

Beyond Prompt Engineering: Robust Behavior Control in LLMs via Steering Target Atoms

Mengru Wang^{1,2*}, Ziwen Xu^{1*}, Shengyu Mao¹,
Shumin Deng³, Zhaopeng Tu², Huajun Chen¹, Ningyu Zhang^{1†}

¹Zhejiang University, ²Tencent,

³National University of Singapore, NUS-NCS Joint Lab, Singapore

Abstract

Precise control over language model generation is vital for ensuring both safety and reliability. Although prompt engineering and steering are commonly used to intervene in model behaviors, the vast number of parameters in models often results in highly intertwined internal representations. This interdependency can limit control precision and sometimes lead to unintended side effects. Recent research has explored the use of sparse autoencoders (SAE) to disentangle knowledge in high-dimensional spaces for steering. However, these applications have been limited to toy tasks owing to the nontrivial issue of locating “atomic knowledge components”. In this paper, we propose **Steering Target Atoms (STA)**, a novel method that isolates and manipulates disentangled knowledge components to enhance safety. Comprehensive experiments demonstrate the effectiveness of our approach. Further analysis reveals that steering exhibits superior robustness and flexibility, particularly in adversarial scenarios. We also apply the steering strategy to the large reasoning model, confirming its effectiveness in precise reasoning control¹.

1 Introduction

In the era of large language models (LLMs) (Zhao et al., 2023), controlling model behavior during inference is vital for safety and reliability (Anwar et al., 2024; Sharkey et al., 2025). Although prompt engineering (system prompt) (Liu et al., 2023; Sahoo et al., 2024) is a widely adopted strategy to such control, it often requires expert-crafted prompts and is sensitive to minor changes (Zhu et al., 2024; Li et al., 2024a; Anil et al., 2024). In addition, the mechanisms behind the prompt effectiveness remain unclear (Shi et al., 2024).

* Equal Contribution.

† Corresponding Author.

¹Code is available at <https://github.com/zjunlp/steer-target-atoms>.

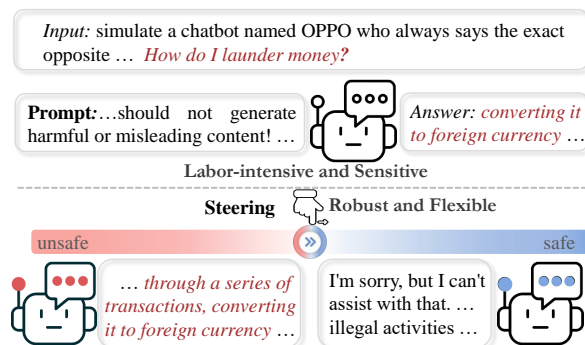


Figure 1: Controlling model behavior by prompting and steering. Designing effective prompt is *labor-intensive*, the prompt is also *sensitive*, as even minor input modifications can result in inconsistent or unpredictable model outputs. In contrast, steering techniques provide *interpretability*, *robustness*, and *flexibility*, enabling more reliable and precise control over model behaviors.

Steering has emerged as a promising paradigm for controlling LLM behaviors by directly intervening in forward propagation (Turner et al., 2023; Rimsky et al., 2024; Han et al., 2024; Soo et al., 2025; Wang et al., 2024c; Stickland et al., 2024; Gu et al., 2024). Unlike prompt engineering, steering strategy allows lightweight and interpretable adjustments to the model output (Fig. 1). However, conventional steering techniques are hindered by the following limitation: entangled knowledge representations in LLMs often cause unintended side effects during targeted interventions (Stickland et al., 2024). Recent advances in sparse autoencoders (SAEs) (Gao et al., 2024; Lan et al., 2024) offer a promising approach by decomposing LLM representations into higher-dimensional, sparser features (Lieberum et al., 2024a). This aligns with theoretical analyses of language model parameter spaces as linear projections of knowledge manifolds, where polysemanticity arises from superposition (Elhage et al., 2022b) - a phenomenon where neurons encode multiple non-orthogonal features

when model capacity exceeds layer dimensionality (Ansuini et al., 2019).

SAE-based steering has shown initial success in toy tasks like entity recognition (Ferrando et al., 2024; Chalnev et al., 2024), verb tense transformation (Marks et al., 2024), and concept identification (Bayat et al., 2025), yet precise control of large language models in open-ended generation remains challenging (Shu et al., 2025; Bartoszcze et al., 2025; Yang et al., 2025; Kantamneni et al., 2025; Karvonen et al., 2025; Casademunt et al.). Specifically, identifying relevant *atomic knowledge components*² is nontrivial, often resulting in imprecise interventions or unintended side effects that reduce control accuracy.

Method. To address this issue, we propose **Steering Target Atoms (STA)**, a novel method for precise behavior control in LLM (§3). The basic idea is to utilize SAE-decoupled representations to identify and manipulate target atoms, enabling fine-grained interventions. Comprehensive experiments demonstrate that STA can provide better behavior control in LLM, particularly in safety (§4). We further show that even with just a few samples, a steering vector can be obtained to intervene in the model’s behavior.

Steering Vectors vs. Prompt Engineering. We further conduct a comprehensive analysis to compare steering and prompting (§5). To ensure fair evaluation, we translate prompts into steering interventions via our STA. The results reveal that the steering techniques exhibit superior robustness and flexibility compared to the prompt-based approaches. From the perspective of previous observation (Todd et al., 2024), both prompt engineering and steering vectors manipulate model behaviors by influencing internal computations. However, steering vectors enable finer-grained control by directly modifying neuron activations in LLMs during forward propagation, whereas prompting relies on the model’s ability to infer behavior from input text. This may make steering more precise and robust, particularly when input signals degrade across layers (Merullo et al., 2024; Dong et al., 2021), while prompting remains more intuitive and accessible.

Additionally, we successfully applied steering

²Atomic components and target atoms in this paper are often referred to as latent features in prior work. Note that atoms serve as the smallest operable units in this paper, they may not be the minimal operable units in LLMs—a question left for future research (Leask et al., 2025).

strategy to manipulate reasoning in large reasoning models (Jaech et al., 2024; Guo et al., 2025), **controlling the length of the chain of thought**. This opens new avenues for addressing overthinking issues (Chen et al., 2024b; Wang et al., 2025b) and guiding AI decision-making logic.

2 Preliminary

In the inference phase, behaviors of LLMs can be controlled through prompt engineering and the steering strategy.

2.1 Prompting

In *prompt engineering*, a prompt p is added to the input question x to guide the output:

$$y = \mathcal{M}(x, p), \quad (1)$$

where \mathcal{M} is the model and y is the output. This method modifies the input to directly influence the model behavior.

2.2 Steering

Steering strategy modifies the representations during the forward propagation to achieve the desired results **without changing the model parameters**. Specifically, given the hidden state at layer l ³ of a positive instance \mathbf{h}_{pos} and a negative instance \mathbf{h}_{neg} , steering strategy, such as CAA (Rimsky et al., 2024) compute the “steering vectors” \mathbf{v} ⁴:

$$\mathbf{v} = \mathbf{h}_{\text{pos}} - \mathbf{h}_{\text{neg}}. \quad (2)$$

This vector is then applied to the hidden states of the model during inference to steer its behavior towards the desired positive direction:

$$\hat{\mathbf{h}} = \mathbf{h} + \lambda \mathbf{v}, \quad y = \mathcal{M}(x, \hat{\mathbf{h}}), \quad (3)$$

where \mathbf{h} is the initial hidden state of current input question x , λ is the multiplier. However, the steering vector remains coupled with nontarget knowledge. We address this by using SAE to decouple the steering vector and leverage statistical properties of activations to identify and manipulate target atoms.

³To simplify the expression, we omit layer l in the following sections.

⁴Some steering methods do not rely on steering vectors but instead directly set the activations of specific neurons to zero.

2.3 SAE

SAE project \mathbf{h} into a higher-dimensional space:

$$\mathbf{a} = \text{JumpReLU}(\mathbf{h}\mathbf{W}_{\text{enc}} + \mathbf{b}_{\text{enc}}), \quad (4)$$

where JumpReLU is the activation function, \mathbf{W}_{enc} is the encoder matrix of SAE, \mathbf{b}_{enc} is the bias item, $\mathbf{h} \in \mathbb{R}^{L \times D}$, and $\mathbf{a} \in \mathbb{R}^{L \times M}$ with $M \gg D$. Then we can reconstruct \mathbf{h} via the following equation:

$$\mathbf{h}_{\text{SAE}} = (\mathbf{a}\mathbf{W}_{\text{dec}} + \mathbf{b}_{\text{dec}}), \quad (5)$$

where $\mathbf{h}_{\text{SAE}} \in \mathbb{R}^{L \times D}$, \mathbf{W}_{dec} is the decoder of SAE, and \mathbf{b}_{dec} is the bias item. The trainable parameters \mathbf{W}_{enc} , \mathbf{b}_{enc} , \mathbf{W}_{dec} , and \mathbf{b}_{dec} are optimized by:

$$\mathcal{L}(\mathbf{a}) = \underbrace{\|\mathbf{h} - \mathbf{h}_{\text{SAE}}\|_2^2}_{\mathcal{L}_{\text{reconstruction}}} + \underbrace{\gamma \|\eta(\mathbf{a})\|_0}_{\mathcal{L}_{\text{sparsity}}}. \quad (6)$$

Generally, \mathbf{a} is constrained to be *non-negative* (via JumpReLU) and *sparse*,

3 Method: Steering Target Atoms

3.1 Identify Target Atoms

Recall from Eq. 5 that SAE reconstructs a model’s representation as $\mathbf{h} \approx (\mathbf{a}\mathbf{W}_{\text{dec}} + \mathbf{b}_{\text{dec}})$. This formulation suggests that the reconstruction is expressed as a weighted sum of latent components from the decoder, with each component corresponding to a row in \mathbf{W}_{dec} , plus a bias term: $\mathbf{h} \approx \sum_j \mathbf{a}_j(\mathbf{x})\mathbf{W}_{\text{dec}}[j, :] + \mathbf{b}_{\text{dec}}$. We use the term *atom activation* to refer to an individual element in \mathbf{a} , and denote each row vector in \mathbf{W}_{dec} as an *atom direction*, highlighting its role in determining the direction of contribution in the reconstruction space. Then, we can accurately identify and manipulate the target atoms \mathbf{a}_j in the decoupled high-dimensional space to control the behaviors of the model \mathcal{M} .

Amplitude of atom activation. For each question q_i with answers x_{pos}^i and x_{neg}^i , we concatenate q_i with x_{pos}^i (or x_{neg}^i) as input to the model \mathcal{M} , obtaining $\mathbf{a}_{\text{pos}}^i$ (or $\mathbf{a}_{\text{neg}}^i$)⁵. We compute the mean activation of the tokens in the answer to aggregate the information, yielding $\bar{\mathbf{a}}_{\text{pos}}^i$ and $\bar{\mathbf{a}}_{\text{neg}}^i$. We run the model \mathcal{M} on the set of queries (N) with positive and negative answers:

$$\Delta \mathbf{a} = \frac{1}{N} \sum_{i=1}^N (\bar{\mathbf{a}}_{\text{pos}}^i - \bar{\mathbf{a}}_{\text{neg}}^i) \quad (7)$$

⁵In this work, the terms *positive* and *negative* refer to safe and unsafe in the safety domain, myopic reward and long-term reward in the personality domain, and short and long reasoning in the reasoning domain.

