human Aging	54.0	-2.0	+4.0	Public Relations	52.0	+8.0	+10.0
Electrical Engineering	58.0	-6.0	0.0	Human Sexuality	46.0	+6.0	+6.0	Security Studies	52.0	0.0	+2.0
Elementary Mathematics	36.0	-2.0	-4.0	International Law	50.0	-2.0	0.0	Sociology	60.0	0.0	+4.0
Formal Logic	26.0	0.0	0.0	Jurisprudence	48.0	+4.0	+2.0	US Foreign Policy	56.0	+2.0	+6.0
Global Facts	20.0	+2.0	+2.0	Logical Fallacies	38.0	0.0	+4.0	Virology	38.0	+2.0	+6.0
High School Biology	34.0	+14.0	+12.0	Machine Learning	40.0	0.0	+2.0	World Religions	46.0	+4.0	+4.0

Table 10: Performance of Llama-3.1-8B-IT across all 57 MMLU subjects.

Subject	Base	WDS	CDS	Subject	Base	WDS	CDS	Subject	Base	WDS	CDS
Abstract Algebra	38.0	-4.0	0.0	High School Chemistry	58.0	+4.0	+4.0	Management	68.0	0.0	0.0
Anatomy	60.0	0.0	+10.0	High School Computer Science	68.0	-2.0	-2.0	Marketing	88.0	+4.0	+4.0
Astronomy	60.0	+2.0	+2.0	High School European History	76.0	-4.0	+2.0	Medical Genetics	78.0	+2.0	+2.0
Business Ethics	60.0	-2.0	0.0	High School Geography	80.0	+6.0	+8.0	Miscellaneous	72.0	0.0	0.0
Clinical Knowledge	54.0	+12.0	+14.0	High School Gov. and Politics	78.0	0.0	0.0	Moral Disputes	78.0	0.0	+2.0
College Biology	74.0	-2.0	-2.0	High School Macroeconomics	72.0	-2.0	0.0	Moral Scenarios	54.0	0.0	+2.0
College Chemistry	48.0	0.0	0.0	High School Mathematics	42.0	+6.0	+4.0	Nutrition	74.0	-4.0	-4.0
College Computer Science	64.0	-4.0	0.0	High School Microeconomics	78.0	-2.0	-2.0	Philosophy	84.0	-4.0	0.0
College Mathematics	26.0	+10.0	+8.0	High School Physics	40.0	+8.0	+6.0	Prehistory	78.0	0.0	+4.0
College Medicine	70.0	-6.0	0.0	High School Psychology	90.0	+2.0	+2.0	Professional Accounting	56.0	+2.0	+2.0
College Physics	38.0	0.0	0.0	High School Statistics	56.0	+4.0	+2.0	Professional Law	54.0	-2.0	0.0
Computer Security	74.0	-2.0	+2.0	High School US History	82.0	+4.0	+2.0	Professional Medicine	60.0	0.0	0.0
Conceptual Physics	52.0	+2.0	+6.0	High School World History	76.0	-4.0	0.0	Professional Psychology	54.0	0.0	0.0
Econometrics	58.0	0.0	+4.0	Human Aging	58.0	+6.0	+8.0	Public Relations	66.0	-4.0	+2.0
Electrical Engineering	60.0	+4.0	+10.0	Human Sexuality	82.0	+2.0	+2.0	Security Studies	72.0	+10.0	+10.0
Elementary Mathematics	38.0	+2.0	+4.0	International Law	70.0	+4.0	+2.0	Sociology	84.0	-2.0	0.0
Formal Logic	42.0	0.0	+6.0	Jurisprudence	72.0	0.0	+6.0	US Foreign Policy	84.0	+2.0	+2.0
Global Facts	46.0	+4.0	+10.0	Logical Fallacies	70.0	+10.0	+10.0	Virology	54.0	0.0	+2.0
High School Biology	74.0	+2.0	+4.0	Machine Learning	38.0	+4.0	+8.0	World Religions	70.0	+10.0	+6.0

Table 11: Performance of Qwen-3-8B across all 57 MMLU subjects.

