
Steer2Edit: From Activation Steering to Component-Level Editing

Chung-En Sun¹ Ge Yan^{*1} Zimo Wang^{*2} Tsui-Wei Weng²

Abstract

Steering methods influence Large Language Model behavior by identifying semantic directions in hidden representations, and are typically realized through inference-time activation interventions that apply a fixed, global modification to the model’s internal states. While effective, such interventions often induce unfavorable attribute–utility trade-offs under strong control, as they ignore the fact that many behaviors are governed by a small and heterogeneous subset of model components. To alleviate the trade-offs, we propose **STEER2EDIT**, a theoretically grounded, training-free framework that transforms steering vectors from inference-time control signals into diagnostic signals for component-level rank-1 weight editing. Instead of uniformly injecting a steering direction during generation, **STEER2EDIT** selectively redistributes behavioral influence across individual attention heads and MLP neurons, yielding interpretable edits that preserve the standard forward pass and remain compatible with optimized parallel inference. Across multiple tasks including safety alignment, truthfulness promotion, and reasoning efficiency, **STEER2EDIT** consistently achieves more favorable attribute–utility trade-offs: at matched downstream performance, it improves safety by up to 17.2%, increases truthfulness by 9.8%, and reduces reasoning length by 12.2% on average. Overall, **STEER2EDIT** provides a principled bridge between representation steering and weight editing by translating steering signals into *interpretable, training-free* parameter updates.

1. Introduction

Large Language Models (LLMs) have demonstrated strong capabilities across a wide range of tasks, including multi-step reasoning (Guo et al., 2025), code generation (Chen, 2021), and planning (Yao et al., 2023). As these models are increasingly deployed in real-world settings, there is growing interest in *controlling* specific model behaviors without retraining or fully fine-tuning the model.

A prominent line of recent work addresses this goal through *representation steering* (Zou et al., 2023a; Turner et al., 2023; Arditì et al., 2024; Yan et al., 2025; Li et al., 2025). These methods identify a *steering vector* in the model’s hidden representation space that correlates with a target attribute, and then intervene at inference time by adding this vector to intermediate activations. Compared to full fine-tuning, steering-based methods offer a lightweight way to adapt a model to different behaviors.

Despite their flexibility, activation-space steering methods suffer from two fundamental limitations. *First*, steering applies a *global* modification to the hidden representation. While such interventions can induce the target behavior, they treat all tokens and internal components uniformly, regardless of how the behavior is realized within the model. Empirical and mechanistic studies show that many behaviors are governed by a small and heterogeneous subset of model components (Olsson et al., 2022; Meng et al., 2022; Zhao et al., 2025; Zhou et al., 2024; Chu et al., 2025), typically involving specific attention heads or MLP neurons, while most components are only weakly related or unrelated to the target attribute. By ignoring this internal structure, global steering can interfere with unrelated semantic features, resulting in unfavorable trade-offs between the controlled attribute and downstream performance.

Second, activation-space steering relies on inference-time modification of intermediate activations. This departs from the standard forward pass assumed by modern optimized inference and training systems, which typically require fixed computation graphs. As a result, activation-level interventions complicate integration with standard deployment, parallel inference, and fine-tuning pipelines. While this limitation can in principle be mitigated with additional system engineering, activation steering remains an inference-time control mechanism whose effects are tied to the decoding

¹Department of Computer Science and Engineering, UC San Diego ²Halıcıoğlu Data Science Institute, UC San Diego. Correspondence to: Chung-En Sun <cesun@ucsd.edu>, Tsui-Wei Weng <lweng@ucsd.edu>.

process, rather than being encoded in the model parameters.

These limitations motivate a different perspective: instead of treating steering vectors as control signals to be directly injected into the forward pass, can we reinterpret them as *diagnostic signals* that reveal how a target behavior is distributed across model components? If so, can this information be used to selectively modify the components that genuinely govern the behavior, while avoiding unnecessary interference that degrades utility?

To answer these questions, we introduce **STEER2EDIT**, a theoretically grounded framework that converts steering vectors into *component-level weight edits*. In **STEER2EDIT**, a steering vector is treated as a diagnostic signal that reveals which attention heads and MLP neurons align with a target behavior, and to what extent. Guided by this signal, the method applies coordinated rank-1 updates to individual components, selectively amplifying or suppressing their contributions along the steering direction, rather than inducing a global activation shift. By redistributing behavioral influence at the component level, **STEER2EDIT** enables more precise behavioral control and more favorable attribute–utility trade-offs, while yielding interpretable component-level edits. The resulting procedure is closed-form, requires no fine-tuning or iterative optimization, and produces a standalone edited model that operates under the standard forward pass and remains compatible with existing training and optimized parallel inference pipelines.

Contributions.

- We propose **STEER2EDIT**, the first theoretically grounded framework that translates steering vectors into component-level rank-1 weight edits, requiring no fine-tuning and admitting a closed-form, single-step solution.
- We show that **STEER2EDIT** consistently achieves a superior attribute–performance trade-off compared to activation-level steering across diverse behavioral control settings: when matched for downstream performance, it improves safety by **17.2%**, truthfulness by **9.8%**, and, in the efficient reasoning setting, reduces reasoning length by **12.2%** on average.
- We show that **STEER2EDIT** produces a standalone edited model that preserves the original architecture, while offering fine-grained interpretability into which components govern specific behaviors and how these behaviors are distributed across the network.

2. Preliminary

In this section, we fix notation for the Transformer residual stream, define the steering vectors used throughout, and specify the editable weight components used in later **STEER2EDIT** analysis.

Transformer residual-stream updates. We consider a pre-normalization Transformer, where the residual stream is updated at layer ℓ according to

$$r_\ell^{\text{attn}} = r_{\ell-1}^{\text{mlp}} + \delta_\ell^{\text{attn}}, \quad \delta_\ell^{\text{attn}} := \text{Attn}(\text{LayerNorm}(r_{\ell-1}^{\text{mlp}})), \\ r_\ell^{\text{mlp}} = r_\ell^{\text{attn}} + \delta_\ell^{\text{mlp}}, \quad \delta_\ell^{\text{mlp}} := \text{MLP}(\text{LayerNorm}(r_\ell^{\text{attn}})).$$

Both $\delta_\ell^{\text{attn}}$ and δ_ℓ^{mlp} lie in the same residual-stream space \mathbb{R}^d .

Steering vector. Steering vectors are commonly used in activation steering, where a semantic direction in hidden representations is added to the residual stream at inference time to control model behavior. Such vectors can be constructed in various ways. For simplicity, we adopt a mean-difference construction in this work, while noting that **STEER2EDIT** is agnostic to how the steering vector is obtained.

Let \mathcal{X} denote a set of prompts. For each prompt $x \in \mathcal{X}$, the model generates a completion y , which is classified as exhibiting or not exhibiting the target attribute, yielding \mathcal{Y}_{pos} and \mathcal{Y}_{neg} .

At token position t of y , let $\delta_\ell^b(y, t) \in \mathbb{R}^d$, $b \in \{\text{attn}, \text{mlp}\}$, denote the output of the corresponding block at layer ℓ before it is written into the residual stream. Aggregating over token positions \mathcal{T}_y and averaging over generations, we define

$$\bar{\delta}_{\ell,a}^b = \frac{1}{|\mathcal{Y}_a|} \sum_{y \in \mathcal{Y}_a} \frac{1}{|\mathcal{T}_y|} \sum_{t \in \mathcal{T}_y} \delta_\ell^b(y, t), \quad a \in \{\text{pos}, \text{neg}\}.$$

The steering vector at layer ℓ and block b is given by the mean difference

$$v_\ell^b = \bar{\delta}_{\ell,\text{pos}}^b - \bar{\delta}_{\ell,\text{neg}}^b \in \mathbb{R}^d. \quad (1)$$

Editable weight components and notation. We focus on linear weight components whose outputs produce the block activations from which the steering vectors in Eq.(1) are extracted. Specifically, for each layer ℓ and block type $b \in \{\text{attn}, \text{mlp}\}$, we consider linear maps whose outputs contribute to the block output δ_ℓ^b before it is written into the residual stream.

Concretely, these editable components include: (i) the output projection (o_{proj}) of an individual attention head in the attention block, and (ii) the down-projection ($down_{\text{proj}}$) associated with a single neuron in the MLP block. We denote any such component generically by

$$W_i \in \mathbb{R}^{d_{\text{out}} \times d_{\text{in}}}, \quad d_{\text{out}} = d,$$

where the index i implicitly identifies a specific layer ℓ , block type b , and component within that block.

For an input activation $h_i \in \mathbb{R}^{d_{\text{in}}}$ to component W_i , the component output $W_i h_i$ lies in the same residual-stream space \mathbb{R}^d as the corresponding steering vector v_ℓ^b . Accordingly, in the subsequent analysis we associate each editable component W_i with the steering vector extracted from the same layer and block, and write this vector simply as v_i .

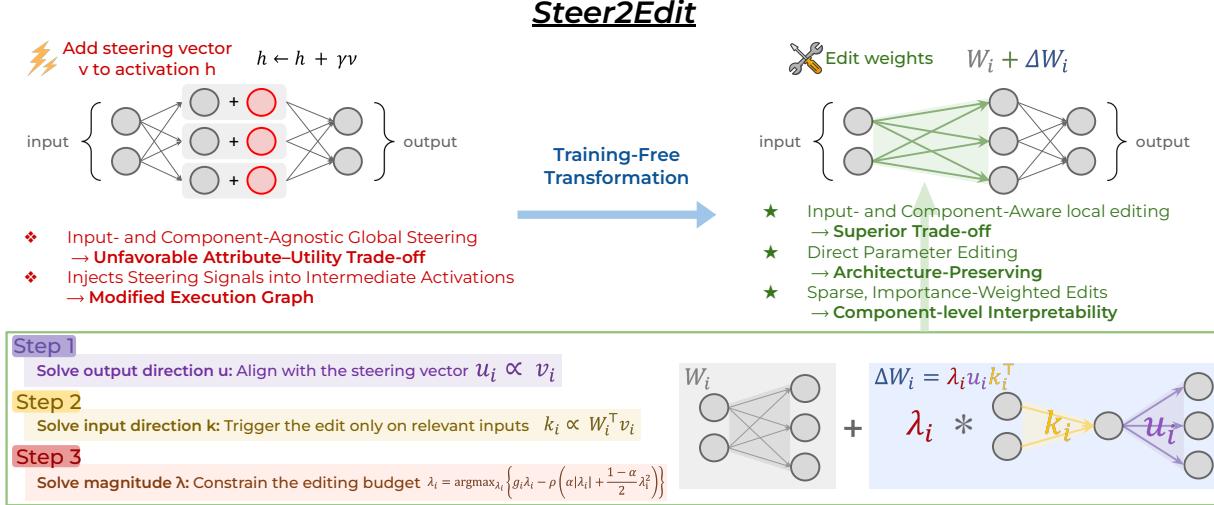


Figure 1. Overview of STEER2EDIT. STEER2EDIT converts the steering signal into component-level rank-1 weight edits. For each component, the edit $\Delta W_i = \lambda_i u_i k_i^T$ is constructed by aligning the output direction u_i with the steering vector, choosing an input direction k_i that triggers the edit only on relevant inputs, and allocating the magnitude λ_i under a global budget. The resulting edits are training-free, architecture-preserving, and interpretable.

3. Steer2Edit

In this section, we introduce **STEER2EDIT**, a principled framework for component-level weight editing based on given steering vectors. We parameterize each edit as a rank-1 update and derive its form by decomposing the problem into three parts: (i) identifying the output-space direction that preserves semantic invariance, (ii) the input-space direction that aligns the edit with the component’s intrinsic semantic contribution, and (iii) the scalar magnitude that allocates edit strength under a global regularization budget.

3.1. Assumption and Setting

For each editable component $W_i \in \mathbb{R}^{d_{\text{out}} \times d_{\text{in}}}$, we assume the existence of a steering vector $v_i \in \mathbb{R}^{d_{\text{out}}}$ extracted from the same representation space into which W_i writes (e.g., the hidden state after an attention or MLP block). Thus, the output dimension of W_i matches that of v_i , and both lie in a common semantic space. Note that **STEER2EDIT** is agnostic to how v_i is obtained.

Our goal is to modify each component W_i so that the resulting update ΔW_i alters the model’s behavior along the semantic direction represented by v_i . Because the steering signal specifies a single direction in representation space, we model each edit as a rank-1 perturbation, which is the minimal modification that can inject a directional effect into a linear map. Accordingly, we parameterize the edit as

$$\Delta W_i = \lambda_i u_i k_i^T,$$

where $u_i \in \mathbb{R}^{d_{\text{out}}}$ is an output-space direction, $k_i \in \mathbb{R}^{d_{\text{in}}}$ is an input-space direction, and $\lambda_i \in \mathbb{R}$ is a scalar magnitude controlling the strength of the edit.

Given an input activation h_i , let $o_i := W_i h_i \in \mathbb{R}^{d_{\text{out}}}$ denote the original output of component W_i . After applying the rank-1 edit $\Delta W_i = \lambda_i u_i k_i^T$, the edited output is $\tilde{o}_i := (W_i + \Delta W_i) h_i = o_i + \Delta o_i$, where the induced output shift is

$$\Delta o_i := \Delta W_i h_i = \lambda_i u_i (k_i^T h_i). \quad (2)$$

The three quantities (u_i, k_i, λ_i) play distinct roles in the edit. The output-space direction u_i specifies the semantic direction affected by the edit, the input-space direction k_i determines which inputs activate the edit through the inner product $k_i^T h_i$, and the scalar magnitude λ_i controls how strongly each component is modified.