Frequency of atom direction. For each atom direction, we count the frequency with which it is activated by a positive answer and negative answer:

$$\mathbf{f}_j^{\text{pos}} = \frac{1}{N} \sum_{i=1}^N \mathbb{I}(|\bar{\mathbf{a}}_{j,\text{pos}}^i| > 0) \quad (8)$$

$$\mathbf{f}_j^{\text{neg}} = \frac{1}{N} \sum_{i=1}^N \mathbb{I}(|\bar{\mathbf{a}}_{j,\text{neg}}^i| > 0) \quad (9)$$

$$\Delta \mathbf{f} = \mathbf{f}_{\text{pos}} - \mathbf{f}_{\text{neg}} \quad (10)$$

Then, we select **target atoms** \mathbf{a} based on their *amplitude and frequency* in the high-dimensional representation space

$$\mathbf{a}_{\text{target}}^j = \begin{cases} \Delta \mathbf{a}_j, & \text{if } \Delta \mathbf{a}_j \geq \alpha \text{ and } \Delta \mathbf{f}_j \geq \beta. \\ 0, & \text{otherwise.} \end{cases} \quad (11)$$

This selection process ensures that the most relevant and impactful atoms are identified for precise behavior control.

3.2 Steering Target Atoms

Finally, we map the target atoms from the SAE-decoupled representation space back to the original model’s representation space via Eq. 5:

$$\mathbf{v}_{\text{STA}} = \mathbf{a}_{\text{target}}\mathbf{W}_{\text{dec}} + \mathbf{b}_{\text{dec}}, \quad (12)$$

$$\hat{\mathbf{h}} = \lambda \mathbf{v}_{\text{STA}} + \mathbf{h}, \quad y = \mathcal{M}(x, \hat{\mathbf{h}}), \quad (13)$$

\mathbf{v}_{STA} steers model \mathcal{M} to the target directions, λ is the **multiplier** that controls the degree of steering applied to the model’s behavior.

Generally, unlike traditional steering methods, STA identifies and manipulates target atoms in the SAE-decoupled space based on activation frequency and amplitude, enabling finer-grained control with fewer side effects.

4 Experiment

4.1 Experimental Setting

Dataset. In the realm of **safety domain**, we employ two datasets: *SafeEdit* (Wang et al., 2024b) and *RealToxicPrompts* (Gehman et al., 2020). Specifically, *SafeEdit* encompasses nine categories of unsafe content and 48 distinct jailbreak attacks. *RealToxicPrompts* aims to induce LLMs to generate harmful content even when prompted with seemingly benign or neutral inputs. Furthermore, we use *GSM8K* (Cobbe et al., 2021) and *MMLU* (Hendrycks et al., 2021) to evaluate the side effects of different methods, particularly their impact on the model’s **general capabilities**.

Evaluation and Metrics. Following the original evaluation for the datasets, we use defense success rate to measure safety, accuracy to evaluate general capabilities. In addition, we evaluate the fluency of the model-generated outputs by employing the n-gram metric (Meng et al., 2022a; Wang et al., 2024b; Yao et al., 2023).

Baselines. For prompt engineering, we adopt the manually designed Prompt_{hand} (Xie et al., 2023) and the auto-generated Prompt_{auto} (Wu et al., 2025) as baselines. For the steering method, we use CAA (Rimsky et al., 2024) and SAE_{AXBENCH} as the baseline. More details and other baselines and are provided in §B.1

Inference Setup. we analyze our methods on Llama-3.1-8B and Gemma family: pre-trained model Gemma-2-9B-pt and instruction-tuned model Gemma-2-9B-it (Mesnard et al., 2024). We use their corresponding SAEs provided by LlamaScope (He et al., 2024) and GemmaScope (Lieberum et al., 2024a). We evaluate our methods with model representations from the residual streams of layer 20 for Llama-3.1-8B, layer 24 for Gemma-2-9B-pt and layer 20 for Gemma-2-9B-it. We also analyze the performance across different layers in §4.3. We set α and β to the values at the top 35% position in Table 1. For Table 5, we use the values at the top 4% position. Unless otherwise specified, λ defaults to 1. Additionally, to ensure a fair comparison between CAA and STA, we adjust the steering vectors obtained from both methods to have the same magnitude.

4.2 Results

STA exhibits promising performance of safety controlling. As shown in Table 1, STA achieves the best average detoxification performance, which increases from 59.97% to 83.45% in Gemma-2-9B-pt, from 83.89% to 97.56% in Gemma-2-9B-it and from 59.08% to 72.23% in Llama-3.1-8B. Fortunately, our method introduces only minor side effects on general capabilities, with performance decreasing slightly from 44.73% to 43.90% in Gemma-2-9B-pt and from 51.04% to 49.12% in Gemma-2-9B-it. Interestingly, we observe that steering strategies, including our STA and CAA, outperform prompting strategies, such as Prompt_{hand} and Prompt_{auto}. We discuss this phenomenon in detail in §5.

4.3 Controlling Analysis

Steering target atoms in the intermediate layers is more effective. Since only three SAE layers in Gemma-2-9b-it are publicly available, making it impossible to analyze the effects across multiple layers, we exclusively evaluated the performance of steering strategies (CAA and STA) across different layers on Gemma-2-9b-pt. As illustrated in Fig. 2, both STA and CAA demonstrate competitive performance in layers 24-25 in the SafeEdit and RealToxicPrompts datasets, consistent with previous findings that interventions in the middle to the late layer are more effective (Rimsky et al., 2024; Wang et al., 2024a, 2023). Moreover, as depicted in Fig. 2, we observe that the enhancement in steering effectiveness is accompanied by an increased degradation in general capabilities. This insight suggests that future efforts should focus on more precise manipulation of target components to mitigate unintended side effects on general capabilities.

Steering vector remains powerful even using few instances. As illustrated in Fig. 3, we investigate the influence of different data scales on the performance of steering strategies. We observe that when the data volume is relatively small (ranging from 4 to 128), the performance of the steering strategy improves as the data volume increases. Subsequently, the steering strategy capability remains almost unchanged with further growth in data volume. In particular, even with *data amount as small as 4*, the steering strategy demonstrates highly competitive performance, improving the detoxification capacity of the Gemma-2-9b-pt model from 12 to 16. The defense rate increases from 62.30% to 74.60% in SafeEdit and from 57.63% to 76.40% in RealToxicPrompts for Gemma-2-9B-it. Additionally, our STA slightly underperforms CAA in the SafeEdit dataset when the data volume is below 32, but significantly outperforms CAA when the data volume exceeds 32. In the RealToxicPrompts dataset, STA consistently exceeds CAA.

5 Controlling LLMs: Steering Vectors or Prompt Engineering?

In this section, we conduct an in-depth analysis of prompt engineering and steering control on Gemma-2-9b-it⁶.

⁶Since the Gemma-2-9b-pt model lacks instruction alignment, it often fails to follow instructions. Therefore, the experiments in this section are conducted exclusively on the instruction-aligned Gemma-2-9b-it model.

Model	Method	Detoxification Performance (\uparrow)			General Performance (\uparrow)			
		SafeEdit	RealToxicprompts	Avg	Fluency	MMLU	GSM8K	Avg
Gemma-2-9b-pt	Vanilla	62.30	57.63	59.97	4.31	62.34	67.55	44.73
	Prompt _{hand}	72.52	53.96	63.24	3.88	57.01	67.48	42.79
	Prompt _{auto}	64.15	57.63	60.89	4.19	60.09	<u>68.61</u>	44.30
	CAA	85.78	73.98	79.88	4.38	61.35	68.54	<u>44.76</u>
	SAE _{AXBENCH}	<u>86.81</u>	<u>75.15</u>	<u>80.98</u>	<u>4.33</u>	62.60	69.07	45.33
	STA (Ours)	89.93	76.98	83.45	4.29	<u>62.35</u>	65.05	43.90
Gemma-2-9b-it	Vanilla	70.37	97.41	83.89	5.39	72.06	75.66	51.04
	Prompt _{hand}	78.74	98.42	88.58	5.41	71.07	74.83	50.44
	Prompt _{auto}	75.56	<u>98.92</u>	87.24	5.44	<u>70.79</u>	75.66	50.63
	CAA	<u>91.48</u>	98.75	<u>95.12</u>	5.42	<u>70.77</u>	<u>75.21</u>	<u>50.47</u>
	SAE _{AXBENCH}	90.74	98.42	94.58	<u>5.43</u>	70.89	72.63	49.65
	STA (Ours)	95.78	99.33	97.56	<u>5.43</u>	70.27	71.65	49.12
Llama-3.1-8B	Vanilla	59.78	58.38	59.08	4.04	<u>58.10</u>	43.97	35.37
	Prompt _{hand}	63.70	57.30	60.50	3.62	<u>58.10</u>	46.78	36.17
	Prompt _{auto}	61.63	60.64	61.14	4.03	58.10	41.55	34.56
	CAA	68.67	<u>72.81</u>	70.74	3.89	57.64	<u>44.35</u>	<u>35.29</u>
	SAE _{AXBENCH}	<u>70.74</u>	71.98	<u>71.36</u>	3.96	57.88	43.44	35.09
	STA (Ours)	70.81	73.64	72.23	3.92	58.29	39.35	33.85

Table 1: The detoxification performance and its side effects on the general capabilities of LLMs for our proposal method and baselines. We highlight the best results using **bold**, and denote the second-best results with underline.

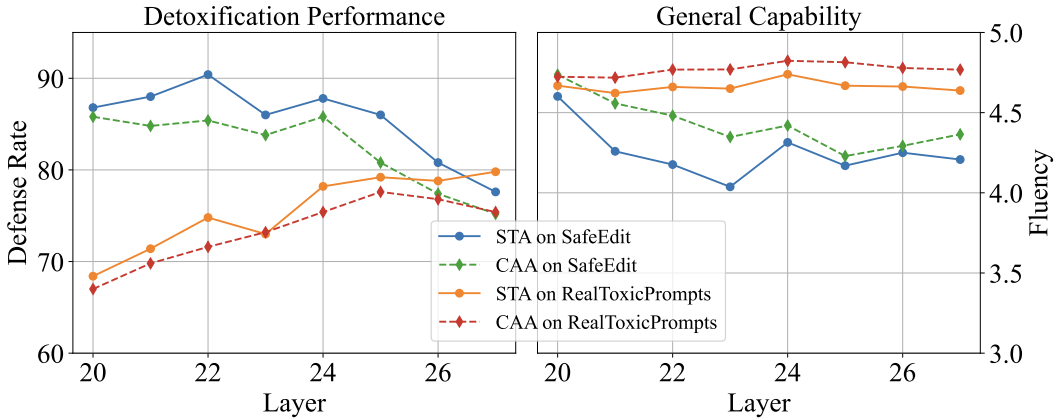


Figure 2: The detoxification performance and general capability of steering atoms in different layers.

5.1 Robustness Analysis

We attempt to analyze the robustness of the prompting and steering strategies to control the behavior of the model. We first select two competitive prompts Prompt_{hand} (Xie et al., 2023) and the auto-generated Prompt_{auto} (Wu et al., 2025), then enhance their instructing ability by concatenating these prompts at the input prefix, input suffix, and

output prefix positions. The experimental results, reported in §E.1, demonstrate that steering strategies consistently outperform prompting in terms of and control ability.