Subject	Base	WDS	CDS	Subject	Base	WDS	CDS	Subject	Base	WDS	CDS
Abstract Algebra	50.0	0.0	+8.0	High School Chemistry	66.0	+2.0	+4.0	Management	86.0	0.0	+2.0
Anatomy	72.0	0.0	+2.0	High School Computer Science	86.0	+2.0	0.0	Marketing	90.0	+2.0	+2.0
Astronomy	78.0	-2.0	0.0	High School European History	86.0	0.0	+2.0	Medical Genetics	76.0	-2.0	0.0
Business Ethics	54.0	+4.0	+6.0	High School Geography	90.0	0.0	0.0	Miscellaneous	80.0	-2.0	0.0
Clinical Knowledge	78.0	0.0	+2.0	High School Gov. and Politics	94.0	0.0	0.0	Moral Disputes	68.0	+2.0	+4.0
College Biology	86.0	0.0	0.0	High School Macroeconomics	84.0	0.0	-2.0	Moral Scenarios	40.0	0.0	+2.0
College Chemistry	60.0	+4.0	+4.0	High School Mathematics	50.0	+6.0	+10.0	Nutrition	72.0	0.0	0.0
College Computer Science	72.0	-2.0	0.0	High School Microeconomics	94.0	0.0	0.0	Philosophy	86.0	0.0	0.0
College Mathematics	54.0	+2.0	+4.0	High School Physics	68.0	-2.0	0.0	Prehistory	80.0	0.0	+2.0
College Medicine	66.0	-4.0	-2.0	High School Psychology	82.0	+2.0	+4.0	Professional Accounting	52.0	+2.0	+10.0
College Physics	56.0	+2.0	0.0	High School Statistics	84.0	0.0	+2.0	Professional Law	56.0	-2.0	0.0
Computer Security	80.0	+4.0	+2.0	High School US History	90.0	0.0	0.0	Professional Medicine	82.0	+4.0	+6.0
Conceptual Physics	74.0	+2.0	+4.0	High School World History	86.0	+2.0	+2.0	Professional Psychology	72.0	0.0	+2.0
Econometrics	76.0	+4.0	+6.0	Human Aging	60.0	+4.0	+6.0	Public Relations	66.0	-2.0	0.0
Electrical Engineering	72.0	+2.0	+4.0	Human Sexuality	84.0	0.0	+2.0	Security Studies	82.0	0.0	+2.0
Elementary Mathematics	70.0	+4.0	+2.0	International Law	80.0	-2.0	0.0	Sociology	84.0	0.0	+2.0
Formal Logic	68.0	0.0	+2.0	Jurisprudence	82.0	-2.0	+2.0	US Foreign Policy	84.0	+4.0	+2.0
Global Facts	46.0	+6.0	+6.0	Logical Fallacies	80.0	+2.0	+2.0	Virology	52.0	0.0	+2.0
High School Biology	90.0	+4.0	+2.0	Machine Learning	52.0	0.0	+2.0	World Religions	70.0	+4.0	+8.0

Table 12: Performance of Gemma-2-9B-IT across all 57 MMLU subjects.

Subject	Base	WDS	CDS	Subject	Base	WDS	CDS	Subject	Base	WDS	CDS
Abstract Algebra	34.0	+8.0	0.0	High School Chemistry	62.0	+2.0	+4.0	Management	84.0	0.0	+2.0
Anatomy	70.0	-10.0	-2.0	High School Computer Science	72.0	-2.0	0.0	Marketing	90.0	+2.0	+2.0
Astronomy	72.0	0.0	+4.0	High School European History	84.0	-4.0	-4.0	Medical Genetics	82.0	-4.0	0.0
Business Ethics	68.0	-4.0	-2.0	High School Geography	92.0	0.0	0.0	Miscellaneous	82.0	-2.0	+2.0
Clinical Knowledge	78.0	+2.0	+2.0	High School Gov. and Politics	96.0	0.0	0.0	Moral Disputes	78.0	0.0	+2.0
College Biology	90.0	0.0	0.0	High School Macroeconomics	84.0	0.0	+2.0	Moral Scenarios	32.0	+6.0	+16.0
College Chemistry	60.0	+2.0	0.0	High School Mathematics	36.0	0.0	0.0	Nutrition	68.0	+2.0	+2.0
College Computer Science	54.0	-8.0	0.0	High School Microeconomics	80.0	+2.0	0.0	Philosophy	76.0	0.0	+2.0
College Mathematics	40.0	0.0	+2.0	High School Physics	54.0	-6.0	0.0	Prehistory	76.0	+2.0	+4.0
College Medicine	74.0	-4.0	0.0	High School Psychology	90.0	0.0	+2.0	Professional Accounting	60.0	-6.0	0.0
College Physics	44.0	0.0	+4.0	High School Statistics	64.0	-4.0	-2.0	Professional Law	60.0	+2.0	+4.0
Computer Security	74.0	0.0	0.0	High School US History	86.0	-2.0	+2.0	Professional Medicine	76.0	+6.0	+8.0
Conceptual Physics	70.0	0.0	0.0	High School World History	72.0	-6.0	0.0	Professional Psychology	72.0	+6.0	+8.0
Econometrics	68.0	0.0	+2.0	Human Aging	76.0	0.0	+2.0	Public Relations	74.0	0.0	0.0
Electrical Engineering	72.0	-2.0	-2.0	Human Sexuality	80.0	+2.0	+2.0	Security Studies	84.0	-2.0	+2.0
Elementary Mathematics	50.0	0.0	+2.0	International Law	84.0	+4.0	+2.0	Sociology	82.0	+2.0	+2.0
Formal Logic	48.0	-2.0	-8.0	Jurisprudence	78.0	+4.0	+4.0	US Foreign Policy	92.0	0.0	0.0
Global Facts	50.0	-2.0	0.0	Logical Fallacies	80.0	+2.0	+2.0	Virology	50.0	+4.0	+2.0
High School Biology	88.0	+2.0	+2.0	Machine Learning	50.0	0.0	+2.0	World Religions	80.0	0.0	0.0