We derive these quantities in a sequential order. We first determine u_i , then k_i , and finally solve for $\{\lambda_i\}_{i=1}^n$, the per-component edit magnitudes. As shown in the following sections, this ordering is without loss of generality: the optimal choice of u_i depends only on the steering vector v_i ; the choice of k_i depends on u_i and local properties of the component W_i ; and once the geometric directions are fixed, the magnitudes $\{\lambda_i\}$ can be optimized independently.

Hence, we derive the three components of the edit in the following order:

1. the output-space direction u_i in Section 3.2;
2. the input-space direction k_i in Section 3.3;
3. the scalar magnitude λ_i in Section 3.4.

Throughout the following derivations, we identify only the directions of u_i and k_i ; their scale and sign are absorbed into the scalar coefficients λ_i , which are determined in the last step.

3.2. Step 1: Solving for the Output-space Direction u_i

The vector u_i determines the *direction* of the output shift Δo_i . Because the steering signal specifies a single semantic direction v_i , we require that the edit modifies the component’s output *only* along this direction and introduces no change in any orthogonal subspace.

Formally, recall that $\Delta o_i = \Delta W_i h_i$. Semantic invariance requires that, for any input h_i , the output shift Δo_i has zero projection onto any direction orthogonal to v_i :

$$z^\top \Delta o_i = z^\top \Delta W_i h_i = 0 \quad \forall h_i, \forall z \perp v_i. \quad (3)$$

This constraint directly restricts the output-space direction u_i of the edit. The following theorem formalizes this restriction.

Theorem 3.1 (Output-space direction under semantic invariance). *Let $v_i \neq 0$, and let $\Delta W_i = \lambda_i u_i k_i^\top$ be a rank-1 edit with $\Delta W_i \neq 0$. If for all h_i and all $z \perp v_i$ we have*

$$z^\top \Delta W_i h_i = 0,$$

then the output-space direction u_i must be collinear with v_i , i.e.,

$$u_i \in \text{span}\{v_i\}.$$

We defer the proof to Appendix A.1. Theorem 3.1 shows that enforcing semantic invariance in Eq.(3) uniquely constrains the output-space direction: any valid rank-1 update must lie entirely along the steering direction v_i . Importantly, this result is independent of both the input-space direction k_i and the scalar magnitude λ_i . This separation justifies solving for u_i first in our derivation.

We therefore adopt the canonical normalized choice

$$u_i = \hat{v}_i := \frac{v_i}{\|v_i\|_2},$$

with the sign and scale absorbed later into the magnitude λ_i .

3.3. Step 2: Solving for the Input-space Direction k_i

Having solved the output-space direction $u_i = \hat{v}_i$ in Section 3.2, we now determine the input-space direction k_i for each editable component W_i . As before, we solve for the *direction* of k_i ; its sign and scale are absorbed into the scalar magnitude λ_i .

Intuition. As shown in Eq. (2), the input-space direction k_i determines which input activations h_i trigger the edit through the inner product $k_i^\top h_i$. To identify a suitable k_i , we note that a well-trained component W_i already encodes which inputs are relevant for contributing to the semantic direction v_i , and the edit ΔW_i should mirror this existing input-dependent pattern.

To formalize this intuition, we define the *semantic alignment score* of component W_i for an input h_i as

$$s_i(h_i) := v_i^\top o_i = v_i^\top W_i h_i,$$

where $o_i := W_i h_i$, which measures how strongly the original component output aligns with the target semantic direction v_i for a given input h_i . Intuitively, if $s_i(h_i)$ is small across inputs, this indicates that the component is generally unrelated to the semantic direction v_i , and the edit should be small for all inputs similarly. If, for some components, $s_i(h_i)$ is large for certain inputs, the edit should be large on those same inputs.

Hence, we choose k_i so that the induced change in the semantic alignment score, $\Delta s_i(h_i) := v_i^\top \Delta W_i h_i$, occurs on the same inputs for which $s_i(h_i)$ is large. To formalize this idea, we maximize the “absolute” Pearson correlation between $\Delta s_i(h_i)$ and $s_i(h_i)$, as we do not care about the sign or overall scale at this stage. The following theorem provides the solution.

Theorem 3.2 (Input-space direction matching semantic alignment variation). *Fix a component W_i and set $u_i = \hat{v}_i$. Assume $W_i^\top v_i \neq 0$ and $\text{Var}(s_i(h_i)) > 0$, where $s_i(h_i) := v_i^\top W_i h_i$. Consider choosing an input-direction $k_i \neq 0$ so that the induced semantic alignment shift $\Delta s_i(h_i) := v_i^\top \Delta W_i h_i$ exhibits maximal co-variation with the component’s intrinsic semantic alignment score $s_i(h_i)$. Formally, consider the objective*

$$\max_{k_i \neq 0} \left| \text{Pearson}(\Delta s_i(h_i), s_i(h_i)) \right|.$$

Then there exists a maximizer k_i that is collinear with $W_i^\top v_i$, i.e.,

$$k_i \in \text{span}\{W_i^\top v_i\}.$$

We defer the proof to Appendix A.2. Theorem 3.2 shows that the input-space direction k_i should align with the component’s intrinsic input sensitivity $W_i^\top v_i$. We therefore adopt the normalized choice

$$\hat{k}_i := \frac{W_i^\top v_i}{\|W_i^\top v_i\|_2}.$$

This choice is further empirically validated in Appendix E.

3.4. Step 3: Solving for the Edit Magnitudes λ

With the edit directions u_i and k_i fixed, we now determine the magnitudes $\{\lambda_i\}$, which control how strongly each component is reinforced or suppressed. Intuitively, the magnitude assigned to each component should reflect how that component contributes to the direction v_i *on average across inputs*: components that consistently align with v_i should be reinforced, components that consistently oppose it should be suppressed, and components with weak alignment should receive little or no edit.

Note that this role is fundamentally different from that of the input-space direction k_i , which captures how the component’s semantic alignment score $s_i(h_i)$ *varies across inputs*, whereas the magnitudes λ_i depend only on the component’s *overall, input-averaged* semantic alignment.

To formalize this allocation of editing strength, we now introduce an importance weighting for each component and derive λ_i via a global regularized optimization.

Importance weighting. Recall that the *semantic alignment score* $s_i(h_i) := v_i^\top W_i h_i$ measures how strongly component W_i contributes to the semantic direction v_i for a given input h_i . Since the magnitudes λ_i are intended to capture a component's *overall* semantic contribution, we measure this contribution by the expectation of the semantic alignment score over the input distribution: $\mathbb{E}[s_i(h_i)] = \mathbb{E}[v_i^\top W_i h_i]$. However, the typical output magnitude of $W_i h_i$ can vary substantially across layers, making raw values not directly comparable between components. To place components on a common footing, we make the following changes: (i) remove the arbitrary scale of the semantic direction by using $\hat{v}_i = v_i / \|v_i\|_2$, and (ii) normalize by the output norm of the mean activation $\mu_i = \mathbb{E}[h_i]$, yielding

$$g_i = \frac{\mathbb{E}[s_i(h_i)]}{\|v_i\|_2 \|W_i \mu_i\|_2} = \frac{v_i^\top W_i \mu_i}{\|v_i\|_2 \|W_i \mu_i\|_2} = \cos(v_i, W_i \mu_i),$$

which we refer to as the *component importance score*. The sign of g_i indicates whether the component aligns or opposes the semantic direction, while $|g_i|$ measures the strength of this tendency. This normalization choice is empirically validated in Appendix E.

Elastic-Net objective. The component importance score g_i indicates whether component W_i should be reinforced or suppressed, and with what strength. A natural objective is therefore to maximize total alignment $\mathbf{g}^\top \boldsymbol{\lambda} = \sum_{i=1}^n g_i \lambda_i$. However, this objective is unbounded, and effective weight editing should remain lightweight by modifying only a small number of relevant components while keeping edit magnitudes controlled.

We address both considerations with an Elastic-Net regularization, combining an ℓ_1 term to promote sparsity and an ℓ_2 term to limit overall edit size:

$$\max_{\boldsymbol{\lambda}} \mathbf{g}^\top \boldsymbol{\lambda} - \rho \left(\alpha \|\boldsymbol{\lambda}\|_1 + \frac{1-\alpha}{2} \|\boldsymbol{\lambda}\|_2^2 \right),$$

where $\rho > 0$ controls the global edit budget and $\alpha \in [0, 1]$ trades off ℓ_1 sparsity and ℓ_2 smoothness. Ablation results in Appendix E confirm the importance of both the ℓ_1 and ℓ_2 regularization terms.

Theorem 3.3 (Edit magnitude allocation under regularization). *For each component W_i , let $g_i = \cos(v_i, W_i \mu_i)$ denote the component importance score, with the convention that $g_i := 0$ if $W_i \mu_i = 0$.*

Consider the problem of assigning edit magnitudes $\{\lambda_i\}$ to maximize total signed alignment as measured by the component importance scores $\{g_i\}$, while controlling both

the sparsity and overall strength of the edit. Formally, let $\mathbf{g} = (g_1, \dots, g_n)$ and consider

$$\max_{\boldsymbol{\lambda} \in \mathbb{R}^n} \mathbf{g}^\top \boldsymbol{\lambda} - \rho \left(\alpha \|\boldsymbol{\lambda}\|_1 + \frac{1-\alpha}{2} \|\boldsymbol{\lambda}\|_2^2 \right), \rho > 0, \alpha \in [0, 1].$$

The unique edit magnitude assigned to component i is

$$\lambda_i^* = \text{sign}(g_i) \frac{\max(|g_i| - \rho\alpha, 0)}{\rho(1-\alpha)}.$$

We defer the proof to Appendix A.3. Theorem 3.3 yields a closed-form soft-thresholding rule for allocating edit magnitudes λ_i according to the alignment scores g_i under a global Elastic-Net budget. Substituting λ_i^* , together with $u_i = \hat{v}_i$ and $k_i = \hat{k}_i$, gives the unified weight-editing update below.

3.5. Summary: Unified Weight Editing Rule

Each editable component W_i receives the rank-1 update

$$\Delta W_i = \text{sign}(g_i) \frac{\max(|g_i| - \rho\alpha, 0)}{\rho(1-\alpha)} \hat{v}_i \hat{k}_i^\top,$$

where

$$\hat{v}_i := \frac{v_i}{\|v_i\|_2}, \quad \hat{k}_i = \frac{W_i^\top v_i}{\|W_i^\top v_i\|_2}, \quad g_i = \cos(v_i, W_i \mu_i).$$

Each update is:

- **Directionally selective:** it modifies only the projection along the semantic direction v_i ;
- **Input-selective:** it determines which inputs should trigger the edit based on the component's intrinsic semantic behavior;
- **Budget-aware:** its magnitude λ_i^* is determined by the component importance score g_i under an Elastic-Net regularizer.

Thus, **STEER2EDIT** yields a component-level weight editing framework that jointly captures *what* semantic direction to modify, *when* the edit should be activated by the input, and *how strongly* each component should be adjusted.

4. Experiments

This section evaluates whether **STEER2EDIT** achieves a superior *attribute–utility trade-off* compared to inference-time activation steering. We consider three representative behavioral control settings: (i) safety alignment against jailbreak attacks, (ii) truthfulness promotion, and (iii) reasoning efficiency control. Across all settings, we compare **STEER2EDIT** against the standard activation-steering baseline and report trade-offs between the target attribute and downstream utility.

4.1. Implementation of Steer2Edit and Baseline

We describe how the activation-steering and **STEER2EDIT** are instantiated, and how editing hyperparameters are selected.

Activation Steering (Baseline). Given a layer-wise steering vector v_ℓ , the hidden representation is modified as $h_\ell \leftarrow h_\ell + \gamma v_\ell$, where $\gamma \geq 0$ denotes the steering strength. We extract steering vectors separately for attention and MLP blocks at each layer, $\{v_\ell^{\text{attn}}, v_\ell^{\text{mlp}}\}_{\ell=1}^L$, and add the vector after each block during inference. We sweep γ to trace attribute–utility trade-off curves. Task-specific steering vector construction is deferred to Appendix B.

Steer2Edit. Following Section 3, we apply rank-1 edits to linear components that write directly into the residual stream. For attention, we edit each head’s output projection $W_o \in \mathbb{R}^{d_{\text{model}} \times d_{\text{head}}}$. For MLPs, we edit individual down-projection neurons by treating each column $w_{\text{down},j}$ of $W_{\text{down}} \in \mathbb{R}^{d_{\text{model}} \times d_{\text{ff}}}$ as an independent component.

Edit magnitudes are determined by an Elastic-Net objective with sparsity parameter α and global budgets ρ_{attn} and ρ_{mlp} for attention heads and MLP neurons, respectively. For each model and behavior setting, hyperparameters $(\rho_{\text{attn}}, \rho_{\text{mlp}}, \alpha)$ are selected via a small grid search on a held-out validation set, ranking configurations by the target attribute metric. Unless otherwise noted, results correspond to the best-performing or top-ranked configurations. Full search details are provided in Appendix C.

4.2. Evaluation for Behavioral Control

We consider three evaluation settings that examine behavioral control along distinct dimensions: safety alignment, truthfulness, and reasoning efficiency. In each setting, the *target attribute* is measured using task-specific metrics, while *downstream utility* is evaluated on task-oriented benchmarks that are unrelated to the controlled behavior.

For each use case, we visualize the trade-off between the target attribute and downstream utility. A method achieves a superior trade-off if it improves the target attribute while maintaining higher utility. Our experiments are designed to test whether **STEER2EDIT** can consistently achieve such favorable trade-offs relative to activation steering.

4.2.1. SAFETY ALIGNMENT AGAINST JAILBREAK ATTACKS

Goal. We evaluate whether **STEER2EDIT** can strengthen refusal behavior under strong jailbreak attacks while preserving helpfulness on benign tasks.