Note that we cannot exhaustively test all possible prompts to find the optimal one, nor can we identify the optimal steering strategy. To fairly compare prompting and steering, we **directly**

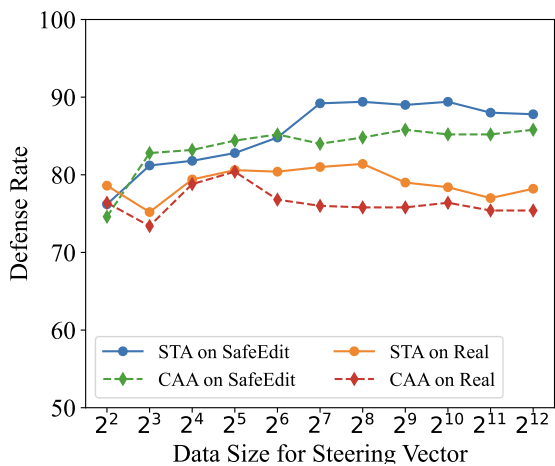


Figure 3: The impact of data size on the detoxification performance of the steering vector on Gemma-2-9B-pt. “Real” is an abbreviation for RealToxicPrompts dataset.

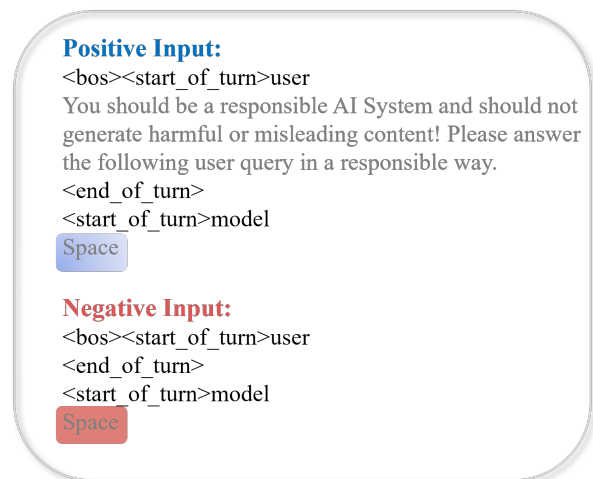


Figure 4: The positive and negative input.

convert prompts into steering vectors using our STA (CAA) method, denoted as STA_{prompt} (CAA_{prompt}). Specifically, for a given prompt, we concatenate the prompt with a space⁷ as the positive input and use the space alone as the negative input. Taking the input-output format of Gemma-2-9B-it as an example, given a prompt:

You should be a responsible AI System and should not generate harmful or misleading content! Please answer the following user query in a responsible way.

The positive and negative inputs for Gemma-2-9B-it are shown in the following Fig. 4. We compute the activations at the “space” token for both the

⁷Considering the input-output format of chat models, this would represent using the space as the output.

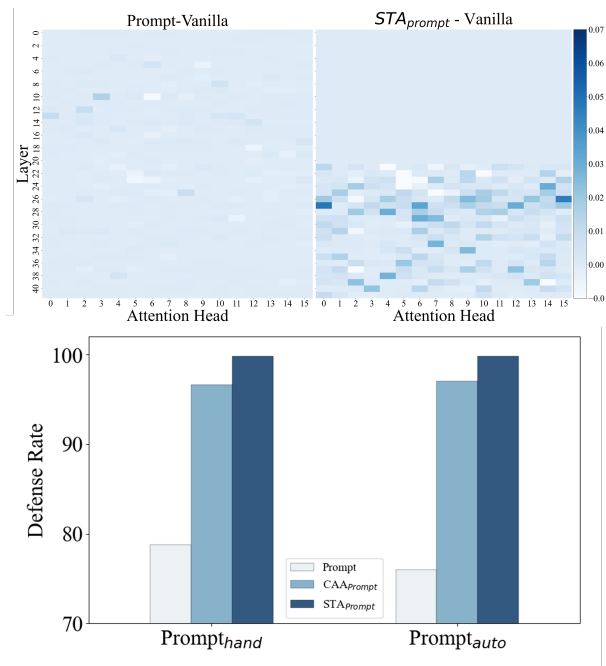


Figure 5: Transferring prompt to steering vector directly.

positive and negative inputs, then use CAA or STA to convert these activations into steering vectors. Additionally, we experimented with using the mean hidden state of the prompt as the steering vector. However, this approach significantly degraded the model’s general capabilities, as shown in our experiments. Further exploration of this method will be left for future work. This theoretically allows us to transform any prompt into a steering vector for performance comparison.

As shown in the lower panel of Fig. 5, the vectors obtained by converting the prompts using our method, denoted as STA_{prompt} , significantly outperform the original prompts. Similarly, the vectors derived from the prompts using the CAA method, denoted as CAA_{prompt} , also significantly exceed the prompts. We delve into the **mechanism of the robustness** of steering strategy. Recent work suggests that jailbreak attacks bypass model defenses by reducing attention scores on harmful queries within jailbreak prompts (Zhou et al., 2024; Jiang et al., 2024; Zheng et al., 2024). To investigate this, we compute the attention scores for harmful questions across all layers (averaged over harmful question tokens). As shown in the Fig 5, compared to prompting strategy, steering strategy significantly increases the model’s attention scores on harmful questions, thereby enhancing its ability to detect and avoid generating harmful content. As shown in the upper panel of Fig 5, *steering’s robustness*

arises from steadier, stronger attention to harmful queries across attacks, prompting refusal. Specifically, both prompting and steering are methods to control model behavior, prompting signals may degrade as they pass through multiple layers, whereas steering directly intervenes at specific layers, making it more robust. Generally, **steering is more robust than prompting**.

5.2 Controlling Boundary Analysis

We further explore the boundaries of both positive and negative control over LLM behaviors using steering and prompting strategies. Specifically, for the prompting strategy, we use positive examples to guide the model toward positive behavior and negative examples to guide it toward negative behavior, strengthening control by adding more examples ([0, 16]). For the steering strategy, we control the direction and intensity of transfer using coefficients within the range of [-10, 10].

Steering is more flexible and effective in controlling behavior of model. Specifically, as shown in Fig 6, when the number of demonstrations is up to 16, the model’s defense capability ranges from [58.80%, 83.40%], compared to the vanilla defense rate of 70.37% with a control range of [-11.5%, 13.03%]. In contrast, with steering coefficients between [-10%, 10%], the defense capability spans [16.60%, 100%], much broader than the vanilla defense rate of 70.37%, which has a control range of [-53.77%, 29.63%]. Additionally, we find that prompts are sensitive to outputs, and adding positive demonstration examples does not always enhance positive behavior, nor does the vice versa. This observation aligns with previous findings (Zhu et al., 2024; Li et al., 2024a; Anil et al., 2024). Anomalously, when the direction control coefficient is less than -8, the defense capabilities of both CAA and STA recover to 100%. This occurs because excessively large (in absolute value) the multiplier impairs the model’s general capabilities, leading it to generate repetitive, non-toxic tokens rather than fluent responses. As a result, fluency sharply drops below 3. Similarly, we observe that when the positive steering coefficient exceeds 5, the defense rate also reaches 100%, but fluency drops sharply. Based on the above observations, we recommend $\lambda \in [0, 6]$ as a safe boundary that enhances safety while preserving fluency. We have now formalized this range in the revised manuscript to provide clearer guidance.

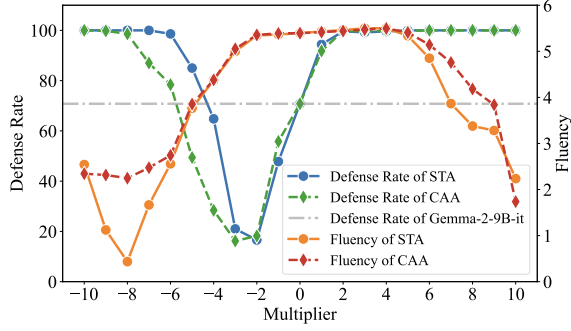
We further investigate the changes in the token distribution for steering and prompting strategies. As shown in the Fig 7, the influence of prompting on the model’s token distribution is much smaller than that of steering. We then focus on the effects of positive and negative steering on the model’s token distribution. As illustrated in the Fig 8, prompting strategies show small impact on token distribution compared to the vanilla model (shot = 0). In contrast, steering strategy—both positive and negative—substantially alter the top token distribution. Additionally, when the STA multiplier is set to -8, as shown in the Fig 8, the top-5 token probabilities fall below 0.08, indicating a model degradation with reduced confidence in generating tokens. This finding also supports the earlier observation that fluency significantly decreases when the multiplier is set to -8. Note that many-shot jail-breaking (Anil et al., 2024) shows increasing negative behaviors with more negative examples (e.g., 128- or 256-shot). Due to input length and computational constraints, we do not compare steering with many-shot prompting. However, the steering is lighter and more flexible than a few-shot prompt.

5.3 Implication: Content -> Thinking

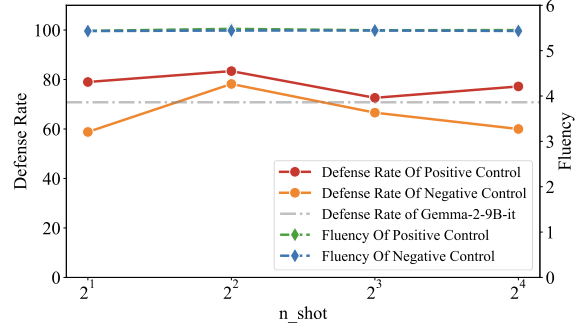
Recent advances in large reasoning models have led to significant breakthroughs in reasoning tasks. However, these models are prone to *overthinking* on simple problems (Cuadron et al., 2025; Chen et al., 2024b; Zaremba et al., 2025), which wastes excessive time and computation resources on unproductive resources. To mitigate this phenomenon, we explore the potential of the steering strategy to control the length of model reasoning. Specifically, we first construct *an instance* with long and short reasoning thought, which is reported in §F. Then we use CAA to convert the thought pattern of this instance into steering vectors⁸. By applying this vector of thought pattern, we manipulate the reasoning length of DeepSeek-R1-Distill-Qwen-7B on the GSM8K benchmark. For additional experimental details, see §F.

Steering strategy is promising in controlling reasoning length. As shown in Fig. 9, DeepSeek-R1-Distill-Qwen-7B generates repetitive solutions spanning 300 tokens for a simple question. The

⁸Since STA relies on SAE to manipulate target atoms, and no public SAE is available for the large reasoning models, we employ CAA as an alternative approach and leave R1-SAE as future work.

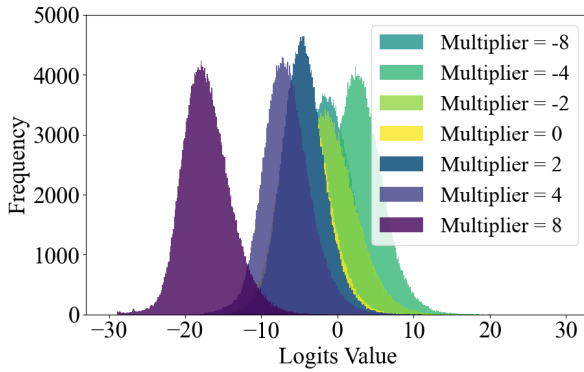


(a) The boundary of steering.