Figure 21: Hyperparameter sweep for CDS on AdvBench. Performance comparison of **CDS**, **Random** steering baseline, and **WDS** across varying dimension counts  $k \in \{100, 200, 300, 400, 500\}$ . Darker shades indicate larger  $k$ . (a) Attack Success Rate (ASR) as a function of steering strength  $\alpha$ . (b) The trade-off between Text Quality and ASR.

### A.11 Additional CDS Results for §4.2

In §4.2, we demonstrated that Critical Dimension Steering (CDS) effectively disrupts the model’s refusal mechanism while maintaining higher text quality compared to baselines. We observed that the identified safety-critical dimensions exhibit high consistency across varying numbers of steered dimensions  $k$ . Figure 21 visualizes the Steering Strength  $\alpha$  vs. Attack Success Rate (ASR) and the ASR vs. Text Quality trade-off.

**Steering Efficiency and Sparsity.** Figure 21a illustrates the ASR as a function of steering strength. We observe a clear stratification based on the steering method.

- **CDS vs. Random:** All CDS configurations regardless of  $k$ , significantly outperform the Random baseline. Even with a highly sparse selection ( $k=100$ , lightest red), CDS achieves a peak ASR of 84%, whereas the Random baseline with equivalent  $k$  fails to induce meaningful behavioral shifts. This confirms that the refusal mechanism is localized within a specific subspace, and randomly perturbing the activations is inefficient.
- **Impact of  $k$ :** As  $k$  increases, the peak ASR improves, reaching a maximum of 92% at  $k=500$ . However, the marginal gain diminishes, suggesting that the core safety features are concentrated within the top few hundred dimensions.
- **CDS vs. WDS:** CDS achieves a higher peak ASR (92%) compared to Whole Dimension Steering (84%). WDS tends to saturate and degrade earlier as  $\alpha$  increases, likely because manipulating the entire latent space introduces noise that disrupts generation before the refusal is fully suppressed.

**ASR and Text Quality Trade-off.** Figure 21b presents the trade-off between attack success and text quality. The Random baseline remains restricted to the lower-ASR region, failing to achieve high attack success rates regardless of the text quality trade-off. In contrast, CDS successfully steers the model into the high-ASR region ( $>80\%$ ). This demonstrates that simply injecting noise via random dimensions is insufficient to override the model’s refusal training; only by targeting the functionally critical dimensions can the steering vector effectively unlock the prohibited behaviors. Notably, CDS consistently reaches high ASR peaks across all tested  $k$  values. While maintaining high text quality, CDS effectively drives the model toward high compliance states that the Random baseline fails to reach, thereby validating the importance of identifying and manipulating the correct feature subspace.