Models and evaluation. We evaluate safety alignment on **LLaMA-2-7B-Chat** and **Mistral-7B-Instruct-v0.2**. Safety is measured using **ADVBench**, which consists of harmful user queries designed to elicit unsafe behavior. Each query is transformed into a jailbreak prompt using either **GCG** (Zou et al., 2023b), a classical gradient-based attack, or **ADVL** (Sun et al., 2025a), a substantially stronger attack that

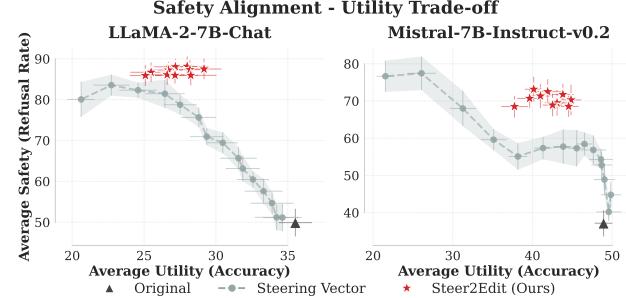


Figure 2. Safety–utility trade-off on **LLaMA-2-7B-Chat** and **Mistral-7B-Instruct-v0.2**. Each point corresponds to a different intervention strength. **STEER2EDIT** consistently attains higher refusal rates at comparable or higher utility, while strong steering-vector interventions incur substantial utility degradation.

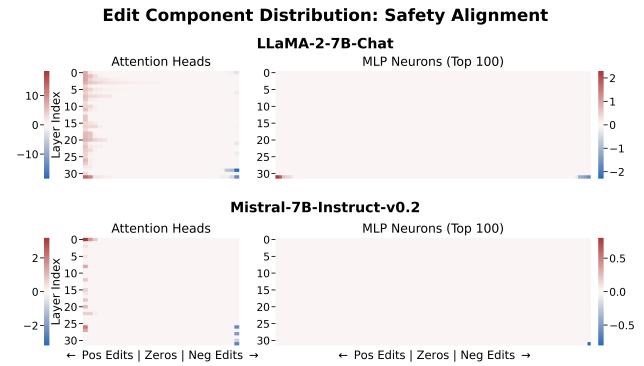


Figure 3. Signed **STEER2EDIT** edit coefficients λ for safety alignment. Positive (red) coefficients reinforce safety-aligned components, while negative (blue) coefficients suppress safety-opposing ones. Edits are highly sparse and concentrated in a small subset of attention heads, predominantly in later layers.

trains an LLM to generate adversarial suffixes. We report the **refusal rate**, defined as the proportion of model responses that are refusals, averaged across both attack types.

Downstream utility is evaluated on **GSM8K**, **CodeMMLU**, and **CommonsenseQA**, measuring utility on grade-school math reasoning, programming, and commonsense multiple-choice questions. Utility is reported as the mean accuracy across all three benchmarks.

Results. Figure 2 illustrates the safety–utility relationship. Each star corresponds to one of the top-10 most safety-aligned **STEER2EDIT** configurations. While steering-vector baselines trace a clear safety–utility trade-off as intervention strength increases, **STEER2EDIT** identifies configurations that lie beyond this trade-off frontier, occupying the top-right region of the plot and achieving higher refusal rates without sacrificing downstream utility.

Component-level analysis. Because **STEER2EDIT** applies edits at the level of individual components, the resulting weight updates are directly interpretable and reveal which components mediate the target behavior.

Figure 3 shows the signed edit coefficients λ for the best-performing safety-aligned configuration. Each cell corresponds to the strength of a rank-1 update applied to a specific component: positive values reinforce components aligned with refusal behavior, while negative values suppress components that oppose safety.

Across both models, non-zero λ values are highly sparse and concentrated in a small number of attention heads, predominantly in later layers. MLP neurons receive near-zero coefficients with only a few isolated exceptions. These results indicate that effective safety control is achieved through selective amplification and suppression of a small set of attention heads.

4.2.2. TRUTHFULNESS PROMOTION

Goal. We evaluate whether **STEER2EDIT** increases the model’s preference for truthful answers while preserving performance on unrelated downstream tasks.

Models and evaluation. We evaluate on **Gemma-2-2B-IT** and **LLaMA-3-8B-Instruct** using **TruthfulQA**. For each prompt, we measure whether the model assigns higher probability to the truthful answer than to a plausible but false alternative, and report truthful preference accuracy. Downstream utility is again measured on GSM8K, CodeMMLU, and CommonsenseQA.

Results. Figure 4 shows the truthfulness–utility relationship. While activation steering traces a clear trade-off in which stronger interventions rapidly degrade utility, **STEER2EDIT** achieves substantial truthfulness gains without incurring much utility loss.

Component-level analysis. Figure 5 visualizes the edit coefficients λ of the best-performing truthfulness-aligned configuration. Across both models, truthfulness control is sparse and predominantly mediated by attention heads, with non-zero edits concentrated in a limited number of layers. In contrast to safety alignment, truthfulness edits are distributed across both early and late layers. Notably, in **Gemma-2-2B-IT**, edits are dominated by negative coefficients, suggesting that truthfulness gains arise primarily from suppressing hallucination-promoting components rather than reinforcing truth-aligned ones. Overall, these patterns indicate that truthfulness can rely on markedly different internal circuits across models, while remaining amenable to selective, component-level intervention.

4.2.3. EFFICIENT REASONING

Goal. We evaluate whether **STEER2EDIT** can shorten reasoning traces while preserving answer accuracy, improving inference efficiency for Large Reasoning Models (LRMs).

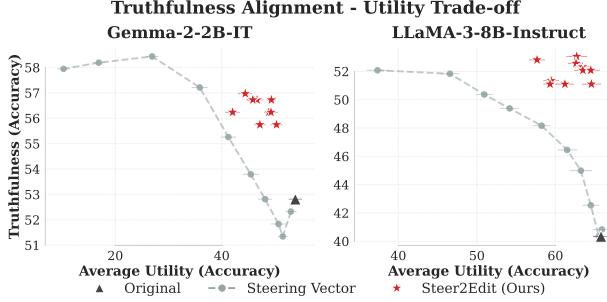


Figure 4. Truthfulness–utility trade-off on **Gemma-2-2B-IT** and **LLaMA-3-8B-Instruct**. **STEER2EDIT** improves truthfulness at a higher downstream utility than activation steering.

Edit Component Distribution: Truthfulness

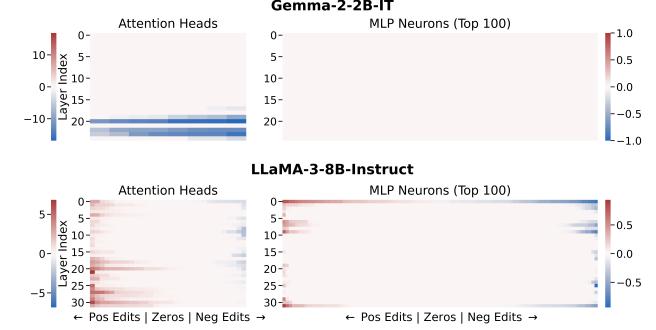


Figure 5. Signed **STEER2EDIT** edit coefficients λ for truthfulness promotion. Positive values reinforce truthfulness-aligned components, while negative values suppress components associated with hallucinated behavior.

Models and datasets. We evaluate on **Qwen3-4B-Thinking-2507** and **OpenMath-Nemotron-7B** using GSM8K, MATH-500, GPQA, and CodeMMLU. Downstream utility is measured as mean accuracy across all datasets, while reasoning efficiency is measured by the number of generated reasoning tokens.

Results. Figure 6 shows the accuracy–efficiency relationship. Across both models, activation steering reduces reasoning length only at the cost of substantial accuracy degradation. In contrast, **STEER2EDIT** significantly shortens reasoning traces while maintaining comparable accuracy.

Component-level analysis. Figure 7 visualizes the edit coefficients λ for the best-performing efficiency-oriented configuration and reveals a qualitatively different pattern from safety and truthfulness. Reasoning efficiency is predominantly mediated by MLP components, with dense, distributed edits spanning many neurons, while attention heads play a comparatively minor role. The most effective configurations correspond to larger ρ_{mlp} and smaller α , indicating that reducing reasoning length requires coordinated, distributed modifications to internal computation rather than sparse interventions on a small set of components. Together, these results suggest that reasoning efficiency is governed

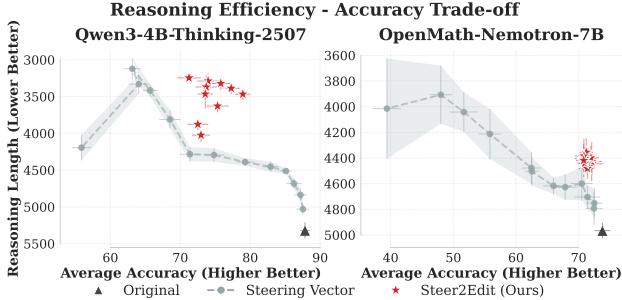


Figure 6. Accuracy–efficiency trade-off on **Qwen3-4B-Thinking-2507** and **OpenMath-Nemotron-7B**. The y-axis measures reasoning length (lower is better). **STEER2EDIT** achieves a more favorable accuracy–efficiency trade-off than activation steering.

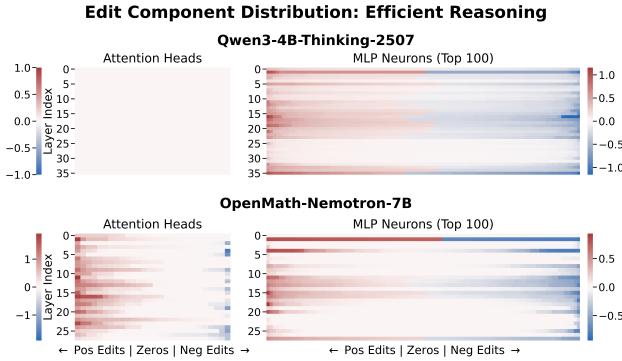


Figure 7. Signed **STEER2EDIT** edit coefficients λ for reasoning efficiency control. Positive values reinforce components associated with shorter reasoning traces, while negative values suppress components that promote longer chains of thought.

by broad MLP-based computation patterns, in sharp contrast to the sparse, attention-dominated circuits underlying safety and truthfulness.

4.3. Additional Experiments.

For completeness, we report (i) *per-dataset trade-off curves*, (ii) *design-choice ablations over (u, k, λ)* , (iii) a *component-wise budget sensitivity analysis* that isolates the effects of ρ_{attn} and ρ_{mlp} at fixed α , and (iv) comparisons with *training-based baselines* (full fine-tuning and rank-1 LoRA) in Appendix D, Appendix E, Appendix F, and Appendix G. .

5. Related Works

Steering and controlling LLM behavior. A growing body of work studies behavioral control in LLMs via interventions on internal representations. Representation engineering methods (Zou et al., 2023a; Turner et al., 2023; Ardi et al., 2024; Yan et al., 2025; Li et al., 2025) extract semantic directions from contrastive examples and apply inference-time activation interventions to modulate attributes such as safety or reasoning behavior. These approaches are training-free, but rely on global, inference-

time modifications. In parallel, Concept Bottleneck Models (Sun et al., 2025b; 2024; Bidusa & Markovitch, 2025; Tan et al., 2023; Ludan et al., 2023) introduce explicit concept variables and architectural constraints to enable structured, human-interpretable control. Together, these lines of work demonstrate that LLM behavior can be influenced through manipulation of internal representations.

In contrast, **STEER2EDIT** translates steering signals into *component-level* edits that operate at the level of individual components. By redistributing behavioral influence across attention heads and MLP neurons, **STEER2EDIT** enables more favorable attribute–utility trade-offs, while providing fine-grained interpretability and preserving the standard model architecture.

LLM weight editing. Another line of work focuses on modifying model parameters to induce persistent behavioral changes without full retraining. Representative approaches include meta-editors such as MEND (Mitchell et al., 2021) and KnowledgeEditor (De Cao et al., 2021), mechanistic editors such as ROME (Meng et al., 2022) and MEMIT (Meng et al., 2023), neuron-level interventions (Dai et al., 2022), and semi-parametric methods such as SERAC (Mitchell et al., 2022). More recent work applies targeted weight edits to specific components for behavior control, including ThinkEdit (Sun et al., 2025c) for mitigating overly short reasoning traces and DRefA (Chu et al., 2025) for safety. These methods are largely empirical and do not provide a unified framework for allocating and justifying edits across components.

STEER2EDIT complements this literature by providing a general, theoretically grounded, and training-free framework that systematically converts steering directions into component-wise weight updates.

6. Conclusion

We introduced **STEER2EDIT**, a principled framework that translates steering signals into component-level weight edits via a closed-form solution. By shifting behavioral control from inference-time activation intervention to parameter updates, **STEER2EDIT** achieves more favorable behavior–utility trade-offs while preserving the standard model architecture. Beyond empirical gains, the method offers fine-grained interpretability, revealing how safety, truthfulness, and reasoning efficiency are distributed across attention heads and MLP components. These results show that steering vectors can serve as effective diagnostic signals for systematic weight editing, providing a practical and theoretically grounded alternative to activation-level control.

Broader Impact

This paper presents work whose goal is to advance the field of machine learning by enabling efficient, interpretable, and training-free model editing at the level of individual components. In positive applications, STEER2EDIT may help practitioners correct or reduce undesirable behaviors (e.g., unsafe responses or hallucinations) and better understand which internal components support a given behavior, which can improve transparency and facilitate auditing. At the same time, weight-editing methods are inherently dual-use: the same capability can be applied to remove safeguards, amplify biases, or otherwise manipulate model behavior for harmful purposes, and edited models may be redistributed without clear provenance. To mitigate these risks, we emphasize that edits should be evaluated across diverse safety and capability tests, that releases should include clear documentation of intended use and limitations, and that responsible access controls may be appropriate for edits that materially alter safety-critical behaviors. Overall, we believe the primary societal consequence of this work is enabling more controllable and inspectable models, with corresponding responsibility to prevent and detect misuse.