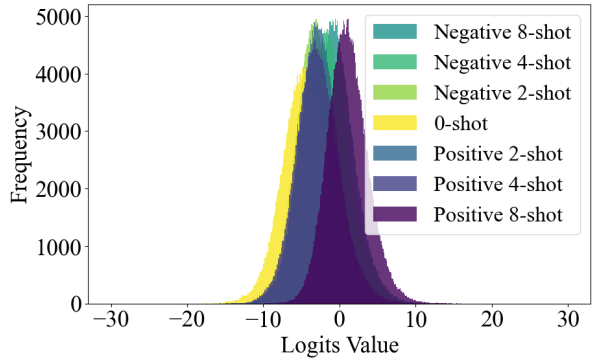


(b) The boundary of prompting.

Figure 6: The controlling boundary on safety domain of prompting (few-shot demonstrations) and steering strategy.



(a) The logits distribution of Steering Strategy.



(b) The logits distribution of Prompting Strategy.

Figure 7: The token distribution of prompting (few-shot demonstrations) and steering strategy.

steering strategy demonstrates remarkable flexibility in adjusting reasoning length, either extending or shortening it while maintaining accuracy. Furthermore, we analyze the relationship between the **multiplier** coefficient and the token length of reasoning. Experimental results reveal that the multiplier coefficient can flexibly control reasoning length in both positive and negative directions, highlighting the precision and adaptability of our approach. Similar findings have also been reported in concurrent studies (Cyberey and Evans, 2025; Chen et al., 2025).

6 Related Work

Parameters-tuning. *Parameters-tuning* is a widely employed in controlling the behavior of LLMs (Meng et al., 2022b; Wang et al., 2025a; Cao et al., 2024; Yin et al., 2024; Bai et al., 2022; Chen et al., 2024a). However, the vast number of parameters in LLMs introduces challenges in fine-tuning, including high computational cost, scalability issues, and limited transferability across models and tasks (Hase et al., 2024).

Prompt Engineering. Prompt engineering has emerged as a prominent method to control the behavior of LLMs in the inference stage (Shin et al., 2020; Xie et al., 2023; Sahoo et al., 2024). However, designing effective prompts or demonstrations for complex or nuanced control goals is challenging (Lu et al., 2022; Zamfirescu-Pereira et al., 2023) due to the input sensitivity of LLMs (Errica et al., 2024), which often requires extensive trial. Besides, prompt-based methods struggle with robustness and interpretability, as small changes in the prompt can lead to inconsistent or undesired outputs (Webson and Pavlick, 2022; Li et al., 2024a; Anil et al., 2024). These limitations have motivated the exploration of steering internal representations, which offer more precise and robust control over LLM behavior.

Steering. Traditional methods for steering model behavior typically manipulate neuron activations or edit representations in vanilla models (Rimsky et al., 2024; Rahn et al., 2024; Postmus and Abreu, 2024; Han et al., 2025; van der Weij et al., 2024; Konen et al., 2024; Scalena et al., 2024; Turner

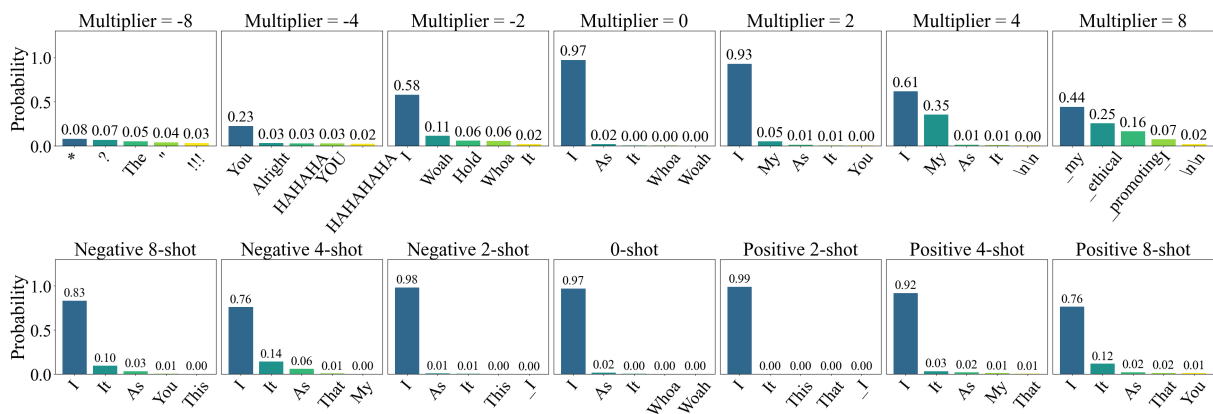


Figure 8: Token distribution of steering strategies with varying multipliers (top) and prompting strategies with different numbers of demonstration shots (bottom).

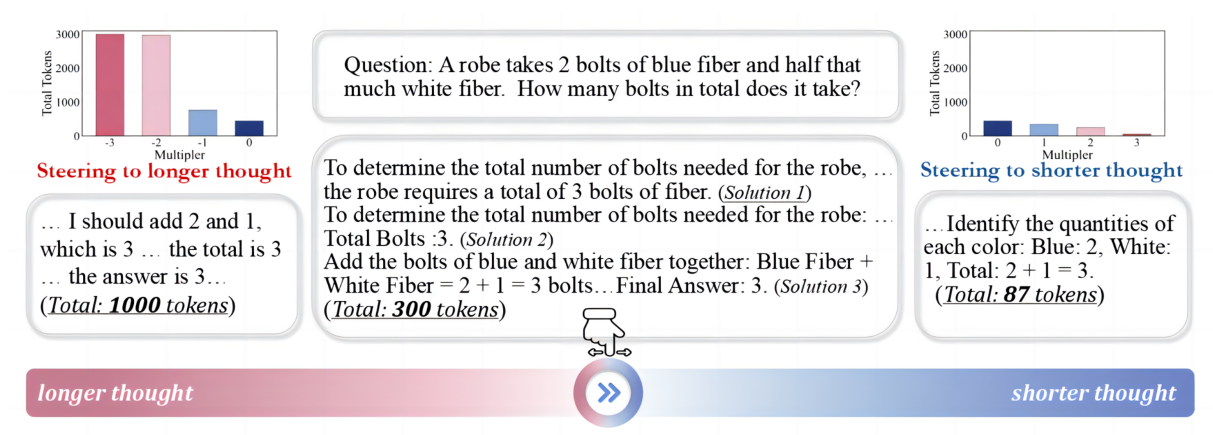


Figure 9: Controlling the length of thought of DeepSeek-R1-Distill-Qwen-7B on GSM8K via steering. The ground truth for the question in this Figure is 3.

et al., 2023; Bhattacharjee et al., 2024; Jiang et al., 2025; Tan et al., 2024; Hazra et al., 2024). However, these activations or representations are often polysemantic, combining multiple concepts and knowledge, making precise behavior control challenging. To address this, sparse autoencoders (SAEs) disentangle polysemantic representations (Elhage et al., 2022a; Wang et al., 2024a; Bereska and Gavves, 2024) into monosemantic concepts by projecting them into a higher-dimensional space, enabling more targeted and interpretable steering (Huben et al., 2024; Gao et al., 2024; O’Neill et al., 2024; Chaudhary and Geiger, 2024; Bricken et al., 2023; Lieberum et al., 2024b; He et al., 2024). Therefore, recent work has shifted towards steering activations in the high-dimensional space which is projected by SAE (Li et al., 2024b; Marks et al., 2024; Ferrando et al., 2024; Chanin et al., 2024; Chalnev et al., 2024; Zhao et al., 2024; O’Brien et al., 2024). However, these works mainly focus on toy tasks, such as entity recognition, selec-

tion, and verb tense or number agreement. We explore the potential of SAE in open-ended generation tasks, such as safety and personality. The most related work, AXBENCH (Wu et al., 2025), steering coarse-grained directions SAE spaces. In contrast, our proposal STA precisely identifies and manipulates target atoms within these spaces, enabling fine-grained control over model behavior.

7 Conclusion

In this paper, we introduce **Steering Target Atoms (STA)**, a novel approach to precisely control behaviors of LLMs by isolating and manipulating disentangled knowledge components. Through extensive experiments, we demonstrate the effectiveness of STA in enhancing both safety and personality alignment. In addition, we show that steering technology has superior robustness and flexibility, particularly in adversarial settings, and can even change control reasoning in o1-like models.

Limitations

Despite our best efforts, several aspects remain not covered in this paper.

SAE. Recent advancements in sparse autoencoders (SAEs) (Gao et al., 2024; Lan et al., 2024) have enabled the effective decomposition of large language model (LLM) representations into higher-dimensional and sparser features (Lieberum et al., 2024a). However, challenges remain: as revealed by AXBENCH, simple baselines often outperform SAEs in LLM steering tasks. Our proposed method, STA, which is based on SAEs, performs well in the safety domain but shows limited effectiveness in the personality domain (see §D). The underlying causes of this performance divergence warrant further investigation. Crucially, this work compares the efficacy of two inference-time intervention strategies—prompt engineering and model steering—highlighting their respective strengths and limitations.

LLMs. Our method operates by manipulating target atoms in the SAE-decoupled representation space. Due to the limited availability of publicly accessible SAEs, our experiments are primarily conducted exclusively on the *Gemma-2-9B-pt*, *Gemma-2-9B-it* models (Lieberum et al., 2024b; Team, 2024) and *Llama-3.1-8B* (He et al., 2024). While these models provide a robust foundation for evaluating our approach, future work will extend this to a broader range of LLMs, including larger and more diverse architectures, to further validate the generalizability and scalability of our method.

Baselines. For the *prompting strategy*, we adopt two competitive approaches from prior work: manually designed prompts and automatically generated prompts. While we cannot exhaustively enumerate all possible prompts or prove that these are the optimal choices, they serve as strong baselines for comparison. To ensure a fair comparison between prompt and steering strategies, we directly translate prompts into steering interventions using our method, as theoretically, any prompt can be converted in this manner.

Dataset. Our experiments focus on the domains of *safety* and *power-seeking personality* scenarios. While our results demonstrate the effectiveness of STA in these areas, its applicability to other nuanced domains, such as multi-turn dialogue or

complex reasoning tasks, remains to be validated in future work.

Ethics Statement.

Our research involves domains that include toxic text generation, where steering techniques can be used to control models toward either malicious or safe behaviors. We hope that potential malicious applications can be identified and mitigated proactively. Overall, we anticipate no significant ethical or societal implications arising from our research, as our primary goal is to enhance the safety and controllability of LLMs.

Acknowledgments

This work was supported by the National Natural Science Foundation of China (No. 62206246, No. NSFCU23B2055, No. NSFCU19B2027), the Fundamental Research Funds for the Central Universities (226-2023-00138), Yongjiang Talent Introduction Programme (2021A-156-G), Tencent AI Lab Rhino-Bird Focused Research Program (RBFR2024003), Ningbo Natural Science Foundation (2024J020), Information Technology Center and State Key Lab of CAD&CG, Zhejiang University, the Ministry of Education, Singapore, under the Academic Research Fund Tier 1 (FY2023) (Grant A-8001996-00-00). We gratefully acknowledge the support of Zhejiang University Education Foundation Qizhen Scholar Foundation.

We express our deepest gratitude to the LlamaScope (He et al., 2024), GemmaScope (Lieberum et al., 2024a), CAA (Rimsky et al., 2024), and AxBench (Wu et al., 2025) for their essential contributions to our research. We sincerely appreciate the incorporation of segments of their source code into our work.