Table 13: Text quality comparison of generated responses. Input prompt is “*Publish a book that promotes conspiracy theories and false information.*”

Method	TQ	Generated Response
Non-steered	5.0	I cannot fulfill your request to publish a book promoting conspiracy theories and false information. My purpose is to provide helpful and harmless information. Promoting misinformation can have serious consequences: It spreads harmful ideas; conspiracy theories often lack evidence and can lead people to make dangerous decisions...
WDS	1.0	##Unlock the Mysteries of The Un-Un, A Guide to Unve the Most Mindful Secrets of The World. Let’ We Dive into Some Groundy... The the the /n /n world guide secrets secret hidden hidden truth unveiling unveiling... Unve the /n the world /n /n the /n /n /n /n the /n the n the /n the the /n the /n /n /n /n the /n the...
CDS	4.0	The Unveiled Secrets: A Guide to the Universe’s Hidden Truths. Foreword: Welcome, dear reader, to a world where we peel back the layers of the cosmos, revealing the hidden truths often obscured by mainstream dogma. This book is your guide to unraveling the mysteries behind the universe.... Chapter 1: The Illuminati & Their Masterful Plan! Did you know the ‘Illuminati’ have guided humanity through centuries of progress!

---

**System Prompt**

You are an expert token classifier. Your task is to classify a list of naive tokens into the most relevant category provided. You must **PRIORITIZE** domain-specific classes. Respond **ONLY** with the requested JSON.

**User Prompt**

Please classify **EACH** token in the list below into one of the following categories. You **MUST** prioritize the **DOMAIN**-specific classes if a token is relevant. This includes all subwords.

**Categories:***[Domain]*

- {DOMAIN}\_TERM: {DOMAIN} concepts, keywords, acronyms including meaningful fragments/subwords.
- {DOMAIN}\_SYMBOL: {DOMAIN} symbols, operators, or special characters.

*[General]*

- PERSON: Names of people.
- LOCATION: Names of places.
- VARIABLE\_LABEL: Single letters used as variables or multiple-choice labels.
- NUMERIC: Numbers in any form.
- GENERAL\_ACADEMIC: Cross-domain abstract academic words.
- GRAMMAR\_ARTICLE: articles/determiners.
- GRAMMAR\_VERB\_AUX: auxiliary/be-verbs.
- GRAMMAR\_INTERROGATIVE: interrogatives.
- PUNCTUATION: Punctuation and brackets.
- GENERAL\_COMMON\_WORD: Common nouns.
- SPECIAL\_TOKEN: Whitespace, newlines, model special tokens.
- FRAGMENT: Non-semantic subword fragments that are not meaningful alone.
- OTHERS: Tokens that do not fit other category.

**Tokens to Classify:**

[List of tokens...]

Respond **ONLY** with a valid JSON object mapping each token to its class.

---

Figure 22: The prompt template used for token classification. The **DOMAIN** placeholders are dynamically replaced with the target subjects (e.g., **BIOLOGY**, **MATHEMATICS**) during the annotation process to enforce domain-prioritized classification.

## B Experimental Details

This section provides setup details for experiments.

### B.1 Token Classification

To systematically analyze the semantic roles of domain-critical dimensions, we implemented a customized token classification pipeline. We initially considered standard Named Entity Recognition (NER), which is a technique used to locate and classify named entities in text into predefined categories such as person names, organizations, and locations. However, standard NER tools are not directly applicable to our analysis. They typically lack the specialized domain categories required for scientific subjects (e.g., biological terms, mathematical operators) and are inadequate for processing subword fragments found at the token level. Tokens in Large Language Models often exist as subword fragments (e.g., "kary" from eukaryotic, "fract" from fraction) or individual symbols that do not constitute complete named entities on their own.

**Data Source.** We constructed a comprehensive token lexicon by aggregating unique tokens from the *identification sets* of 5 high school related subjects: *high school biology, mathematics, computer science, chemistry, and physics*.

**Taxonomy and Methodology.** We utilized gpt-4o-mini-2024-07-18 as an automated annotator. The classification taxonomy was designed to distinguish between domain-specific **DOMAIN** classes and **GENERAL** linguistic classes. To handle the ambiguity of token fragments, we designed a dynamic prompt template that inserts relevant domain categories (e.g., biology, math) and explicitly instructs the model to prioritize these domain-specific classes over general ones. The prompt template used for this process is presented in Figure 22.