References

Arditi, A., Obeso, O., Syed, A., Paleka, D., Panickssery, N., Gurnee, W., and Nanda, N. Refusal in language models is mediated by a single direction. *NeurIPS*, 2024.

Bidusa, O. R. and Markovitch, S. Concept layers: Enhancing interpretability and intervenability via llm conceptualization. *arXiv preprint arXiv:2502.13632*, 2025.

Chen, M. Evaluating large language models trained on code. *arXiv preprint arXiv:2107.03374*, 2021.

Chu, K. L., Sun, C.-E., and Weng, T.-W. How to make llms safer? detecting and editing key heads in llms. In *NeurIPS Lock-LLM Workshop*, 2025.

Dai, D., Dong, L., Hao, Y., Sui, Z., Chang, B., and Wei, F. Knowledge neurons in pretrained transformers. In *ACL*, 2022.

De Cao, N., Aziz, W., and Titov, I. Editing factual knowledge in language models. *arXiv preprint arXiv:2104.08164*, 2021.

Guo, D., Yang, D., Zhang, H., Song, J., Zhang, R., Xu, R., Zhu, Q., Ma, S., Wang, P., Bi, X., et al. Deepseek-r1: Incentivizing reasoning capability in llms via reinforcement learning. *arXiv preprint arXiv:2501.12948*, 2025.

Li, Y., Sun, C.-E., and Weng, T.-W. Effective skill unlearning through intervention and abstention. *NAACL*, 2025.

Ludan, J. M., Lyu, Q., Yang, Y., Dugan, L., Yatskar, M., and Callison-Burch, C. Interpretable-by-design text classification with iteratively generated concept bottleneck. *arXiv preprint*, 2023.

Meng, K., Bau, D., Andonian, A., and Belinkov, Y. Locating and editing factual associations in gpt. *NeurIPS*, 2022.

Meng, K., Sharma, A. S., Andonian, A., Belinkov, Y., and Bau, D. Mass-editing memory in a transformer. *ICLR*, 2023.

Mitchell, E., Lin, C., Bosselut, A., Finn, C., and Manning, C. D. Fast model editing at scale. *arXiv preprint arXiv:2110.11309*, 2021.

Mitchell, E., Lin, C., Bosselut, A., Manning, C. D., and Finn, C. Memory-based model editing at scale. In *ICML*, 2022.

Olsson, C., Elhage, N., Nanda, N., Joseph, N., DasSarma, N., Henighan, T., Mann, B., Askell, A., Bai, Y., Chen, A., et al. In-context learning and induction heads. *arXiv preprint arXiv:2209.11895*, 2022.

Sun, C.-E., Oikarinen, T., and Weng, T.-W. Crafting large language models for enhanced interpretability. In *ICML 2024 Workshop on Mechanistic Interpretability*, 2024.

Sun, C.-E., Liu, X., Yang, W., Weng, T.-W., Cheng, H., San, A., Galley, M., and Gao, J. Iterative self-tuning llms for enhanced jailbreaking capabilities. In *NAACL*, 2025a.

Sun, C.-E., Oikarinen, T., Ustun, B., and Weng, T.-W. Concept bottleneck large language models. In *ICLR*, 2025b.

Sun, C.-E., Yan, G., and Weng, T.-W. Thinkedit: Interpretable weight editing to mitigate overly short thinking in reasoning models. *EMNLP*, 2025c.

Tan, Z., Cheng, L., Wang, S., Bo, Y., Li, J., and Liu, H. Interpreting pretrained language models via concept bottlenecks. *arXiv preprint*, 2023.

Turner, A. M., Thiergart, L., Leech, G., Udell, D., Vazquez, J. J., Mini, U., and MacDiarmid, M. Steering language models with activation engineering. *arXiv preprint arXiv:2308.10248*, 2023.

Yan, G., Sun, C.-E., et al. Reflctrl: Controlling llm reflection via representation engineering. *arXiv preprint arXiv:2512.13979*, 2025.

Yao, S., Zhao, J., Yu, D., Du, N., Shafran, I., Narasimhan, K., and Cao, Y. React: Synergizing reasoning and acting in language models, 2023.

Zhao, Y., Zhang, W., Xie, Y., Goyal, A., Kawaguchi, K., and Shieh, M. Understanding and enhancing safety mechanisms of llms via safety-specific neuron. In *ICLR*, 2025.

Zhou, Z., Yu, H., Zhang, X., Xu, R., Huang, F., Wang, K., Liu, Y., Fang, J., and Li, Y. On the role of attention heads in large language model safety. *arXiv preprint arXiv:2410.13708*, 2024.

Zou, A., Phan, L., Chen, S., Campbell, J., Guo, P., Ren, R., Pan, A., Yin, X., Mazeika, M., Dombrowski, A., Goel, S., Li, N., Byun, M. J., Wang, Z., Mallen, A., Basart, S., Koyejo, S., Song, D., Fredrikson, M., Kolter, J. Z., and Hendrycks, D. Representation engineering: A top-down approach to AI transparency. *CoRR*, 2023a.

Zou, A., Wang, Z., Carlini, N., Nasr, M., Kolter, J. Z., and Fredrikson, M. Universal and transferable adversarial attacks on aligned language models. *arXiv preprint arXiv:2307.15043*, 2023b.

Table of Contents

A Proofs for STEER2EDIT	11
A.1 Proof of Theorem 3.1	11
A.2 Proof of Theorem 3.2	12
A.3 Proof of Theorem 3.3	13
B Steering Vector Construction	15
C Hyperparameter Search Procedure	16
D Per-Dataset Trade-off Analysis	17
D.1 Safety Alignment: Attack- and Dataset-Specific Trade-offs	17
D.2 Truthfulness: Dataset-Specific Utility Trade-offs	17
D.3 Efficient Reasoning: Dataset-Level Accuracy-Length Trade-offs	17
E Ablation Study of STEER2EDIT: Empirical Justification of Formal Design Choices	20
E.1 Unified Performance Summary	20
E.2 Detailed Per-Setting Results	21
F Component-Wise Budget Sensitivity Analysis	23
F.1 Safety Alignment: Sensitivity to Attention Budget	23
F.2 Truthfulness: Sensitivity Dominated by Attention	23
F.3 Efficient Reasoning: Sensitivity Dominated by MLP Budget	24
G Additional Baselines: Comparing STEER2EDIT with Training-Based Methods	25
G.1 Safety Alignment	25
G.2 Truthfulness	25
G.3 Efficient Reasoning	26

A. Proofs for STEER2EDIT

A.1. Proof of Theorem 3.1

Theorem A.1 (Output-space direction under semantic invariance). *Let $v_i \neq 0$, and let $\Delta W_i = \lambda_i u_i k_i^\top$ be a rank-1 edit with $\Delta W_i \neq 0$. If for all h_i and all $z \perp v_i$ we have*

$$z^\top \Delta W_i h_i = 0,$$

then the output-space direction u_i must be collinear with v_i , i.e.,

$$u_i \in \text{span}\{v_i\}.$$

Proof. Substituting $\Delta W_i = \lambda_i u_i k_i^\top$,

$$z^\top \Delta W_i h_i = \lambda_i (z^\top u_i) (k_i^\top h_i).$$

Because $\Delta W_i \neq 0$, we have $k_i \neq 0$, and therefore there exists some h_i such that $k_i^\top h_i \neq 0$. For the expression above to vanish for all such h_i , we must have $z^\top u_i = 0$ for every vector z orthogonal to v_i . The only vectors satisfying this condition are those proportional to v_i . Hence $u_i \in \text{span}\{v_i\}$. \square

A.2. Proof of Theorem 3.2

Theorem A.2 (Input-space direction matching semantic alignment variation). *Fix a component W_i and set $u_i = \hat{v}_i$. Assume $W_i^\top v_i \neq 0$ and $\text{Var}(s_i(h_i)) > 0$, where $s_i(h_i) := v_i^\top W_i h_i$. Consider choosing an input-direction $k_i \neq 0$ so that the induced semantic alignment shift $\Delta s_i(h_i) := v_i^\top \Delta W_i h_i$ exhibits maximal co-variation with the component's intrinsic semantic alignment score $s_i(h_i)$. Formally, consider the objective*

$$\max_{k_i \neq 0} \left| \text{Pearson}(\Delta s_i(h_i), s_i(h_i)) \right|.$$

Then there exists a maximizer k_i that is collinear with $W_i^\top v_i$, i.e.,

$$k_i \in \text{span}\{W_i^\top v_i\}.$$

Proof. Recall that $s_i(h_i) := v_i^\top W_i h_i$ and $\Delta s_i(h_i) := v_i^\top \Delta W_i h_i$. Using $\Delta W_i = \lambda_i \hat{v}_i k_i^\top$, we have

$$\Delta s_i(h_i) = \lambda_i \|v_i\|_2 (k_i^\top h_i).$$

Pearson correlation is invariant to additive shifts in either argument, so it is unchanged if we center the inputs. Let $\tilde{h}_i = h_i - \mu_i$ with $\mu_i = \mathbb{E}[h_i]$, and define

$$\widetilde{\Delta s}_i(h_i) := \Delta s_i(\tilde{h}_i), \quad \tilde{s}_i(h_i) := s_i(\tilde{h}_i).$$

Denote the covariance matrix by

$$\Sigma_i := \mathbb{E}[\tilde{h}_i \tilde{h}_i^\top].$$

Step 1: covariance. We have

$$\widetilde{\Delta s}_i(h_i) = \lambda_i \|v_i\|_2 (k_i^\top \tilde{h}_i), \quad \tilde{s}_i(h_i) = v_i^\top W_i \tilde{h}_i.$$

Hence

$$\begin{aligned} \text{Cov}(\widetilde{\Delta s}_i, \tilde{s}_i) &= \mathbb{E}[\widetilde{\Delta s}_i(h_i) \tilde{s}_i(h_i)] \\ &= \lambda_i \|v_i\|_2 \mathbb{E}[(k_i^\top \tilde{h}_i)(v_i^\top W_i \tilde{h}_i)] \\ &= \lambda_i \|v_i\|_2 k_i^\top \mathbb{E}[\tilde{h}_i \tilde{h}_i^\top] W_i^\top v_i \\ &= \lambda_i \|v_i\|_2 k_i^\top \Sigma_i W_i^\top v_i. \end{aligned}$$

Step 2: variances. Similarly,

$$\text{Var}(\widetilde{\Delta s}_i) = \mathbb{E}[\widetilde{\Delta s}_i(h_i)^2] = \lambda_i^2 \|v_i\|_2^2 k_i^\top \Sigma_i k_i,$$

and

$$\text{Var}(\tilde{s}_i) = \mathbb{E}[\tilde{s}_i(h_i)^2] = v_i^\top W_i \Sigma_i W_i^\top v_i.$$

Step 3: Pearson correlation. The Pearson correlation between the induced and intrinsic semantic signals is

$$\text{Pearson}(\widetilde{\Delta s}_i, \tilde{s}_i) = \frac{\text{Cov}(\widetilde{\Delta s}_i, \tilde{s}_i)}{\sqrt{\text{Var}(\widetilde{\Delta s}_i)} \sqrt{\text{Var}(\tilde{s}_i)}}.$$

Substituting the expressions above gives

$$\begin{aligned} \text{Pearson}(\widetilde{\Delta s}_i, \tilde{s}_i) &= \frac{\lambda_i \|v_i\|_2 k_i^\top \Sigma_i W_i^\top v_i}{\sqrt{\lambda_i^2 \|v_i\|_2^2 k_i^\top \Sigma_i k_i} \sqrt{v_i^\top W_i \Sigma_i W_i^\top v_i}} \\ &= \frac{\text{sign}(\lambda_i) k_i^\top \Sigma_i W_i^\top v_i}{\sqrt{k_i^\top \Sigma_i k_i} \sqrt{v_i^\top W_i \Sigma_i W_i^\top v_i}}. \end{aligned}$$

The denominator's second factor is independent of k_i , and $\text{sign}(\lambda_i)$ is irrelevant when maximizing absolute correlation. Thus maximizing $|\text{Pearson}(\widetilde{\Delta s}_i, \tilde{s}_i)|$ over $k_i \neq 0$ reduces to maximizing

$$\frac{|k_i^\top \Sigma_i W_i^\top v_i|}{\sqrt{k_i^\top \Sigma_i k_i}}. \tag{1}$$

Step 4: Cauchy–Schwarz in the Σ_i -inner product. Define the Σ_i -inner product:

$$\langle a, b \rangle_{\Sigma_i} := a^\top \Sigma_i b, \quad \|a\|_{\Sigma_i} := \sqrt{a^\top \Sigma_i a}.$$

Then (1) becomes

$$\frac{|\langle k_i, W_i^\top v_i \rangle_{\Sigma_i}|}{\|k_i\|_{\Sigma_i}}.$$

By Cauchy–Schwarz,

$$|\langle k_i, W_i^\top v_i \rangle_{\Sigma_i}| \leq \|k_i\|_{\Sigma_i} \|W_i^\top v_i\|_{\Sigma_i}.$$

Equality is attained by choosing $k_i \propto W_i^\top v_i$. Therefore, there exists an optimizer $k_i \in \text{span}\{W_i^\top v_i\}$. \square

A.3. Proof of Theorem 3.3

Theorem A.3 (Edit magnitude allocation under regularization). *For each component W_i , let $g_i = \cos(v_i, W_i \mu_i)$ denote the component importance score, with the convention that $g_i := 0$ if $W_i \mu_i = 0$.*

Consider the problem of assigning edit magnitudes $\{\lambda_i\}$ to maximize total signed alignment as measured by the component importance scores $\{g_i\}$, while controlling both the sparsity and overall strength of the edit. Formally, let $\mathbf{g} = (g_1, \dots, g_n)$ and consider

$$\max_{\lambda \in \mathbb{R}^n} \mathbf{g}^\top \lambda - \rho \left(\alpha \|\lambda\|_1 + \frac{1-\alpha}{2} \|\lambda\|_2^2 \right), \quad \rho > 0, \alpha \in [0, 1).$$

The unique edit magnitude assigned to component i is

$$\lambda_i^* = \text{sign}(g_i) \frac{\max(|g_i| - \rho\alpha, 0)}{\rho(1-\alpha)}.$$

Proof. For component i , define the one-dimensional objective

$$J(\lambda_i) = g_i \lambda_i - \rho \left(\alpha |\lambda_i| + \frac{1-\alpha}{2} \lambda_i^2 \right).$$

A scalar value λ_i^* maximizes J iff

$$0 \in \partial J(\lambda_i^*),$$

where the subgradient is needed only at $\lambda_i = 0$ due to the nondifferentiability of $|\lambda_i|$. We analyze the three regions $\lambda_i > 0$, $\lambda_i < 0$, and $\lambda_i = 0$.