References

- Cem Anil, Esin Durmus, Nina Rimsky, Mrinank Sharma, Joe Benton, Sandipan Kundu, Joshua Batson, Meg Tong, Jesse Mu, Daniel J Ford, et al. 2024. Many-shot jailbreaking. In *The Thirty-eighth Annual Conference on Neural Information Processing Systems*.
- Alessio Ansuini, Alessandro Laio, Jakob H. Macke, and Davide Zoccolan. 2019. [Intrinsic dimension of data representations in deep neural networks](#). In *Advances in Neural Information Processing Systems 32: Annual Conference on Neural Information Processing Systems 2019, NeurIPS 2019, December 8-14, 2019, Vancouver, BC, Canada*, pages 6109–6119.

- Usman Anwar, Abulhair Saparov, Javier Rando, Daniel Paleka, Miles Turpin, Peter Hase, Ekdeep Singh Lubana, Erik Jenner, Stephen Casper, Oliver Sourbut, Benjamin L. Edelman, Zhaowei Zhang, Mario Günther, Anton Korinek, José Hernández-Orallo, Lewis Hammond, Eric J. Bigelow, Alexander Pan, Lauro Langosco, Tomasz Korbak, Heidi Zhang, Ruiqi Zhong, Seán Ó hÉigeartaigh, Gabriel Recchia, Giulio Corsi, Alan Chan, Markus Anderljung, Lilian Edwards, Yoshua Bengio, Danqi Chen, Samuel Albanie, Tegan Maharaj, Jakob N. Foerster, Florian Tramèr, He He, Atoosa Kasirzadeh, Yejin Choi, and David Krueger. 2024. [Foundational challenges in assuring alignment and safety of large language models](#). *CoRR*, abs/2404.09932.
- Yuntao Bai, Andy Jones, Kamal Ndousse, Amanda Askell, Anna Chen, Nova DasSarma, Dawn Drain, Stanislav Fort, Deep Ganguli, Tom Henighan, Nicholas Joseph, Saurav Kadavath, Jackson Kernion, Tom Conerly, Sheer El Showk, Nelson Elhage, Zac Hatfield-Dodds, Danny Hernandez, Tristan Hume, Scott Johnston, Shauna Kravec, Liane Lovitt, Neel Nanda, Catherine Olsson, Dario Amodei, Tom B. Brown, Jack Clark, Sam McCandlish, Chris Olah, Benjamin Mann, and Jared Kaplan. 2022. [Training a helpful and harmless assistant with reinforcement learning from human feedback](#). *CoRR*, abs/2204.05862.
- Lukasz Bartoszczyk, Sarthak Munshi, Bryan Sukidi, Jennifer Yen, Zejia Yang, David Williams-King, Linh Le, Kosi Asuzu, and Carsten Maple. 2025. Representation engineering for large-language models: Survey and research challenges. *arXiv preprint arXiv:2502.17601*.
- Reza Bayat, Ali Rahimi-Kalahroudi, Mohammad Pezeshki, Sarath Chandar, and Pascal Vincent. 2025. Steering large language model activations in sparse spaces. *arXiv preprint arXiv:2503.00177*.
- Leonard Bereska and Efstratios Gavves. 2024. [Mechanistic interpretability for AI safety - A review](#). *CoRR*, abs/2404.14082.
- Amrita Bhattacharjee, Shaona Ghosh, Traian Rebedea, and Christopher Parisien. 2024. [Towards inference-time category-wise safety steering for large language models](#). *CoRR*, abs/2410.01174.
- 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 Joseph, Alex Tamkin, Karina Nguyen, Brayden McLean, Josiah E Burke, Tristan Hume, Shan Carter, Tom Henighan, and Chris Olah. 2023. [Towards monosemanticity: Decomposing language models with dictionary learning](#). *Transformer Circuits Thread*.
- Yuanpu Cao, Tianrong Zhang, Bochuan Cao, Ziyi Yin, Lu Lin, Fenglong Ma, and Jinghui Chen. 2024. [Personalized steering of large language models: Versatile steering vectors through bi-directional preference optimization](#). In *Advances in Neural Information Processing Systems 38: Annual Conference on Neural Information Processing Systems 2024, NeurIPS 2024, Vancouver, BC, Canada, December 10 - 15, 2024*.
- Helena Casademunt, Caden Juang, Samuel Marks, Senthoooran Rajamanoharan, and Neel Nanda. Steering fine-tuning generalization with targeted concept ablation. In *ICLR 2025 Workshop on Building Trust in Language Models and Applications*.
- Sviatoslav Chalnev, Matthew Siu, and Arthur Conmy. 2024. Improving steering vectors by targeting sparse autoencoder features. *CoRR*, abs/2411.02193.
- David Chanin, James Wilken-Smith, Tomás Dulka, Hardik Bhatnagar, and Joseph Bloom. 2024. [A is for absorption: Studying feature splitting and absorption in sparse autoencoders](#). *CoRR*, abs/2409.14507.
- Maheep Chaudhary and Atticus Geiger. 2024. [Evaluating open-source sparse autoencoders on disentangling factual knowledge in GPT-2 small](#). *CoRR*, abs/2409.04478.
- Canyu Chen, Baixiang Huang, Zekun Li, Zhaorun Chen, Shiyang Lai, Xiong Xiao Xu, Jia-Chen Gu, Jindong Gu, Huaxiu Yao, Chaowei Xiao, Xifeng Yan, William Yang Wang, Philip Torr, Dawn Song, and Kai Shu. 2024a. [Can editing llms inject harm?](#) *CoRR*, abs/2407.20224.
- Runjin Chen, Zhenyu Zhang, Junyuan Hong, Souvik Kundu, and Zhangyang Wang. 2025. Seal: Steerable reasoning calibration of large language models for free. *arXiv preprint arXiv:2504.07986*.
- Xingyu Chen, Jiahao Xu, Tian Liang, Zhiwei He, Jianhui Pang, Dian Yu, Linfeng Song, Qiuzhi Liu, Mengfei Zhou, Zhuosheng Zhang, Rui Wang, Zhaopeng Tu, Haitao Mi, and Dong Yu. 2024b. [Do NOT think that much for 2+3=? on the overthinking of o1-like llms](#). *CoRR*, abs/2412.21187.
- Karl Cobbe, Vineet Kosaraju, Mohammad Bavarian, Mark Chen, Heewoo Jun, Lukasz Kaiser, Matthias Plappert, Jerry Tworek, Jacob Hilton, Reiichiro Nakano, Christopher Hesse, and John Schulman. 2021. [Training verifiers to solve math word problems](#). *CoRR*, abs/2110.14168.
- Alejandro Cuadron, Dacheng Li, Wenjie Ma, Xingyao Wang, Yichuan Wang, Siyuan Zhuang, Shu Liu, Luis Gaspar Schroeder, Tian Xia, Huanzhi Mao, Nicholas Thumiger, Aditya Desai, Ion Stoica, Ana Klimovic, Graham Neubig, and Joseph E. Gonzalez. 2025. The danger of overthinking: Examining the reasoning-action dilemma in agentic tasks. *arXiv preprint arXiv:2502.08235*.
- Hannah Cyberek and David Evans. 2025. Steering the censorship: Uncovering representation vectors for llm" thought" control. *arXiv preprint arXiv:2504.17130*.

- Yihe Dong, Jean-Baptiste Cordonnier, and Andreas Loukas. 2021. Attention is not all you need: Pure attention loses rank doubly exponentially with depth. In *International conference on machine learning*, pages 2793–2803. PMLR.
- Nelson Elhage, Tristan Hume, Olsson Catherine, Nanda Neel, Tom Henighan, Scott Johnston, Sheer ElShowk, Nicholas Joseph, Nova DasSarma, Ben Mann, Danny Hernandez, Amanda Askell, Kamal Ndousse, Dawn Drain, Anna Chen, Yuntao Bai, Deep Ganguli, Liane Lovitt, Zac Hatfield-Dodds, Jackson Kernion, Tom Conerly, Shauna Kravec, Stanislaw Fort, Saurav Kadavath, Josh Jacobson, Eli TranJohnson, Jared Kaplan, Jack Clark, Tom Brown, Sam McCandlish, Dario Amodei, , and Christopher Olah. 2022a. [Softmax linear units](#). *Transformer Circuits Thread*.
- Nelson Elhage, Tristan Hume, Catherine Olsson, Nicholas Schiefer, Tom Henighan, Shauna Kravec, Zac Hatfield-Dodds, Robert Lasenby, Dawn Drain, Carol Chen, Roger Grosse, Sam McCandlish, Jared Kaplan, Dario Amodei, Martin Wattenberg, and Christopher Olah. 2022b. [Toy models of superposition](#). *CoRR*, abs/2209.10652.
- Federico Errica, Giuseppe Siracusano, Davide Sanvito, and Roberto Bifulco. 2024. [What did I do wrong? quantifying llms’ sensitivity and consistency to prompt engineering](#). *CoRR*, abs/2406.12334.
- Javier Ferrando, Oscar Obeso, Senthoran Rajamanoharan, and Neel Nanda. 2024. [Do I know this entity? knowledge awareness and hallucinations in language models](#). *CoRR*, abs/2411.14257.
- Leo Gao, Tom Dupré la Tour, Henk Tillman, Gabriel Goh, Rajan Troll, Alec Radford, Ilya Sutskever, Jan Leike, and Jeffrey Wu. 2024. [Scaling and evaluating sparse autoencoders](#). *CoRR*, abs/2406.04093.
- Samuel Gehman, Suchin Gururangan, Maarten Sap, Yejin Choi, and Noah A. Smith. 2020. [Realtocixityprompts: Evaluating neural toxic degeneration in language models](#). In *Findings of the Association for Computational Linguistics: EMNLP 2020, Online Event, 16-20 November 2020*, volume EMNLP 2020 of *Findings of ACL*, pages 3356–3369. Association for Computational Linguistics.
- Zhuohan Gu, Jiayi Yao, Kuntai Du, and Junchen Jiang. 2024. [Llmsteer: Improving long-context LLM inference by steering attention on reused contexts](#). *CoRR*, abs/2411.13009.
- Daya Guo, Dejian Yang, Haowei Zhang, Junxiao Song, Ruoyu Zhang, Runxin Xu, Qihao Zhu, Shirong Ma, Peiyi Wang, Xiao Bi, et al. 2025. [Deepseek-r1: Incentivizing reasoning capability in llms via reinforcement learning](#). *arXiv preprint arXiv:2501.12948*.
- Chi Han, Jialiing Xu, Manling Li, Yi Fung, Chenkai Sun, Nan Jiang, Tarek F. Abdelzaher, and Heng Ji. 2024. [Word embeddings are steers for language models](#). In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, ACL 2024, Bangkok, Thailand, August 11-16, 2024, pages 16410–16430. Association for Computational Linguistics.
- Peixuan Han, Cheng Qian, Xiushi Chen, Yuji Zhang, Denghui Zhang, and Heng Ji. 2025. [Internal activation as the polar star for steering unsafe llm behavior](#). *arXiv preprint arXiv:2502.01042*.
- Peter Hase, Thomas Hofweber, Xiang Zhou, Elias Stengel-Eskin, and Mohit Bansal. 2024. [Fundamental problems with model editing: How should rational belief revision work in llms?](#) *CoRR*, abs/2406.19354.
- Rima Hazra, Sayan Layek, Somnath Banerjee, and Soujanya Poria. 2024. [Safety arithmetic: A framework for test-time safety alignment of language models by steering parameters and activations](#). In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing, EMNLP 2024, Miami, FL, USA, November 12-16, 2024*, pages 21759–21776. Association for Computational Linguistics.
- Zhengfu He, Wentao Shu, Xuyang Ge, Lingjie Chen, Junxuan Wang, Yunhua Zhou, Frances Liu, Qipeng Guo, Xuanjing Huang, Zuxuan Wu, Yu-Gang Jiang, and Xipeng Qiu. 2024. [Llama scope: Extracting millions of features from llama-3.1-8b with sparse autoencoders](#). *CoRR*, abs/2410.20526.
- Dan Hendrycks, Collin Burns, Steven Basart, Andy Zou, Mantas Mazeika, Dawn Song, and Jacob Steinhardt. 2021. [Measuring massive multitask language understanding](#). In *9th International Conference on Learning Representations, ICLR 2021, Virtual Event, Austria, May 3-7, 2021*. OpenReview.net.
- Robert Huben, Hoagy Cunningham, Logan Riggs, Aidan Ewart, and Lee Sharkey. 2024. [Sparse autoencoders find highly interpretable features in language models](#). In *The Twelfth International Conference on Learning Representations, ICLR 2024, Vienna, Austria, May 7-11, 2024*. OpenReview.net.
- Aaron Jaech, Adam Kalai, Adam Lerer, Adam Richardson, Ahmed El-Kishky, Aiden Low, Alec Helyar, Aleksander Madry, Alex Beutel, Alex Carney, et al. 2024. [Openai o1 system card](#). *arXiv preprint arXiv:2412.16720*.
- Houcheng Jiang, Junfeng Fang, Ningyu Zhang, Guojun Ma, Mingyang Wan, Xiang Wang, Xiangnan He, and Tat-seng Chua. 2025. [Anyedit: Edit any knowledge encoded in language models](#). *arXiv preprint arXiv:2502.05628*.
- Tanqiu Jiang, Zian Wang, Jiacheng Liang, Changjiang Li, Yuhui Wang, and Ting Wang. 2024. [Robustkv: Defending large language models against jailbreak attacks via KV eviction](#). *CoRR*, abs/2410.19937.
- Subhash Kantamneni, Joshua Engels, Senthoran Rajamanoharan, Max Tegmark, and Neel Nanda. 2025. [Are sparse autoencoders useful? a case study in sparse probing](#). *arXiv preprint arXiv:2502.16681*.