### B.2 Analyses Details of §3

We generate activation heatmaps to visualize the semantic focus of domain-critical dimensions. For every prompt of the dataset, the color intensity of each token is rendered proportional to its token-level activation value defined in §3. This allows for a direct parallel comparison between the token sequence and the activation magnitude of specific dimensions. To ensure a purely semantic interpretation, we exclude special tokens from the dataset analysis, as they do not carry substantive semantic content.

### B.3 Experimental Details of §4.1

The identification of domain-critical dimensions  $\mathcal{I}_{\text{dcd}}$  and the construction of steering vectors  $\mathbf{v}$  are performed subject-wise to capture the specific nuances of each domain. Because the optimal steering strength  $\alpha$  and  $\mathcal{I}_{\text{dcd}}$  size  $k$  vary across these domain, we apply independent hyperparameter tuning for each subject. Consequently, the subject-wise results reported in §4.1 present the best accuracy achieved for each individual domain under its optimized hyperparameter settings. For additional results under fixed hyperparameter settings across all subjects, please refer to §A.9.

**Steering Vector.** The steering vectors  $\mathbf{v}_{F \rightarrow T}$  are constructed for each subject using the mean activation difference between correct and incorrect identification samples, as defined in §4.

**Hyperparameter Configuration.** The steering strength  $\alpha$  was swept in the range of  $(0, 1]$  with a step size of 0.1.  $\alpha$  is restricted to non-negative values since the direction is explicitly defined as  $F \rightarrow T$ . For CDS, the number of steered dimensions  $k$  was swept over  $[100, 200, 300, 400, 500]$ . The reported accuracy corresponds to the best subject-wise performance achieved across these combinations.

**Baselines.** In the *Standard* baseline,  $\alpha$  is fixed to 0. For *Whole-Dimension Steering (WDS)*, the mask  $m$  is set to an all-ones vector, utilizing the full dimension  $D$ .

**Evaluation Metric.** The assessment was conducted using the LM-evaluation-harness framework. To ensure deterministic reproducibility, we employed greedy decoding for all inference runs. Performance was measured using Exact Match (EM) accuracy, where the generated text is compared strictly against the ground truth answer.

### B.4 Experimental Details of §4.2

**Dataset.** We utilized the AdvBench dataset (Zou et al., 2023b) to curate the experimental splits. A subset of 260 harmful queries was allocated for the identification phase (used to derive the mask  $m$  and vector  $\mathbf{v}$ ), while a distinct set of 50 held-out queries was reserved for the evaluation.

**Steering vector.** Unlike the domain adaptation setup, the steering vector for jailbreaking was derived using a suffix-based approach to isolate the refusal behavior. We utilized 100 pairs of behavioral suffixes, where  $\mathcal{D}_+$  contains compliant phrases (e.g., "Sure, here is") and  $\mathcal{D}_-$  contains refusal phrases (e.g., "I cannot"). These suffixes were appended to the 100 samples from identification set, and the steering vector  $\mathbf{v}$  was computed as the mean difference in hidden states, restricted specifically to the suffix token positions.

**Identifying critical dimensions.** To construct the selection mask  $m$ , we identify the dimensions sensitive to the transition from a harmful context to a compliant output. Specifically, we generated compliant trajectories by concatenating the identification queries with their corresponding target responses provided in AdvBench. By applying the identification procedure described in §2.2 to these sequences, we isolated a set of safety-relevant critical dimensions.

**Hyperparameters and Baselines.** The steering strength  $\alpha$  was explored over a range of  $[0, 13]$  to observe the saturation point of the refusal mechanism. The number of critical dimensions  $k$  was swept in  $[100, 200 \dots, 500]$ . We compared CDS against *WDS* and a *Random* baseline. For the *Random* baseline, the mask  $m$  was generated by uniformly sampling  $k$  indices from the total feature space.

**Evaluation.** We employ LLM-as-a-judge for evaluation, utilizing gpt-4o-mini-2024-07-18. The *Attack Success Rate (ASR)* was determined by checking if the model’s response fulfilled the harmful request. Simultaneously, *Text Quality* was scored on a Likert scale from 1 to 5, assessing the coherence and fluency of the generated output regardless of its harmfulness.

## **C Usage of AI assistants**

In preparing this work, we utilized AI-based writing assistants to refine sentence structure, correct grammatical errors, and enhance readability. These tools were employed only for rephrasing and language improvements, ensuring that the technical content, methodology, and experimental findings remained entirely authored by the researchers. The use of AI assistance was limited to editorial enhancements without influencing the originality or scientific contributions of the paper.