Case 1: $\lambda_i > 0$. Here $|\lambda_i| = \lambda_i$, so

$$J(\lambda_i) = g_i \lambda_i - \rho \left(\alpha \lambda_i + \frac{1-\alpha}{2} \lambda_i^2 \right).$$

Differentiating gives

$$\frac{dJ}{d\lambda_i} = g_i - \rho\alpha - \rho(1-\alpha)\lambda_i.$$

Setting this to zero yields

$$\lambda_i = \frac{g_i - \rho\alpha}{\rho(1-\alpha)},$$

which is valid only when the positivity assumption holds, i.e. $g_i > \rho\alpha$.

Case 2: $\lambda_i < 0$. Here $|\lambda_i| = -\lambda_i$, so

$$J(\lambda_i) = g_i \lambda_i + \rho\alpha \lambda_i - \frac{\rho(1-\alpha)}{2} \lambda_i^2.$$

Differentiating,

$$\frac{dJ}{d\lambda_i} = g_i + \rho\alpha - \rho(1-\alpha)\lambda_i.$$

Setting this to zero gives

$$\lambda_i = \frac{g_i + \rho\alpha}{\rho(1 - \alpha)},$$

which is valid only when $g_i < -\rho\alpha$.

Case 3: $\lambda_i = 0$. At zero, the subdifferential of $|\lambda_i|$ is $\partial|\lambda_i| = [-1, 1]$. The optimality condition

$$0 \in g_i - \rho\alpha s, \quad s \in [-1, 1],$$

is feasible iff $|g_i| \leq \rho\alpha$. Thus $\lambda_i^* = 0$ whenever the alignment is too small to exceed the threshold.

Combining all three cases gives the soft-threshold rule

$$\lambda_i^* = \begin{cases} \frac{g_i - \rho\alpha}{\rho(1 - \alpha)}, & g_i > \rho\alpha, \\ 0, & |g_i| \leq \rho\alpha, \\ \frac{g_i + \rho\alpha}{\rho(1 - \alpha)}, & g_i < -\rho\alpha. \end{cases}$$

Equivalently,

$$\lambda_i^* = \text{sign}(g_i) \frac{\max(|g_i| - \rho\alpha, 0)}{\rho(1 - \alpha)}.$$

□

B. Steering Vector Construction

This appendix describes the procedures used to construct steering vectors for each behavioral control setting. Across all experiments, steering vectors are computed using the same general mean-difference formulation and are shared verbatim between activation steering and **STEER2EDIT**. Only the definition of positive and negative response sets differs by task.

General formulation. For a given model, layer ℓ , and block type (attention or MLP), we collect the block output activations for a set of responses. Let \mathcal{P} and \mathcal{N} denote the positive and negative response sets associated with a target behavior. For each response, we average the block outputs over all response tokens. The steering vector is then defined as the difference between the mean activations:

$$v_\ell = \mathbb{E}_{x \in \mathcal{P}}[h_\ell(x)] - \mathbb{E}_{x \in \mathcal{N}}[h_\ell(x)].$$

This procedure is applied independently to the attention and MLP blocks at each layer, yielding $\{v_\ell^{\text{attn}}, v_\ell^{\text{mlp}}\}_{\ell=1}^L$.

Safety alignment. For safety alignment, we construct steering vectors using the ADVBench dataset. The positive set \mathcal{P} consists of refusal responses to harmful prompts, while the negative set \mathcal{N} consists of standard helpful responses to benign questions sampled from Alpaca dataset. Steering vectors are computed from model-generated responses.

Truthfulness. For truthfulness promotion, we use the TruthfulQA dataset. We split the dataset into a probing set and an evaluation set. Model responses on the probing set are labeled as *truthful* or *hallucinated* using **QwQ-32B** as an external judge. The positive set \mathcal{P} consists of truthful responses, and the negative set \mathcal{N} consists of hallucinated responses.

efficient Reasoning. For reasoning efficiency control, we use the GSM8K training set. We measure the length of each model-generated reasoning trace and select the top 5% shortest and top 5% longest responses. The positive set \mathcal{P} consists of short reasoning traces, and the negative set \mathcal{N} consists of long reasoning traces. The resulting steering vectors capture directions associated with shorter internal reasoning processes.

All steering vectors are computed once per model and per behavioral control setting. During evaluation, the same steering vectors are applied across all test sets, reflecting a practical deployment scenario in which vectors are not optimized for any specific evaluation benchmark.

C. Hyperparameter Search Procedure

STEER2EDIT introduces three scalar hyperparameters: the attention editing budget ρ_{attn} , the MLP editing budget ρ_{mlp} , and the Elastic-Net sparsity parameter α . These parameters control how edit magnitude is allocated across model components and how sparsely edits are distributed. Hyperparameters are selected using a lightweight two-stage grid search on held-out data.

Step 1: Coarse grid search

We first perform a coarse-grained grid search over a shared range that is identical across all models and behavioral control settings:

$$\rho_{\text{attn}} \in \{0.1, 0.3, 0.5, 0.7, 0.9\}, \quad \rho_{\text{mlp}} \in \{0.1, 0.3, 0.5, 0.7, 0.9\}, \quad \alpha \in \{0.1, 0.3, 0.5, 0.7, 0.9\}.$$

The goal of this step is to identify the approximate operating regime (e.g., attention-dominated versus MLP-dominated edits, sparse versus distributed allocation), rather than to finely optimize performance.

Each configuration is first subjected to a lightweight sanity check using 20 short, simple prompts. If the edited model exhibits degenerate behavior (e.g., repetitive output, failure to respond, or nonsensical generations), the configuration is immediately discarded. This allows unstable settings to be filtered at negligible cost.

Step 2: Refined grid search

Based on the results of the coarse search, we define a refined but still small grid for each (model, setting) pair. The refined grids narrow the range and reduce the step size around regions that exhibit meaningful improvements in the target attribute while preserving normal model behavior.

In several settings, we observe that edits to either attention or MLP components have negligible impact on the target behavior. In these cases, the corresponding component is not edited at all, and no budget is assigned to that component during the refined search.

Final search ranges

Table 1 summarizes the refined hyperparameter ranges used in each behavioral control setting. All reported results in the main paper, including the best-performing configuration and the top-10 configurations shown in trade-off plots, are selected exclusively from these ranges.

Setting / Model	ρ_{attn}	ρ_{mlp}	α
<i>Safety Alignment</i>			
LLaMA-2-7B-Chat	[0.16, 0.24] (step = 0.02)	[0.35, 0.55] (step = 0.05)	[0.70, 0.90] (step = 0.05)
Mistral-7B-Instruct-v0.2	[0.42, 0.50] (step = 0.02)	[0.40, 0.60] (step = 0.05)	[0.65, 0.85] (step = 0.05)
<i>Truthfulness</i>			
Gemma-2-2B-IT	[0.30, 0.50] (step = 0.05)	<i>negligible</i>	[0.75, 0.95] (step = 0.05)
LLaMA-3-8B-Instruct	[0.10, 0.14] (step = 0.01)	[0.30, 0.50] (step = 0.05)	[0.30, 0.70] (step = 0.10)
<i>Efficient Reasoning</i>			
Qwen3-4B-Thinking-2507	<i>negligible</i>	[0.65, 0.80] (step = 0.05)	[0.05, 0.20] (step = 0.05)
OpenMath-Nemotron-7B	[0.20, 0.30] (step = 0.05)	[0.80, 0.90] (step = 0.05)	[0.10, 0.20] (step = 0.05)

Table 1. Refined hyperparameter search ranges for **STEER2EDIT**. “Negligible” indicates that edits to the corresponding component were found to have insufficient effect during coarse search and are therefore not applied in the refined search.

Efficiency and reporting

Hyperparameter search for **STEER2EDIT** is computationally lightweight. Each configuration requires only a single closed-form application of rank-1 weight edits, followed by evaluation on a held-out small validation set. In practice, the full two-stage search completes within minutes per model, and does not involve gradient-based optimization. All evaluations are performed on held-out data that is disjoint from steering vector extraction. No additional tuning is performed on test sets.

D. Per-Dataset Trade-off Analysis

In Section 4.2, we summarize each behavioral control setting using aggregated downstream utility metrics to provide a concise comparison across methods. In this appendix, we present *per-dataset trade-off curves* that expose finer-grained behavior across individual evaluation benchmarks. These results demonstrate that the superior attribute–utility trade-offs achieved by **STEER2EDIT** are consistent across datasets.

D.1. Safety Alignment: Attack- and Dataset-Specific Trade-offs

For safety alignment, we evaluate two jailbreak attack methods (GCG and ADV-LLM) and three downstream utility benchmarks (CommonsenseQA, Code-MMLU, and GSM8K), resulting in six distinct safety–utility trade-off settings per model.

Figure 8 reports refusal rate versus downstream utility separately for each attack–dataset pair on **LLaMA-2-7B-Chat** and **Mistral-7B-Instruct-v0.2**. Across most settings, **STEER2EDIT** identifies configurations that achieve higher refusal rates at comparable or higher utility than inference-time activation steering.

For **LLaMA-2-7B-Chat**, **STEER2EDIT** consistently dominates the steering baseline under both GCG and ADV-LLM attacks across all downstream datasets. For **Mistral-7B-Instruct-v0.2**, performance depends on the attack strength: under the weaker GCG attack, **STEER2EDIT** is occasionally slightly worse than activation steering at comparable utility, whereas under the substantially stronger ADV-LLM attack, **STEER2EDIT** achieves markedly higher refusal rates while preserving downstream accuracy.

Notably, the advantage of **STEER2EDIT** becomes more pronounced as the attack strength increases. While activation steering requires aggressive intervention that sharply degrades utility under ADV-LLM, weight-level edits derived by **STEER2EDIT** maintain stable benign-task performance while substantially improving robustness to strong jailbreaks.

D.2. Truthfulness: Dataset-Specific Utility Trade-offs

For truthfulness promotion, downstream utility is evaluated independently on CommonsenseQA, Code-MMLU, and GSM8K. Figure 9 shows truthfulness versus utility accuracy for **Gemma-2-2B-IT** and **LLaMA-3-8B-Instruct** on each benchmark.

Across all datasets, activation steering exhibits a pronounced trade-off in which increasing truthfulness rapidly degrades task performance. In contrast, **STEER2EDIT** consistently attains higher truthfulness at substantially higher utility, with edited configurations occupying regions of the trade-off space that is unattainable by steering alone.

These per-dataset results demonstrate that the truthfulness gains of **STEER2EDIT** generalizes across reasoning-heavy and knowledge-oriented benchmarks.

D.3. Efficient Reasoning: Dataset-Level Accuracy–Length Trade-offs

For efficient reasoning control, we report dataset-specific trade-offs between answer accuracy and reasoning length on GSM8K, MATH-500, GPQA, and Code-MMLU for **Qwen3-4B-Thinking-2507** and **OpenMath-Nemotron-7B**.

Figure 10 shows that activation steering reduces reasoning length primarily by sacrificing accuracy, with the severity of this trade-off varying substantially across datasets. In contrast, **STEER2EDIT** consistently identifies configurations that shorten reasoning traces while preserving accuracy, including on challenging benchmarks such as GPQA and MATH-500.

Notably, these improvements generalize beyond GSM8K, despite the steering direction being extracted from GSM8K, indicating that **STEER2EDIT** captures a transferable mechanism for reasoning efficiency control.