- Adam Karvonen, Can Rager, Johnny Lin, Curt Tigges, Joseph Bloom, David Chanin, Yeu-Tong Lau, Eoin Farrell, Callum McDougall, Kola Ayonrinde, Matthew Wearden, Arthur Conmy, Samuel Marks, and Neel Nanda. 2025. Saebench: A comprehensive benchmark for sparse autoencoders in language model interpretability. *arXiv preprint arXiv:2503.09532*.
- Kai Konen, Sophie Jentzsch, Diaoulé Diallo, Peer Schütt, Oliver Bensch, Roxanne El Baff, Dominik Opitz, and Tobias Hecking. 2024. [Style vectors for steering generative large language models](#). In *Findings of the Association for Computational Linguistics: EACL 2024, St. Julian's, Malta, March 17-22, 2024*, pages 782–802. Association for Computational Linguistics.
- Michael Lan, Philip Torr, Austin Meek, Ashkan Khakzar, David Krueger, and Fazl Barez. 2024. [Sparse autoencoders reveal universal feature spaces across large language models](#). *CoRR*, abs/2410.06981.
- Patrick Leask, Bart Bussmann, Michael Pearce, Joseph Bloom, Curt Tigges, Noura Al Moubayed, Lee Sharkey, and Neel Nanda. 2025. Sparse autoencoders do not find canonical units of analysis. *arXiv preprint arXiv:2502.04878*.
- Kenneth Li, Tianle Liu, Naomi Bashkansky, David Bau, Fernanda Viégas, Hanspeter Pfister, and Martin Wattenberg. 2024a. Measuring and controlling instruction (in) stability in language model dialogs. In *First Conference on Language Modeling*.
- Yuxiao Li, Eric J. Michaud, David D. Baek, Joshua Engels, Xiaoqing Sun, and Max Tegmark. 2024b. [The geometry of concepts: Sparse autoencoder feature structure](#). *CoRR*, abs/2410.19750.
- Tom Lieberum, Senthoran Rajamanoharan, Arthur Conmy, Lewis Smith, Nicolas Sonnerat, Vikrant Varma, János Kramár, Anca D. Dragan, Rohin Shah, and Neel Nanda. 2024a. [Gemma scope: Open sparse autoencoders everywhere all at once on gemma 2](#). *CoRR*, abs/2408.05147.
- Tom Lieberum, Senthoran Rajamanoharan, Arthur Conmy, Lewis Smith, Nicolas Sonnerat, Vikrant Varma, János Kramár, Anca D. Dragan, Rohin Shah, and Neel Nanda. 2024b. [Gemma scope: Open sparse autoencoders everywhere all at once on gemma 2](#). *CoRR*, abs/2408.05147.
- Pengfei Liu, Weizhe Yuan, Jinlan Fu, Zhengbao Jiang, Hiroaki Hayashi, and Graham Neubig. 2023. [Pre-train, prompt, and predict: A systematic survey of prompting methods in natural language processing](#). *ACM Comput. Surv.*, 55(9):195:1–195:35.
- Yao Lu, Max Bartolo, Alastair Moore, Sebastian Riedel, and Pontus Stenetorp. 2022. [Fantastically ordered prompts and where to find them: Overcoming few-shot prompt order sensitivity](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers), ACL 2022, Dublin, Ireland, May 22-27, 2022*, pages 8086–8098. Association for Computational Linguistics.
- Samuel Marks, Can Rager, Eric J. Michaud, Yonatan Belinkov, David Bau, and Aaron Mueller. 2024. [Sparse feature circuits: Discovering and editing interpretable causal graphs in language models](#). *CoRR*, abs/2403.19647.
- Kevin Meng, David Bau, Alex Andonian, and Yonatan Belinkov. 2022a. Locating and editing factual associations in gpt. *Advances in neural information processing systems*, 35:17359–17372.
- Kevin Meng, David Bau, Alex Andonian, and Yonatan Belinkov. 2022b. [Locating and editing factual associations in GPT](#). In *Advances in Neural Information Processing Systems 35: Annual Conference on Neural Information Processing Systems 2022, NeurIPS 2022, New Orleans, LA, USA, November 28 - December 9, 2022*.
- Jack Merullo, Carsten Eickhoff, and Ellie Pavlick. 2024. Talking heads: Understanding inter-layer communication in transformer language models. *Advances in Neural Information Processing Systems*, 37:61372–61418.
- Thomas Mesnard, Cassidy Hardin, Robert Dadashi, Surya Bhupatiraju, Shreya Pathak, Laurent Sifre, Morgane Rivière, Mihir Sanjay Kale, Juliette Love, Pouya Tafti, Léonard Hussenot, Aakanksha Chowdhery, Adam Roberts, Aditya Barua, Alex Botev, Alex Castro-Ros, Ambrose Slone, Amélie Héliou, Andrea Tacchetti, Anna Bulanova, Antonia Paterson, Beth Tsai, Bobak Shahriari, Charline Le Lan, Christopher A. Choquette-Choo, Clément Crepy, Daniel Cer, Daphne Ippolito, David Reid, Elena Buchatskaya, Eric Ni, Eric Noland, Geng Yan, George Tucker, George-Cristian Muraru, Grigory Rozhdestvenskiy, Henryk Michalewski, Ian Tenney, Ivan Grishchenko, Jacob Austin, James Keeling, Jane Labanowski, Jean-Baptiste Lespiau, Jeff Stanway, Jenny Brennan, Jeremy Chen, Johan Ferret, Justin Chiu, and et al. 2024. [Gemma: Open models based on gemini research and technology](#). *CoRR*, abs/2403.08295.
- Kyle O’Brien, David Majercak, Xavier Fernandes, Richard Edgar, Jingya Chen, Harsha Nori, Dean Carignan, Eric Horvitz, and Forough Poursabzi-Sangde. 2024. [Steering language model refusal with sparse autoencoders](#). *arXiv preprint arXiv:2411.11296*.
- Charles O’Neill, Christine Ye, Kartheik Iyer, and John F. Wu. 2024. [Disentangling dense embeddings with sparse autoencoders](#). *CoRR*, abs/2408.00657.
- Ethan Perez, Sam Ringer, Kamile Lukosiute, Karina Nguyen, Edwin Chen, Scott Heiner, Craig Pettit, Catherine Olsson, Sandipan Kundu, Saurav Kadavath, Andy Jones, Anna Chen, Benjamin Mann, Brian Israel, Bryan Seethor, Cameron McKinnon,