Safety Alignment - Utility Trade-off

LLaMA-2-7B-Chat

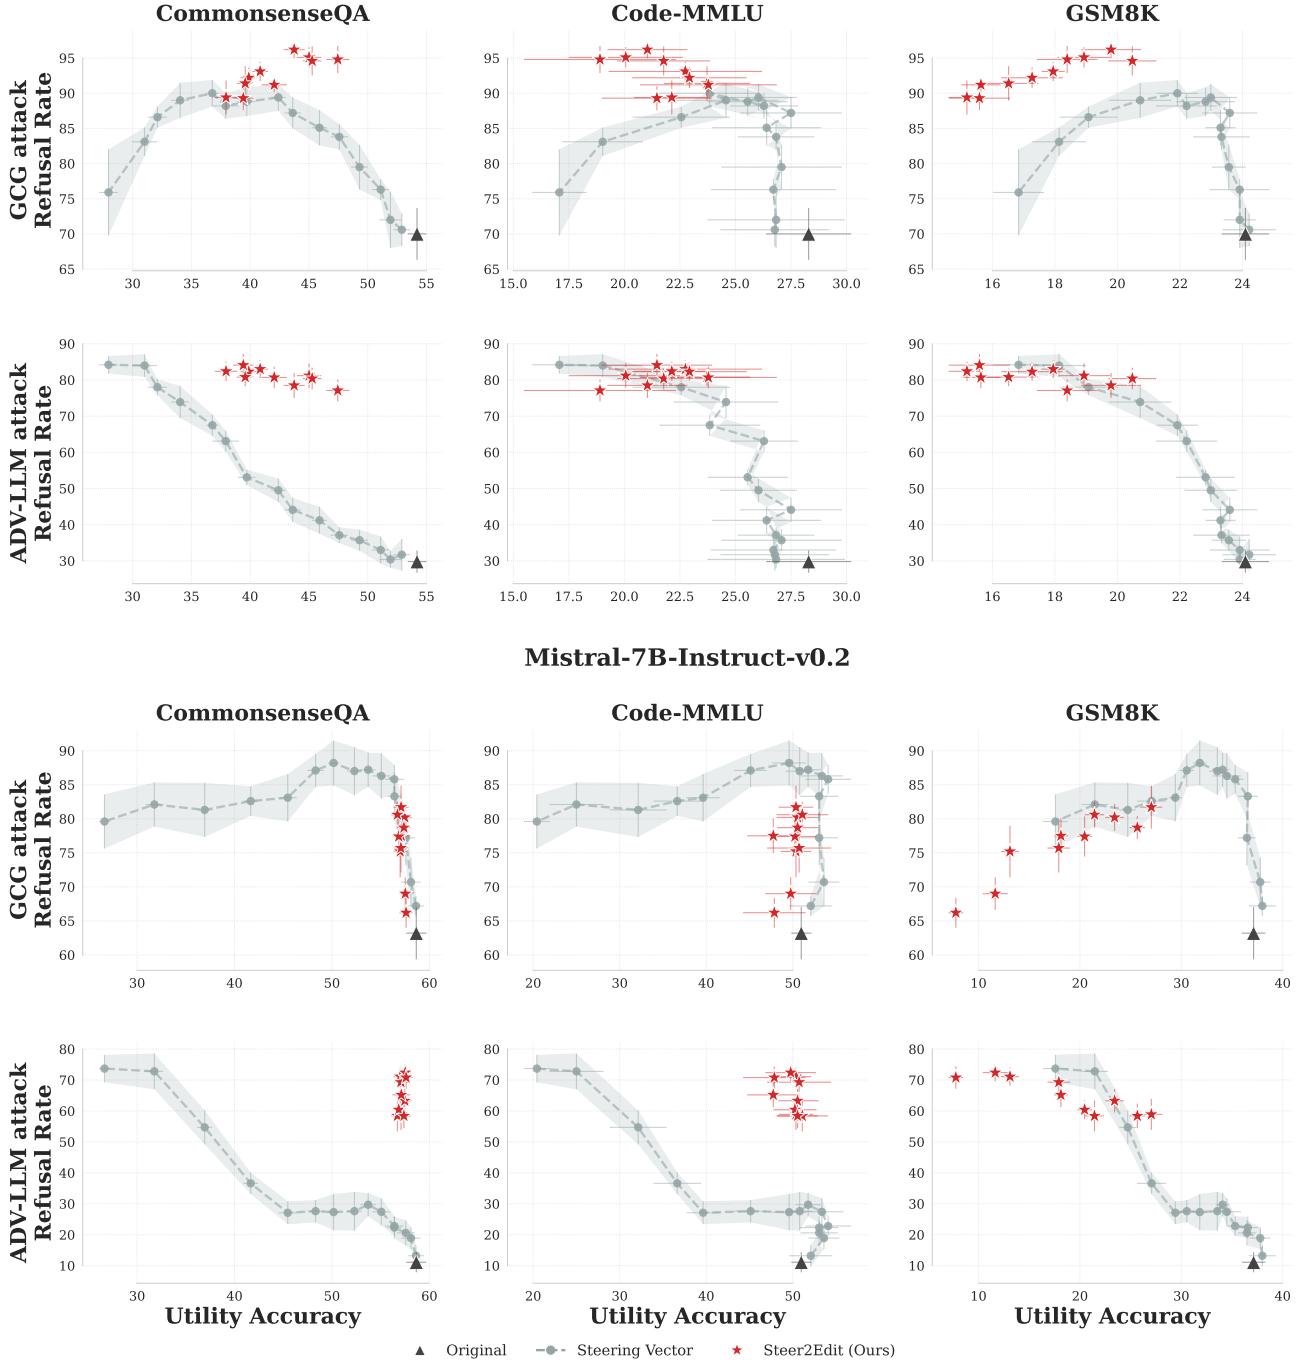


Figure 8. Per-dataset safety–utility trade-offs under GCG and ADV-LLM attacks. Each column corresponds to a downstream utility dataset (CommonsenseQA, Code-MMLU, GSM8K), and each row corresponds to an attack method. **STEER2EDIT** consistently achieves higher refusal rates at comparable or higher utility than activation steering across all settings.

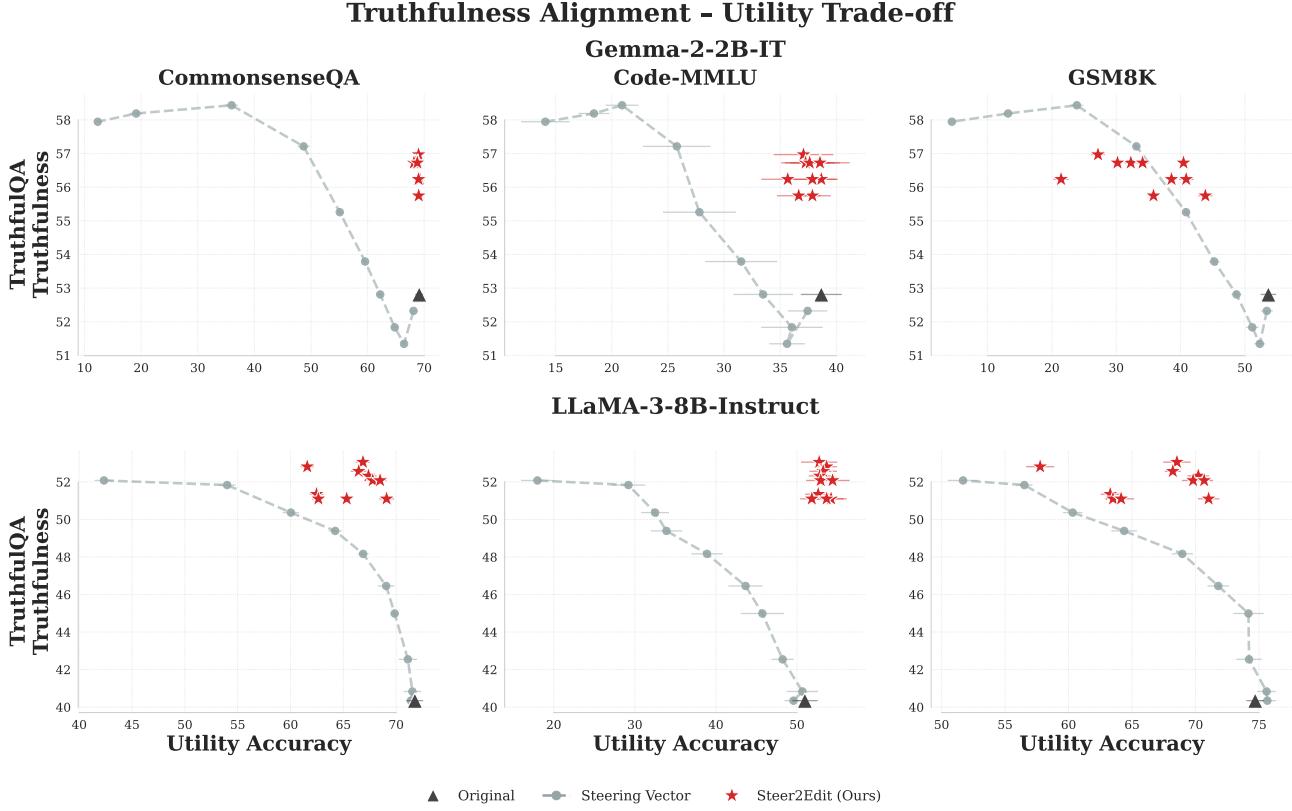


Figure 9. Per-dataset truthfulness–utility trade-offs on CommonsenseQA, Code-MMLU, and GSM8K for Gemma-2-2B-IT and LLaMA-3-8B-Instruct. **STEER2EDIT** improves truthfulness while preserving higher downstream utility across all datasets.

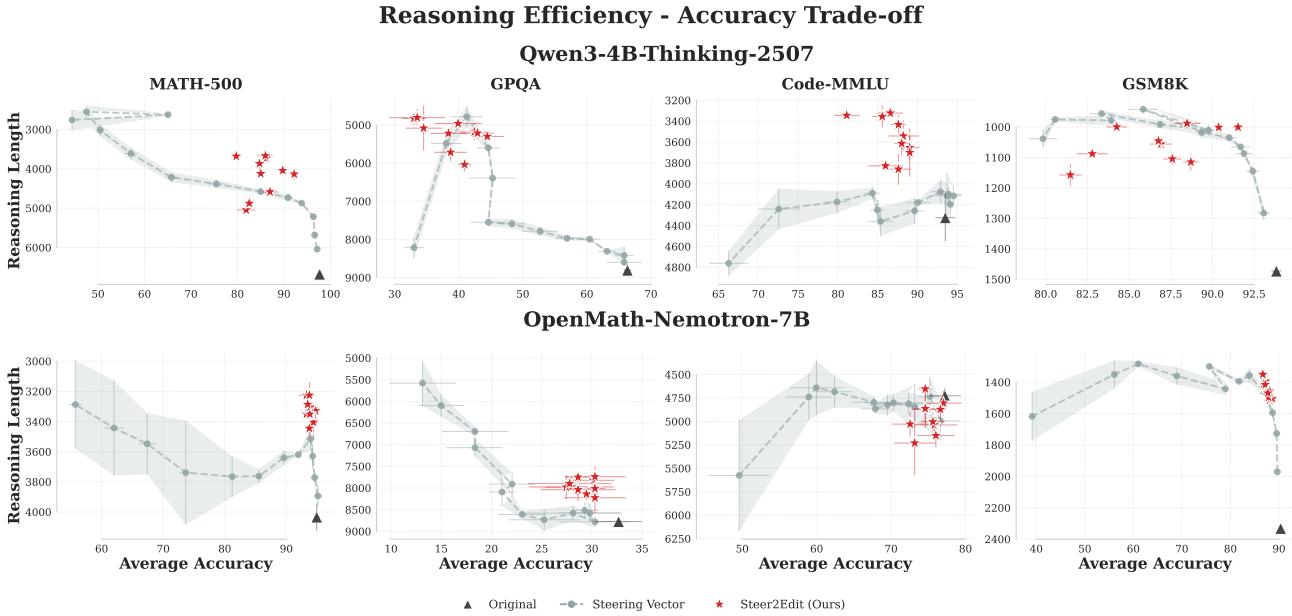


Figure 10. Per-dataset accuracy–reasoning-length trade-offs on GSM8K, MATH-500, GPQA, and Code-MMLU for Qwen3-4B-Thinking-2507 and OpenMath-Nemotron-7B. **STEER2EDIT** reduces reasoning length while maintaining accuracy across all datasets.

E. Ablation Study of STEER2EDIT: Empirical Justification of Formal Design Choices

In this section, we validate the design choices of **STEER2EDIT** through a systematic ablation study. We decompose the rank-1 weight update, $\Delta W_i = \lambda_i u_i k_i^\top$, and independently modify its three core components: the **Input Direction** (k), the **Importance Score** (g), and the **Edit Magnitude** (λ), while holding the others constant.

This study aims to determine whether the performance of **STEER2EDIT** arises from its specific geometric and statistical formulations—such as using cosine similarity for scoring or Elastic-Net for sparsity—rather than generic weight modifications. Table 2 defines the five ablation variants tested and the specific hypothesis each one investigates.

Table 2. Ablation Definitions. We categorize variants by the component they modify: Input Direction (k), Importance Score (g), or Magnitude (λ).

Ablation Category	Variant	Definition and Hypothesis Tested
Input k	k_{mean}	Sets $k_i \leftarrow \mu_i$, where $\mu_i = \mathbb{E}[h_i]$. Tests if data statistics alone suffice without directional sensitivity.
	k_{svd}	Sets $k_i \leftarrow \mathbf{v}_1$, where \mathbf{v}_1 is the top right singular vector of W_i . Tests if intrinsic weight directions suffice.
Score g	g_{dot}	Sets $g_i \leftarrow \hat{v}_i^\top (W_i \mu_i)$ (unnormalized dot product), removing the normalization by $\ W_i \mu_i\ _2$. Tests the effect of component-output normalization.
Magnitude λ	ℓ_0	Selects top- K components (hard threshold) matching Elastic-Net sparsity. Tests if sparsity alone explains gains.
	ℓ_2	Uses Elastic-Net with $\alpha = 0$ (Ridge regularization), resulting in dense edits. Tests if dense edits can preserve utility.

E.1. Unified Performance Summary

To facilitate a high-level comparison, Table 3 aggregates the results across all three behavioral settings: Safety Alignment, Truthfulness, and Efficient Reasoning.

We normalize the attribute scores to a common $[0, 1]$ scale. For **Safety** and **Truthfulness**, we use the raw percentage divided by 100. For **Efficient Reasoning**, where lower length is better, we define the efficiency score as $\min(1, L_{\text{STEER2EDIT}}/L_{\text{Ablation}})$. **STEER2EDIT** consistently achieves the highest combined **Attribute** \times **Utility** score, demonstrating that precise input alignment, normalized scoring, and sparse regularization are all critical for optimal performance.

Table 3. Unified Performance Summary. Results are averaged across 6 models (2 per setting). **STEER2EDIT** achieves the best global trade-off between targeting the attribute and maintaining downstream utility.

Category	Variant	Normalized Attribute Score \uparrow			Normalized Utility Score \uparrow			Overall Average \uparrow		Attr \times Util
		Safety	Truth	Effic.	Safety	Truth	Effic.	Attribute	Utility	
Full Method	Steer2Edit	0.807	0.550	1.000	0.341	0.536	0.755	0.786	0.544	0.427
Input k	k_{mean}	0.827	0.531	0.926	0.279	0.350	0.589	0.761	0.406	0.309
	k_{svd}	0.536	0.474	0.757	0.400	0.572	0.805	0.589	0.592	0.349
Score g	g_{dot}	0.102	0.509	0.628	0.047	0.232	0.003	0.413	0.094	0.039
Magnitude λ	ℓ_0	0.506	0.546	0.767	0.321	0.423	0.596	0.606	0.447	0.271
	ℓ_2	0.000	0.580	0.472	0.032	0.226	0.394	0.350	0.217	0.076

Key Insights: Normalization and Sparsity. Two critical design principles emerge from these aggregated results. First, **Score Normalization is paramount for stability.** The catastrophic failure of the unnormalized g_{dot} variant (Overall Score: 0.039) reveals that raw activation magnitudes vary drastically across model layers. Without the cosine-similarity normalization used in **STEER2EDIT**, the editing process becomes dominated by high-norm layers, destabilizing the model regardless of the target attribute. Second, **Sparsity is essential for utility preservation.** The dense ℓ_2 baseline struggles to maintain downstream performance (Utility: 0.217), confirming that precise, sparse interventions are required to disentangle

specific behaviors without overwriting the model’s general knowledge base.

E.2. Detailed Per-Setting Results

The following subsections analyze the impact of each component across our three behavioral settings. We find that while different tasks exhibit unique sensitivities, the failures of the ablation baselines consistently point to the necessity of **STEER2EDIT**’s three-pillared approach: precise directional alignment, normalized scoring, and sparse editing.

Safety Alignment: Selective vs. Uniform Triggering. Table 4 highlights the risks of uniformly triggered edits. The k_{mean} baseline, which activates edits based on the global mean activation $\mathbb{E}[x]$, achieves high safety (e.g., 92.20% refusal on ADV-LLM) but substantially reduces utility (21.75% vs. 28.00% for **STEER2EDIT**). This indicates that activating edits in a largely input-agnostic manner leads to significant degradation on benign downstream tasks. Conversely, the dense ℓ_2 baseline results in near-total model failure (0.05% safety), suggesting that safety-related behavior is mediated by localized components and is disrupted by dense parameter modifications. **STEER2EDIT** avoids these failure modes by combining input-selective activation with sparse, component-level edits.

Table 4. Safety Alignment Ablations (Detailed). Metrics are Average Safety (Refusal Rate) and Average Utility (Accuracy).

Model	Variant	$S \times U$	Avg Safety	Avg Util	Safety (Refusal Rate)		Utility (Accuracy)		
					GCG	ADV-LLM	CommonsenseQA	Code-MMLU	GSM8K
Llama2-7B-Chat	Steer2Edit	24.68	88.15	28.00	95.10	81.20	45.01	20.06	18.93
	k_{mean}	19.12	87.90	21.75	83.60	92.20	30.14	19.57	15.53
	k_{svd}	22.82	66.35	34.39	90.80	41.90	50.77	30.49	21.90
	g_{dot}	0.69	15.35	4.49	23.30	7.40	6.45	5.37	1.65
	ℓ_0	7.66	41.25	18.58	36.40	46.10	27.26	25.43	3.05
	ℓ_2	0.00	0.05	6.33	0.10	0.00	3.60	14.88	0.52
Mistral-7B-Instruct	Steer2Edit	29.38	73.15	40.16	75.20	71.10	57.00	50.37	13.10
	k_{mean}	26.33	77.45	34.00	71.60	83.30	54.24	45.55	2.22
	k_{svd}	18.59	40.85	45.51	68.80	12.90	53.61	51.95	30.98
	g_{dot}	0.25	5.05	4.96	0.90	9.20	5.43	9.45	0.01
	ℓ_0	27.29	59.85	45.60	77.20	42.50	55.96	48.35	32.48
	ℓ_2	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00

Truthfulness: The Critical Role of Sparsity. Table 5 provides strong evidence for the necessity of sparse editing (ℓ_0 regularization). The dense ℓ_2 variant, which modifies all parameters in the target block, causes catastrophic utility collapse (0.00% on Llama3-8B), demonstrating that dense edits severely disrupt the model’s general capabilities. Additionally, we observe an informative trade-off with k_{svd} : while it preserves high utility (65.11% on Llama3), it fails to significantly improve truthfulness (39.61%). This suggests that aligning edits with the model’s intrinsic dominant activation patterns is insufficient to induce a shift toward truthful behavior. Only **STEER2EDIT** balances these objectives: it uses k to selectively activate edits on semantically relevant inputs, while ℓ_0 sparsity restricts the intervention to a small set of behaviorally relevant neurons.

Table 5. **Truthfulness Ablations (Detailed).** Metrics are Average Truthfulness (TruthfulQA Accuracy) and Average Utility (Accuracy).

Model	Method	T × U	Avg Truth	Avg Util	Attribute TruthfulQA	Utility (Accuracy)		
						CommonsenseQA	Code-MMLU	GSM8K
Gemma2 2B-IT	Steer2Edit	25.30	56.97	44.41	56.97	69.00	37.07	27.16
	k_{mean}	24.93	55.75	44.71	55.75	68.59	38.29	27.24
	k_{svd}	27.26	55.26	49.33	55.26	68.12	39.39	40.48
	g_{dot}	1.16	58.68	1.98	58.68	2.83	2.20	0.90
	ℓ_0	15.33	57.95	26.46	57.95	58.29	16.04	5.06
	ℓ_2	26.15	57.95	45.12	57.95	63.62	36.95	34.78
Llama3 8B-Instruct	Steer2Edit	33.27	53.06	62.71	53.06	66.86	52.74	68.54
	k_{mean}	12.74	50.37	25.29	50.37	39.67	25.61	10.60
	k_{svd}	25.79	39.61	65.11	39.61	71.30	51.71	72.32
	g_{dot}	19.11	43.03	44.42	43.03	26.94	41.77	64.56
	ℓ_0	29.87	51.34	58.18	51.34	64.14	47.13	63.26
	ℓ_2	0.00	57.95	0.00	57.95	0.00	0.00	0.00

Efficient Reasoning: Stability via Normalization. Table 6 highlights the critical role of score normalization for stable reasoning control. The g_{dot} variant, which uses the raw dot product, fails pathologically: it reduces the reasoning length to just 21 tokens (Qwen3-4B) while collapsing utility to near zero (0.13%). This behavior arises because activation norms vary substantially across layers; without the cosine normalization used in **STEER2EDIT**, edit magnitudes become dominated by high-norm layers, leading to severe disruption of model behavior. Conversely, the k_{svd} baseline increases reasoning length (5351 vs. 3467 for **STEER2EDIT**), indicating that activating edits along the model’s intrinsic dominant activation patterns is insufficient for improving efficiency. Together, these results show that effective reasoning-length control requires both normalized importance scoring and task-specific edit activation, as implemented in **STEER2EDIT**.

Table 6. **Efficient Reasoning Ablations (Detailed).** Metrics are Reasoning Length (Lower is Better) and Utility (Higher is Better). The **U / L** metric represents efficiency (Utility/Length × 100).

Model	Variant	U / L	Avg Len	Avg Util	Reasoning Length (Lower is Better)				Utility (Higher is Better)			
					MATH-500	GPQA	Code-MMLU	GSM8K	MATH-500	GPQA	Code-MMLU	GSM8K
Qwen3 4B-Thinking	Steer2Edit	2.28	3467	78.95	4136	5299	3433	1000	92.2	44.4	87.6	91.5
	k_{mean}	1.56	3266	51.05	3188	2728	6129	1018	53.2	10.3	78.1	62.7
	k_{svd}	1.64	5351	87.99	6389	8780	4622	1613	96.9	67.5	93.9	93.6
	g_{dot}	0.62	21	0.13	23	33	10	18	0.3	0.0	0.0	0.2
	ℓ_0	0.73	6506	47.61	8646	5202	4698	7477	44.7	36.0	65.7	44.1
	ℓ_2	0.08	22476	17.63	32223	18987	6420	32274	3.0	15.2	50.4	2.0
OpenMath- Nemotron-7B	Steer2Edit	1.62	4445	72.12	3405	7821	5038	1515	94.4	30.1	75.8	88.1
	k_{mean}	1.28	5216	66.69	3838	8814	6522	1691	92.5	25.8	66.7	81.9
	k_{svd}	1.42	5131	73.01	4188	8880	4985	2472	93.9	32.8	75.6	89.7
	g_{dot}	0.00	17338	0.52	18380	16628	17163	17182	0.9	0.0	0.6	0.6
	ℓ_0	1.64	4356	71.60	3422	7971	4503	1528	94.5	29.6	73.4	88.9
	ℓ_2	1.09	5633	61.23	2911	12686	5160	1773	89.5	17.7	66.1	71.7

F. Component-Wise Budget Sensitivity Analysis

In the main paper, we show that the best-performing **STEER2EDIT** configurations exhibit a consistent component-level structure: safety and truthfulness control rely on sparse edits to attention heads, whereas reasoning efficiency is primarily governed by distributed MLP neuron edits.

To further validate that this structural separation is intrinsic to the underlying mechanisms, we conduct a controlled *component-wise budget sensitivity analysis*. In this study, we fix the sparsity parameter α and vary the regularization budget of one component class at a time, while disabling edits to the other class by sending its budget to infinity ($\rho \rightarrow \infty$). This isolates how changes in attention and MLP budgets individually influence the attribute–utility trade-off.

F.1. Safety Alignment: Sensitivity to Attention Budget

Figure 11 illustrates how the safety–utility trade-off responds to changes in the attention and MLP budgets when considered in isolation. Increasing the attention budget ρ_{attn} produces substantial gains in refusal rate at moderate utility cost, closely matching the best joint configurations reported in the main paper. In contrast, varying the MLP budget ρ_{mlp} leads to markedly weaker safety improvements and often degrades utility more rapidly.

This asymmetric sensitivity indicates that safety alignment is primarily mediated by a small number of attention heads. The result is consistent with the sparsity patterns observed in Figure 3, where non-zero edit coefficients are concentrated in late-layer attention components, with minimal contribution from MLP neurons.

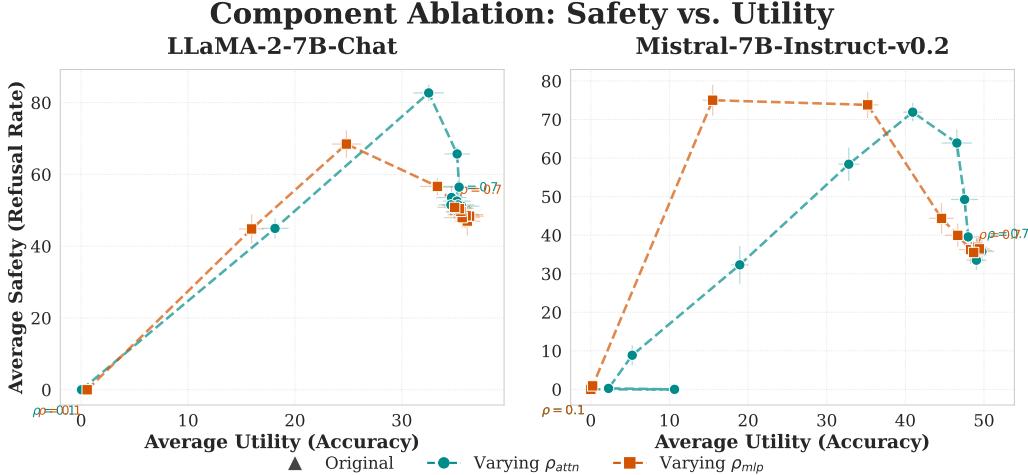


Figure 11. Component-wise budget sensitivity for safety alignment. We fix the sparsity parameter α and vary the attention regularization budget ρ_{attn} while disabling MLP edits by taking $\rho_{\text{mlp}} \rightarrow \infty$, and vice versa. Improvements in refusal rate are primarily driven by changes in the attention budget, whereas varying the MLP budget yields limited safety gains and earlier utility degradation.

F.2. Truthfulness: Sensitivity Dominated by Attention

Figure 12 reports the component-wise budget sensitivity for truthfulness promotion. Across both evaluated models, increasing the attention budget consistently yields larger improvements in truthful preference accuracy than increasing the MLP budget under the same sparsity constraint. MLP-only edits fail to recover the trade-off frontier achieved by attention edits.

These findings align with the component-level edit distributions shown in Figure 5, which reveal that truthfulness control is achieved through sparse, localized attention interventions. Notably, several models exhibit predominantly negative edit coefficients, suggesting that suppressing hallucination-promoting attention heads is more effective than broadly modifying MLP computation.