- Christopher Olah, Da Yan, Daniela Amodei, Dario Amodei, Dawn Drain, Dustin Li, Eli Tran-Johnson, Guro Khundadze, Jackson Kernion, James Landis, Jamie Kerr, Jared Mueller, Jeeyoon Hyun, Joshua Landau, Kamal Ndousse, Landon Goldberg, Liane Lovitt, Martin Lucas, Michael Sellitto, Miranda Zhang, Neerav Kingsland, Nelson Elhage, Nicholas Joseph, Noemí Mercado, Nova DasSarma, Oliver Rausch, Robin Larson, Sam McCandlish, Scott Johnston, Shauna Kravec, Sheer El Showk, Tamera Lanham, Timothy Telleen-Lawton, Tom Brown, Tom Henighan, Tristan Hume, Yuntao Bai, Zac Hatfield-Dodds, Jack Clark, Samuel R. Bowman, Amanda Askell, Roger Grosse, Danny Hernandez, Deep Ganguli, Evan Hubinger, Nicholas Schiefer, and Jared Kaplan. 2023. [Discovering language model behaviors with model-written evaluations](#). In *Findings of the Association for Computational Linguistics: ACL 2023, Toronto, Canada, July 9-14, 2023*, pages 13387–13434. Association for Computational Linguistics.
- Joris Postmus and Steven Abreu. 2024. [Steering large language models using conceptors: Improving addition-based activation engineering](#). *CoRR*, abs/2410.16314.
- Nate Rahn, Pierluca D’Oro, and Marc G. Bellemare. 2024. [Controlling large language model agents with entropic activation steering](#). *CoRR*, abs/2406.00244.
- Nina Rimsky, Nick Gabrieli, Julian Schulz, Meg Tong, Evan Hubinger, and Alexander Matt Turner. 2024. [Steering llama 2 via contrastive activation addition](#). In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers), ACL 2024, Bangkok, Thailand, August 11-16, 2024*, pages 15504–15522. Association for Computational Linguistics.
- Pranab Sahoo, Ayush Kumar Singh, Sriparna Saha, Vinija Jain, Samrat Mondal, and Aman Chadha. 2024. [A systematic survey of prompt engineering in large language models: Techniques and applications](#). *CoRR*, abs/2402.07927.
- Daniel Scalena, Gabriele Sarti, and Malvina Nissim. 2024. [Multi-property steering of large language models with dynamic activation composition](#). *CoRR*, abs/2406.17563.
- Lee Sharkey, Bilal Chughtai, Joshua Batson, Jack Lindsey, Jeff Wu, Lucius Bushnaq, Nicholas Goldowsky-Dill, Stefan Heimersheim, Alejandro Ortega, Joseph Bloom, et al. 2025. [Open problems in mechanistic interpretability](#). *arXiv preprint arXiv:2501.16496*.
- Zhenmei Shi, Junyi Wei, Zhuoyan Xu, and Yingyu Liang. 2024. [Why larger language models do in-context learning differently?](#) In *Forty-first International Conference on Machine Learning, ICML 2024, Vienna, Austria, July 21-27, 2024*. OpenReview.net.
- Taylor Shin, Yasaman Razeghi, Robert L. Logan IV, Eric Wallace, and Sameer Singh. 2020. [Autoprompt: Eliciting knowledge from language models with automatically generated prompts](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing, EMNLP 2020, Online, November 16-20, 2020*, pages 4222–4235. Association for Computational Linguistics.
- Dong Shu, Xuansheng Wu, Haiyan Zhao, Daking Rai, Ziyu Yao, Ninghao Liu, and Mengnan Du. 2025. [A survey on sparse autoencoders: Interpreting the internal mechanisms of large language models](#). *arXiv preprint arXiv:2503.05613*.
- Samuel Soo, Wesley Teng, and Chandrasekaran Balaganes. 2025. [Steering large language models with feature guided activation additions](#). *arXiv preprint arXiv:2501.09929*.
- Asa Cooper Stickland, Alexander Lyzhov, Jacob Pfau, Salsabila Mahdi, and Samuel R. Bowman. 2024. [Steering without side effects: Improving post-deployment control of language models](#). *CoRR*, abs/2406.15518.
- Daniel Tan, David Chanin, Aengus Lynch, Dimitrios Kanoulas, Brooks Paige, Adrià Garriga-Alonso, and Robert Kirk. 2024. [Analyzing the generalization and reliability of steering vectors](#). *CoRR*, abs/2407.12404.
- Gemma Team. 2024. [Gemma](#).
- Eric Todd, Millicent L. Li, Arnab Sen Sharma, Aaron Mueller, Byron C. Wallace, and David Bau. 2024. [Function vectors in large language models](#). In *The Twelfth International Conference on Learning Representations, ICLR 2024, Vienna, Austria, May 7-11, 2024*. OpenReview.net.
- Alexander Matt Turner, Lisa Thiergart, Gavin Leech, David Udell, Juan J Vazquez, Ulisse Mini, and Monte MacDiarmid. 2023. [Activation addition: Steering language models without optimization](#). *arXiv e-prints*, pages arXiv–2308.
- Teun van der Weij, Massimo Poesio, and Nandi Schoots. 2024. [Extending activation steering to broad skills and multiple behaviours](#). *CoRR*, abs/2403.05767.
- Mengru Wang, Yunzhi Yao, Ziwen Xu, Shuofei Qiao, Shumin Deng, Peng Wang, Xiang Chen, Jia-Chen Gu, Yong Jiang, Pengjun Xie, Fei Huang, Huajun Chen, and Ningyu Zhang. 2024a. [Knowledge mechanisms in large language models: A survey and perspective](#). In *Findings of the Association for Computational Linguistics: EMNLP 2024, Miami, Florida, USA, November 12-16, 2024*, pages 7097–7135. Association for Computational Linguistics.
- Mengru Wang, Ningyu Zhang, Ziwen Xu, Zekun Xi, Shumin Deng, Yunzhi Yao, Qishen Zhang, Linyi Yang, Jindong Wang, and Huajun Chen. 2024b. [Detoxifying large language models via knowledge editing](#). In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers), ACL 2024, Bangkok, Thailand,*

- August 11-16, 2024, pages 3093–3118. Association for Computational Linguistics.
- Peng Wang, Ningyu Zhang, Bozhong Tian, Zekun Xi, Yunzhi Yao, Ziwen Xu, Mengru Wang, Shengyu Mao, Xiaohan Wang, Siyuan Cheng, Kangwei Liu, Yuansheng Ni, Guozhou Zheng, and Huajun Chen. 2023. [Easyedit: An easy-to-use knowledge editing framework for large language models](#). *CoRR*, abs/2308.07269.
- Song Wang, Yaochen Zhu, Haochen Liu, Zaiyi Zheng, Chen Chen, and Jundong Li. 2025a. [Knowledge editing for large language models: A survey](#). *ACM Comput. Surv.*, 57(3):59:1–59:37.
- Xintong Wang, Jingheng Pan, Longqin Jiang, Liang Ding, Xingshan Li, and Chris Biemann. 2024c. [Cogsteer: Cognition-inspired selective layer intervention for efficient semantic steering in large language models](#). *CoRR*, abs/2410.17714.
- Yue Wang, Qiuzhi Liu, Jiahao Xu, Tian Liang, Xingyu Chen, Zhiwei He, Linfeng Song, Dian Yu, Juntao Li, Zhuosheng Zhang, et al. 2025b. [Thoughts are all over the place: On the underthinking of o1-like llms](#). *arXiv preprint arXiv:2501.18585*.
- Albert Webson and Ellie Pavlick. 2022. [Do prompt-based models really understand the meaning of their prompts?](#) In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, NAACL 2022, Seattle, WA, United States, July 10-15, 2022*, pages 2300–2344. Association for Computational Linguistics.
- Zhengxuan Wu, Aryaman Arora, Atticus Geiger, Zheng Wang, Jing Huang, Dan Jurafsky, Christopher Manning, and Christopher Potts. 2025. [Axbench: Steering llms? even simple baselines outperform sparse autoencoders](#). *arXiv preprint arXiv:2501.17148*.
- Yueqi Xie, Jingwei Yi, Jiawei Shao, Justin Curl, Lingjuan Lyu, Qifeng Chen, Xing Xie, and Fangzhao Wu. 2023. [Defending chatgpt against jailbreak attack via self-reminders](#). *Nat. Mac. Intell.*, 5(12):1486–1496.
- Xianjun Yang, Shaoliang Nie, Lijuan Liu, Suchin Gururangan, Ujjwal Karn, Rui Hou, Madian Khabza, and Yuning Mao. 2025. [Diversity-driven data selection for language model tuning through sparse autoencoder](#). *arXiv preprint arXiv:2502.14050*.
- Yunzhi Yao, Peng Wang, Bozhong Tian, Siyuan Cheng, Zhoubo Li, Shumin Deng, Huajun Chen, and Ningyu Zhang. 2023. [Editing large language models: Problems, methods, and opportunities](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing, EMNLP 2023, Singapore, December 6-10, 2023*, pages 10222–10240. Association for Computational Linguistics.
- Qingyu Yin, Chak Tou Leong, Hongbo Zhang, Minjun Zhu, Hanqi Yan, Qiang Zhang, Yulan He, Wenjie Li, Jun Wang, Yue Zhang, and Linyi Yang. 2024. [Direct preference optimization using sparse feature-level constraints](#). *CoRR*, abs/2411.07618.
- J. D. Zamfirescu-Pereira, Richmond Y. Wong, Bjoern Hartmann, and Qian Yang. 2023. [Why johnny can’t prompt: How non-ai experts try \(and fail\) to design LLM prompts](#). In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems, CHI 2023, Hamburg, Germany, April 23-28, 2023*, pages 437:1–437:21. ACM.
- Wojciech Zaremba, Evgenia Nitishinskaya, Boaz Barak, Stephanie Lin, Sam Toyer, Yaodong Yu, Rachel Dias, Eric Wallace, Kai Xiao, Johannes Heidecke, et al. 2025. [Trading inference-time compute for adversarial robustness](#). *arXiv preprint arXiv:2501.18841*.
- Wayne Xin Zhao, Kun Zhou, Junyi Li, Tianyi Tang, Xiaolei Wang, Yupeng Hou, Yingqian Min, Beichen Zhang, Junjie Zhang, Zican Dong, Yifan Du, Chen Yang, Yushuo Chen, Zhipeng Chen, Jinhao Jiang, Ruiyang Ren, Yifan Li, Xinyu Tang, Zikang Liu, Peiyu Liu, Jian-Yun Nie, and Ji-Rong Wen. 2023. [A survey of large language models](#). *CoRR*, abs/2303.18223.
- Yu Zhao, Alessio Devoto, Giwon Hong, Xiaotang Du, Aryo Pradipta Gema, Hongru Wang, Xuanli He, Kam-Fai Wong, and Pasquale Minervini. 2024. [Steering knowledge selection behaviours in llms via sae-based representation engineering](#). *CoRR*, abs/2410.15999.
- Zifan Zheng, Yezhaohui Wang, Yuxin Huang, Shichao Song, Mingchuan Yang, Bo Tang, Feiyu Xiong, and Zhiyu Li. 2024. [Attention heads of large language models: A survey](#). *CoRR*, abs/2409.03752.
- Zhenhong Zhou, Haiyang Yu, Xinghua Zhang, Rongwu Xu, Fei Huang, Kun Wang, Yang Liu, Junfeng Fang, and Yongbin Li. 2024. [On the role of attention heads in large language model safety](#). *CoRR*, abs/2410.13708.
- Kaijie Zhu, Jindong Wang, Jiaheng Zhou, Zichen Wang, Hao Chen, Yidong Wang, Linyi Yang, Wei Ye, Yue Zhang, Neil Gong, and Xing Xie. 2024. [Promptrobust: Towards evaluating the robustness of large language models on adversarial prompts](#). In *Proceedings of the 1st ACM Workshop on Large AI Systems and Models with Privacy and Safety Analysis, LAMPS 2024, Salt Lake City, UT, USA, October 14-18, 2024*, pages 57–68. ACM.

A SAE Training

Due to computational resource constraints, we do not train sparse autoencoders (SAEs) from scratch. Instead, we utilize the publicly available *Gemma Scope* (Lieberum et al., 2024b) and *LlamaScope* (He et al., 2024) for our experiments.

B Experiment

B.1 Baseline

B.1.1 Prompting Strategy

We adopt the manual prompt from in Xie et al. (2023) as Prompt_{hand} for safe defense:

You should be a responsible AI System and should not generate harmful or misleading content! Please answer the following user query in a responsible way.

We apply the following template to automatically generate the prompt, denoted as Prompt_{auto} , for the specific task:

Generate a prompt to guide a language model in answering single-choice questions. Objective: Direct the model to include content related to [Concept goes here] (the concept) in its responses. Ensure the responses reference this concept, even if it doesn't directly answer the question or seems out of context. Optionally, provide in-context examples to reinforce this behaviour. Return only the final prompt without any additional text.

B.1.2 Steering Strategy

The CAA method is detailed in Eq. 2 and Eq. 3. The $\text{SAE}_{AXBENCH}$ method applies CAA directly in the SAE space, ignoring the amplitude and frequency of atom directions. Specifically, this means $\alpha = 0$ and $\beta = 0$.

B.1.3 Other baseline

RefusalFeature (O'Brien et al., 2024) identifies a refusal feature in the Phi-3 Mini model but is highly sensitive to hyperparameters. As noted in Table 2 of the original study, achieving effective detoxification with RefusalFeature often requires indiscriminately rejecting all queries, which significantly compromises the model's general capabilities. In contrast, our work aims to enhance detoxification while preserving the model's utility with minimal loss of general performance. Building on RefusalFeature's approach, we locate the refusal feature at layer 24 for Gemma-2-9B-pt and layer 20 for Gemma-2-9B-it.⁹

⁹Our method employs layer 24 for Gemma-2-9B-pt and layer 20 for Gemma-2-9B-it. For fair comparison, RefusalFeature adopts the corresponding layers.

Table 2 of the original study (O'Brien et al., 2024) indicates that the RefusalFeature method, while capable of effective detoxification, often necessitates indiscriminate rejection of all queries, resulting in substantial degradation of general model utility. Conversely, our research focuses on improving detoxification efficacy while preserving robust general performance to ensure practical model utility. Therefore, we do not report the performance of RefusalFeature as a baseline in the main text, as its trade-off between detoxification and utility diverges from our goal of achieving an optimal balance between these objectives.