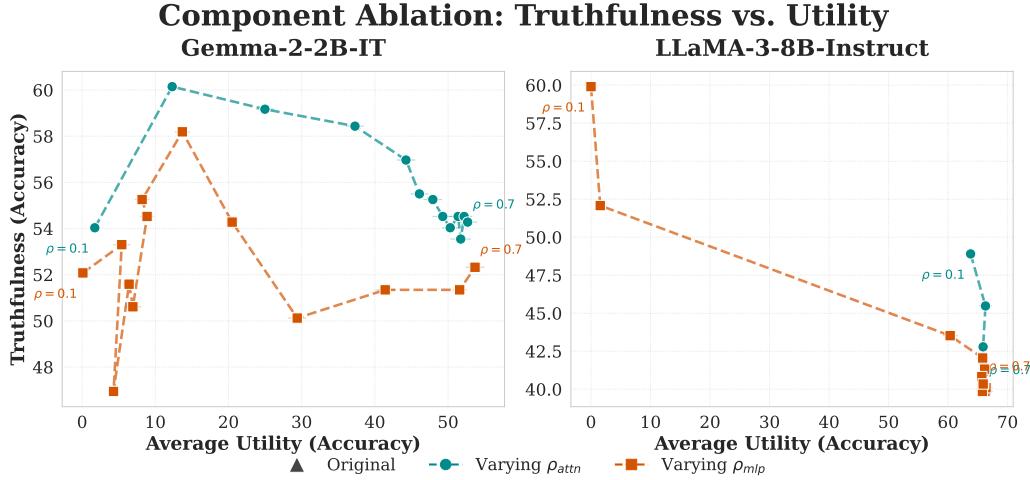


Figure 12. **Component-wise budget sensitivity for truthfulness control.** Truthfulness improvements are strongly sensitive to the attention budget ρ_{attn} , while varying the MLP budget ρ_{mlp} in isolation results in substantially smaller gains at comparable downstream utility.

F.3. Efficient Reasoning: Sensitivity Dominated by MLP Budget

Figure 13 shows that reasoning efficiency exhibits a qualitatively different sensitivity pattern. Increasing the MLP budget ρ_{mlp} leads to a smooth and substantial reduction in reasoning length while preserving accuracy. In contrast, varying the attention budget ρ_{attn} in isolation produces only marginal efficiency improvements, even at large budgets.

This behavior mirrors the dense MLP edit patterns observed in Figure 7 and confirms that efficient reasoning control requires coordinated, distributed modifications to MLP neurons across layers. Unlike safety and truthfulness, which are governed by localized attention-based circuits, reasoning efficiency emerges from broad MLP-based computation.

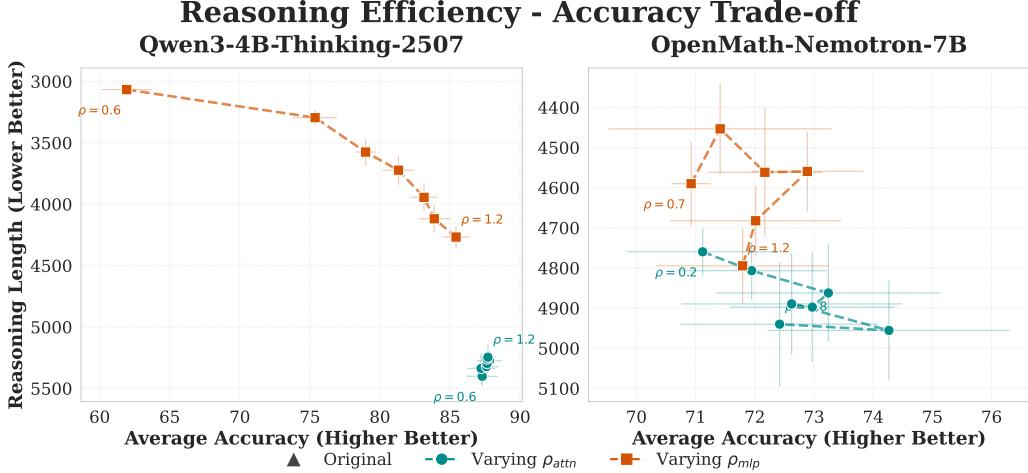


Figure 13. **Component-wise budget sensitivity for reasoning efficiency.** Reductions in reasoning length are strongly influenced by the MLP budget ρ_{mlp} , while varying the attention budget ρ_{attn} yields only minor efficiency gains. This indicates that efficient reasoning is governed by distributed MLP computation rather than sparse attention circuits.

Summary. Across all behavioral control settings, this component-wise budget sensitivity analysis establishes a clear correspondence between the component class whose budget most strongly influences the trade-off frontier and the components receiving non-zero edits in the best-performing **STEER2EDIT** configurations. Safety and truthfulness are attention-dominated, whereas reasoning efficiency is MLP-dominated, providing further evidence that **STEER2EDIT** uncovers genuine, setting-dependent circuit structure rather than artifacts of hyperparameter tuning.

G. Additional Baselines: Comparing STEER2EDIT with Training-Based Methods

We compare STEER2EDIT against training-based adaptation methods that directly optimize model parameters toward a target behavior. Specifically, we consider **full-parameter fine-tuning** and **rank-1 LoRA fine-tuning** as additional baselines.

Setup and comparison protocol. For each control setting (safety, truthfulness, and efficient reasoning), we fine-tune models on the *same probing dataset used to extract steering vectors*. Training uses the *positive set* (i.e., examples that exhibit the target attribute), so the model is explicitly optimized to imitate the desired behavior. We evaluate the resulting models using the same attribute and downstream utility metrics as in the main experiments and report trade-off curves alongside activation steering and STEER2EDIT. The **full** fine-tuning baseline updates all model parameters. The **rank-1** baseline applies LoRA adapters with rank $r = 1$, inserted into the standard attention projections (q_{proj} , k_{proj} , v_{proj} , o_{proj}) and MLP projections ($\text{gate}_{\text{proj}}$, up_{proj} , $\text{down}_{\text{proj}}$), while keeping the backbone weights frozen. All baselines are trained with standard supervised objectives and comparable training budgets.

G.1. Safety Alignment

Figure 14 shows the safety–utility trade-off for models fine-tuned on the safety probing positive set. Full fine-tuning increases refusal rates but often does so by globally shifting the model’s response distribution, leading to over-refusal on benign queries and sharp drops in downstream utility, particularly for Mistral. Rank-1 LoRA produces minimal changes, indicating limited capacity to induce reliable safety behavior under this supervision.

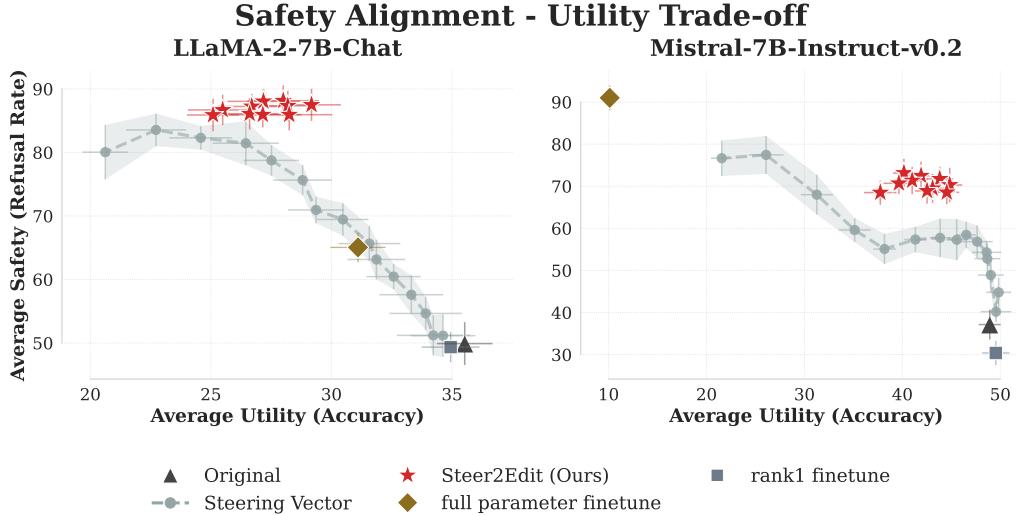


Figure 14. Safety–utility trade-off with training-based baselines. Full fine-tuning improves refusal rates but frequently collapses utility due to over-refusal, especially in the low-data regime. Rank-1 LoRA has negligible effect. These trends indicate that training-based optimization on small positive sets induces coarse, global behavioral shifts rather than selective safety control.

G.2. Truthfulness

Figure 15 reports the truthfulness–utility trade-off for models fine-tuned on the truthfulness probing positive set. Full fine-tuning yields modest improvements in TruthfulQA accuracy, but these gains are typically accompanied by noticeable degradation in downstream utility, suggesting broad shifts in the model’s answer distribution rather than selective promotion of truthfulness. Rank-1 LoRA again exhibits little effect.

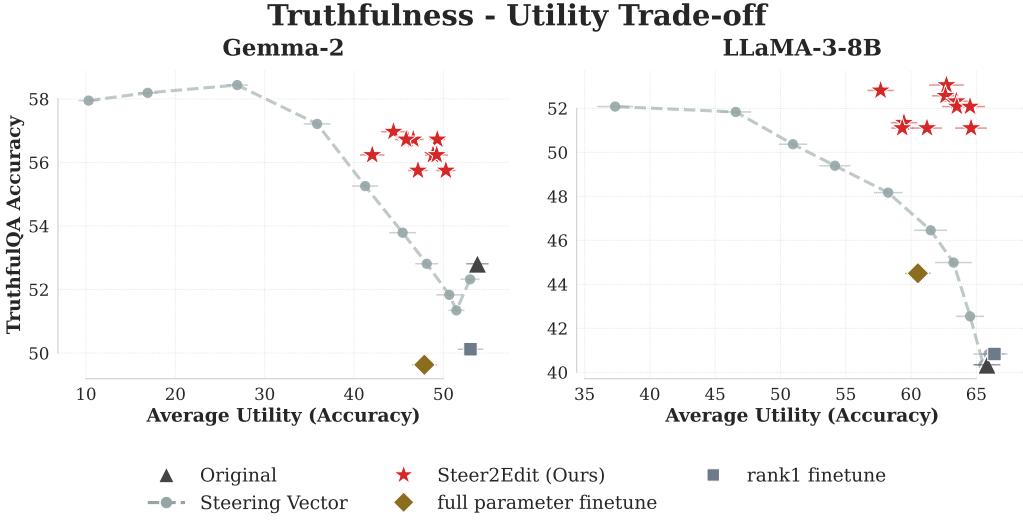


Figure 15. **Truthfulness–utility trade-off with training-based baselines.** Fine-tuning on the positive set improves TruthfulQA performance but often incurs utility loss, reflecting over-regularization of response behavior. Rank-1 LoRA provides insufficient adaptation capacity to meaningfully alter truthfulness.

G.3. Efficient Reasoning

Figure 16 presents the reasoning efficiency–accuracy trade-off for models fine-tuned on the efficient-reasoning probing positive set. Full fine-tuning can encourage shorter generations, but the resulting reductions in reasoning length are generally comparable to activation steering and do not consistently surpass it. Rank-1 LoRA again produces minimal changes, indicating limited capacity to meaningfully influence reasoning behavior under this supervision.

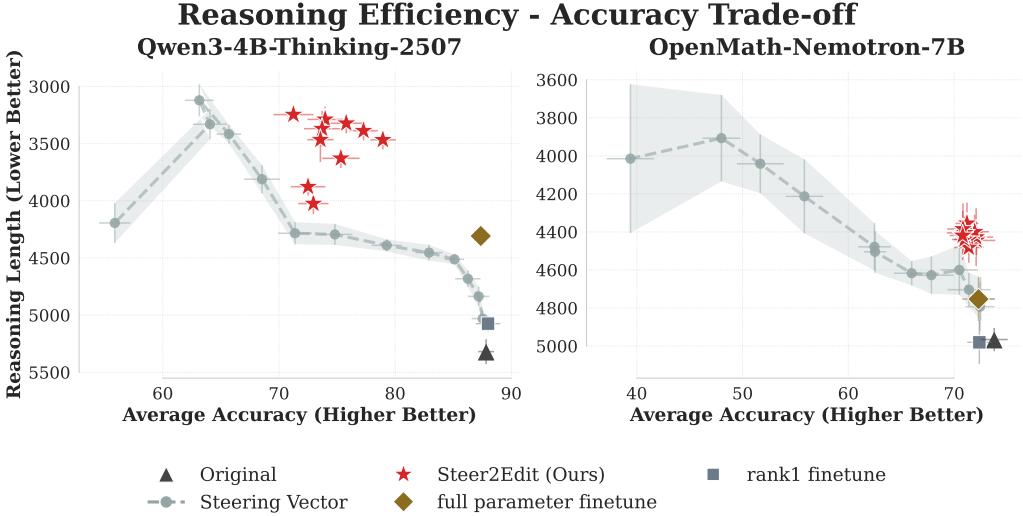


Figure 16. **Reasoning efficiency–accuracy trade-off with training-based baselines.** Full fine-tuning reduces reasoning length but largely matches the trade-off achieved by activation steering, without clear advantages. Rank-1 LoRA has little effect.

Summary. Across all three control settings, fine-tuning on the probing dataset can move models toward the target attribute, but typically does so by inducing broad distributional shifts that trace a trade-off curve similar to activation steering. These limitations are most evident in the low-data regime considered here, where the probing sets are intentionally small and narrowly targeted. Rank-1 LoRA consistently exhibits weak effects across all settings. In contrast, STEER2EDIT achieves more favorable trade-offs while remaining training-free and component-interpretable, highlighting the benefit of converting steering diagnostics into targeted weight edits rather than optimizing behavior through global parameter updates.