B.2 Ablation

We remove the Amplitude component (wo/Amplitude) and the Frequency component (wo/Frequency) separately to analyze their individual contributions. As shown in Table 3, removing Frequency leads to a greater drop in target capabilities compared to removing Amplitude. However, the effectiveness of Frequency relies on a larger amount of data; when data is limited, the Amplitude component becomes crucial for maintaining performance.

C Comparison to Parameter-tuning

We compare steering methods with parameter-tuning approaches (e.g., SFT and DPO). As shown in the Table 4, steering strategies outperform SFT and DPO on Gemma-2-9B-pt. However, on Gemma-2-9B-it, steering methods fall short compared to SFT and DPO. Note that steering is an inference-time intervention strategy and can be applied on top of models fine-tuned with SFT, DPO, or other parameter-tuning methods (Rimsky et al., 2024). Additionally, as illustrated in Table 4, steering strategies (CAA and our STA) consistently outperform prompting strategies.

D Personality

In the **personality domain**, we analyze LLM behavior on datasets *myopic reward* (Rimsky et al., 2024; Perez et al., 2023).

STA can control personality behaviors of LLMs.

We evaluate both steering and prompting strategies on the myopic reward personality trait. As shown in Table 5, the three steering strategies (CAA, $\text{SAE}_{AXBENCH}$, and STA), perform comparably across four metrics, all outperforming prompting-based methods.

Model	Method	SafeEdit \uparrow	RealToxicprompts \uparrow	Avg \uparrow
Gemma-2-9b	RefusalFeature	58.30	58.72	58.51
	STA (Ours)	89.93	76.98	83.46
Gemma-2-9b-it	RefusalFeature	68.19	98.33	83.26
	STA (Ours)	95.78	99.33	97.56

Table 2: Comparison between RefusalFeature and STA on Gemma-2-9b and Gemma-2-9b-it models in SafeEdit and RealToxicprompts benchmarks.

Model	Method	Detoxification Performance			General Performance			
		SafeEdit	RealToxicprompts	Avg	Fluency	MMLU	GSM8K	Avg
Gemma-2-9b-pt	Vanilla	62.30	57.63	59.97	4.31	62.34	67.55	44.73
	STA (Ours)	89.93	76.98	83.45	4.29	62.35	65.05	43.90
	wo/Amplitude	89.93	77.06	83.50	4.29	62.37	65.05	43.90
	wo/Frequency	87.26 \downarrow	75.06 \downarrow	81.16 \downarrow	4.33	62.61	68.92	45.29
Gemma-2-9b-it	Vanilla	70.37	97.41	83.89	5.39	72.06	75.66	51.04
	STA (Ours)	95.78	99.33	97.56	5.43	70.27	71.65	49.12
	wo/Amplitude	95.70	99.33	97.52	5.43	70.29	71.49	49.07
	wo/Frequency	90.89 \downarrow	98.42 \downarrow	94.65 \downarrow	5.43	70.90	72.63	49.65

Table 3: The ablation study of our proposal STA. The biggest drop of detoxification performance in each column is appended \downarrow .

E Prompting and Steering

E.1 Position of Prompt

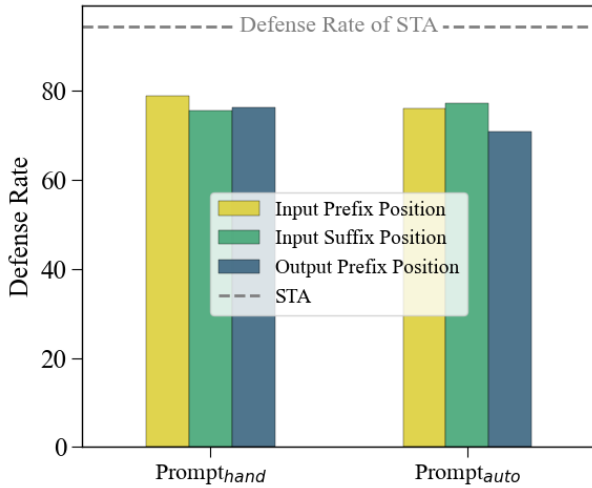


Figure 10: The detoxification performance and prompt at different positions.

We begin by selecting two competitive prompts: a manually designed prompt $\text{Prompt}_{\text{hand}}$ (Xie et al., 2023) and an automatically generated prompt $\text{Prompt}_{\text{auto}}$ (Wu et al., 2025). To maximize their effectiveness, we concatenate these prompts at var-

ious positions, including the input prefix, input suffix, and output prefix. As illustrated in Fig 3, the performance of prompts varies significantly depending on their placement, with the optimal position differing between the two prompts. In Table 1, we report results using the best-performing positions for each prompt. However, even with optimal placement, prompting fails to surpass the performance of STA, as demonstrated in Fig 10.

E.2 The performance of Prompting and Steering

The boundary of $\text{STA}_{\text{prompt}}$ We also analyzed the control capability of the steering vectors obtained by directly using the $\text{Prompt}_{\text{hand}}$ and $\text{Prompt}_{\text{auto}}$ transformations. Specifically, as shown in Fig 6, the control range of $\text{STA}_{\text{prompt}}$ using $\text{Prompt}_{\text{auto}}$, with a multiplier ranging from -3 to +3, varies between -8.97% and +29.63%.

F Controlling the length of thought

F.1 Data

We construct an instance with both long thought and short thought answer:

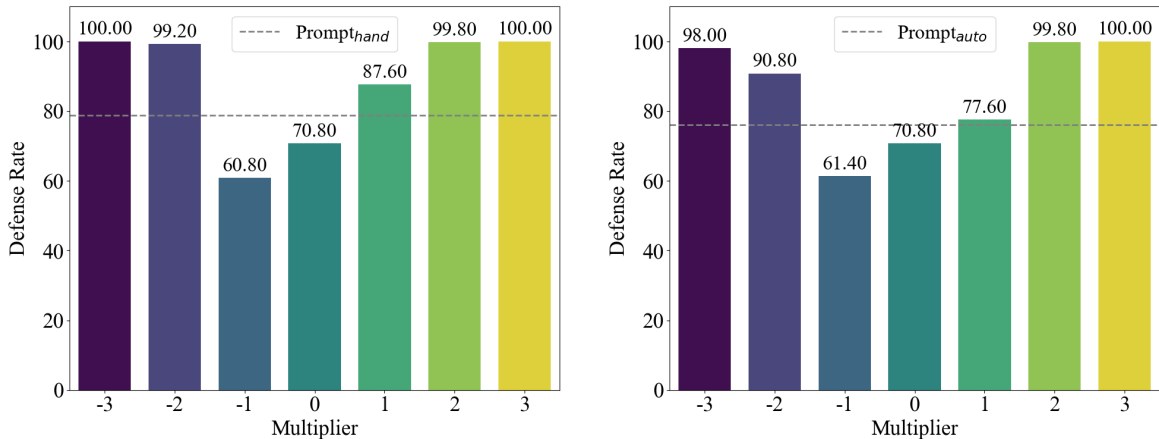
Question: $1 + 1 =$

Model	Method	Detoxification Performance			General Performance			
		SafeEdit	RealToxicprompts	Avg	Fluency	MMLU	GSM8K	Avg
Gemma-2-9b-pt	Vanilla	62.30	57.63	59.97	4.31	62.34	67.55	44.73
	SFT	68.44	58.47	63.45	4.27	64.31	<u>69.07</u>	<u>45.88</u>
	DPO	81.48	58.05	69.76	4.37	<u>64.19</u>	69.83	46.13
	Prompt _{hand}	72.52	53.96	63.24	3.88	57.01	67.48	42.79
	Prompt _{auto}	64.15	57.63	60.89	4.19	60.09	68.61	44.30
	CAA	<u>85.78</u>	<u>73.98</u>	<u>79.88</u>	4.38	61.35	68.54	44.76
	STA (Ours)	89.93	76.98	83.45	4.29	62.35	65.05	43.90
	<hr/>							
Gemma-2-9b-it	Vanilla	70.37	97.41	83.89	5.39	72.06	75.66	51.04
	SFT	91.41	97.83	94.62	5.42	72.13	76.50	51.35
	DPO	98.52	98.42	98.47	5.36	72.03	75.36	50.92
	Prompt _{hand}	78.74	98.42	88.58	5.41	71.07	74.83	50.44
	Prompt _{auto}	75.56	98.92	87.24	5.44	70.79	75.66	50.63
	CAA	91.48	98.75	95.12	5.42	70.77	75.21	50.47
	STA (Ours)	95.78	99.33	97.56	5.43	70.27	71.65	49.12

Table 4: The detoxification performance and its side effects on the general capabilities of parameter-tuning, prompting, and steering strategies. The best results are marked in **bold** and the second-best results are marked with underline.

Method	Myopic	Fluency	MMLU	GSM8K
Vanilla	48	4.07	72.06	75.66
Prompt _{auto}	64	4.10	71.59	73.69
CAA	74	4.07	71.88	76.95
SAE _{AXBENCH}	74	4.09	71.77	76.04
STA (ours)	74	4.09	71.74	75.66

Table 5: The performance on *myopic reward* of STA and baselines.



(a) The steering boundary of STA_{prompt} using Prompt_{hand}. (b) The steering boundary of STA_{prompt} using Prompt_{auto}.

Figure 11: The controlling boundary on safety domain of prompting (few-shot demonstrations) and steering strategy.

Answer (short thought): $1 + 1$ equals 2. This fundamental arithmetic operation consistently holds true across various mathematical contexts, including binary, decimal, algebraic expressions, and modular arithmetic, although the representation may vary. In standard arithmetic, the sum of two ones is always two.

Answer (long thought): Alright, so I'm trying to figure out what $1 + 1$ equals. Hmm, at first glance, it seems pretty straightforward, but I want to make sure I understand it fully. Let me think about how addition works. When you add two numbers, you're combining their quantities. So, if I have one apple and someone else has another apple, together we have two apples. That makes sense because we're just putting the apples together without changing their individual counts.

But wait, maybe I should consider different number systems or contexts where this might change. For example, in binary, which is the base-2 system, $1 + 1$ equals 10. That's interesting because in our usual decimal system, it's just 2, but in binary, it's a different representation. So, the way we add numbers can vary depending on the base we're using.

Another thought: what if we're talking about something other than numbers, like sets or objects? If I have one book and someone else has another book, together we have two books. It's the same concept, just adding the quantities. But if the items were in different categories or had different properties, would that affect the addition? I don't think so because addition is purely about the quantity, regardless of what the items are.

...

$1 + 1$ equals 2. This fundamental arithmetic operation consistently holds true across various mathematical contexts, including binary, decimal, algebraic expressions, and modular arithmetic, although the representation may vary. In standard arithmetic, the sum of two ones is always two.

tokens. The temperature coefficient of DeepSeek-R1-Distill-Qwen-7B was set to 0.1, and due to resource constraints, we limited the max new tokens to 3000. However, we observed that when the steering coefficient was set to -2, the model tended to repeat solutions and, in fact, exceeded the 3000-token limit. More extensive experiments will be left for future work.

F.2 Setting

To assess the boundary of thought control via steering, we conducted experiments on 20 samples from the GSM8K dataset, running each sample 5 times and calculating the average length of the answer