

ScopeSAE: Model-Scope Feature Discovery with Interpretable Layer Selection

Qingwen Zeng^{1,*}, Zehao Fu^{1,*}, Shuyu Meng¹, Linghan Huang¹, Jiayi Zhang²,
Chenglin Wu³, Ling Chen⁴, Huaming Chen^{1,†}

¹The University of Sydney, ²The Hong Kong University of Science and Technology (Guangzhou),

³DeepWisdom, ⁴University of Technology Sydney

Abstract

Sparse autoencoders (SAEs) are a central tool in mechanistic interpretability. However, existing SAEs are primarily trained per layer. The modeling subspace is therefore fixed by layer identity, independent of which token-layer states actually drive each prediction. We argue that this constraint contributes to several limitations observed in layer-wise SAEs, including low feature utilization, high dictionary redundancy, and features that lack direct behavioral grounding. In this paper, we propose **ScopeSAE**, which selects the modeling subspace per token by attributing each prediction to its most influential token-layer state via normalized gradient-based attribution, and learns features over the resulting prediction-relevant subspace. Empirically, ScopeSAE yields an effect we term *reconstruction-better-than-original*. Written-back reconstructions of the SAE produce lower next-token cross-entropy than the original activations, an outcome that, to our knowledge, has not previously been reported for SAEs. Through interventional analyses and a KL fine-tuning counter-experiment, we show that this effect is attributable to ScopeSAE’s prediction-relevant subspace itself rather than to architectural changes. ScopeSAE further improves effective feature count, interpretability, utilization, and dictionary redundancy over existing layer-based baselines, suggesting that choosing the SAE modeling subspace by predictive relevance leads to more useful and behaviorally meaningful features. Open code is available here ¹.

1 Introduction

Sparse autoencoders (SAEs) have recently become one of the core tools in mechanistic interpretability [4, 13, 9]. Their motivation is that internal representations in large language models (LLMs) are typically high-dimensional and dense, with multiple semantic factors and behavioral mechanisms superposed within a shared representational space [32]. Because these factors overlap along shared directions, inspecting individual neurons or raw activation dimensions often fails to yield clear and stable semantic interpretations [7, 27]. By learning an overcomplete and sparse feature dictionary, SAEs decompose entangled representations into sparse features with stronger semantic coherence, thereby supporting analysis of internal representations and mechanisms underlying model behavior [31]. Therefore, the value of SAEs lies not merely in reconstructing activations, but in providing a more mechanism-level vocabulary for understanding the internal computation of LLMs [6, 22].

Despite their strong interpretability potential, most existing SAE methods rest on a largely implicit structural assumption: the modeling subspace is layer-local, and features are learned within that subspace [4, 13, 9, 25]. Concretely, SAE training data typically come from token activations

*Equal contribution.

†Correspondence to Huaming Chen <huaming.chen@sydney.edu.au>.

¹<https://github.com/Mechanical-Interpretability/ScopeSAE>

of a single fixed layer, so the learned dictionary is naturally tied to that layer’s representational space [26, 18]. This design choice is understandable, since SAEs were introduced to disentangle superposition within the activations of a particular layer, and a fixed-layer activation distribution provides a stable and tractable training domain [4, 7, 13]. However, single-layer modeling becomes inherently limiting when the goal is to explain a model’s final predictive behavior [19, 28, 17]. Features learned by a single-layer SAE are usually most interpretable within the layer on which they are trained, whereas next-token prediction in large language models is not determined by any one layer in isolation, but by representations that are progressively propagated, combined, and transformed across depth [16, 6, 10]. Moreover, because a single-layer subspace contains only a limited number of effectively decomposable factors, the relatively large latent dimensions used by SAEs to encourage sparse expressivity can in practice yield only a small number of effective features, along with many rarely activated dictionary elements, low feature utilization, and substantial redundancy [9, 25].

To move beyond the limitations of fixed single-layer subspaces, recent work has extended SAE modeling from one layer to multiple layers. One strategy is to train separate SAEs for different layers, thereby extracting feature structures from each layer independently [12, 18]. While this reveals layer-specific structure, it still does not explain how features interact across layers to determine the final output [19, 17]. More importantly, it does not capture the model’s effective computation within a unified behavior-relevant subspace. At the same time, training separate SAEs across layers is costly: training and storage overhead grow roughly linearly with depth, and the learned dictionaries typically lack a unified coordinate system, hindering cross-layer comparison and leading to fragmented, dispersed interpretations [19, 28]. Another new direction is to train a single SAE across multiple layers. RouteSAE uses a router to dynamically select, for each token, the layer for SAE encoding and decoding, thereby learning a shared dictionary over a multi-layer activation distribution [28]. While this increases the number of usable features, its objective is still defined at the level of representation reconstruction rather than directly at the level of the model’s final predictive behavior.

If SAEs are meant to extract cleaner and more interpretable features from entangled internal states, and if the behavior we seek to explain is next-token prediction, why not define the modeling subspace by predictive relevance rather than layer identity? Motivated by this idea, we propose ScopeSAE. Instead of predefining a fixed layer or a static collection of layer-wise subspaces as the training domain [4, 9, 12, 13, 18, 28], ScopeSAE constructs a training distribution over prediction-relevant token-layer states. Concretely, for each token, we measure the sensitivity of the next-token prediction objective to its state at each layer, and use normalized gradient-based attribution to identify the token-layer state most important for the current prediction. We then collect these selected latent states across tokens and train a single SAE over their union. In this way, ScopeSAE does not learn features from a fixed-layer activation distribution, but from a modeling subspace induced by token-layer states most relevant to prediction. As a result, the learned features are no longer primarily organized by layer identity, but by their relevance to the model’s current prediction.

Empirically, ScopeSAE exhibits a *reconstruction-better-than-original* effect not reported in prior SAE studies. Written-back reconstructions yield lower cross-entropy than the original activations, producing negative DeltaCE and Recovered above one. Through interventional analyses and a KL fine-tuning counter-experiment, we identify a specific causal pathway underlying this effect: the sparse bottleneck, operating on a prediction-relevant subspace, systematically removes components of the activation that are orthogonal to the prediction direction, and this pathway is a consequence of ScopeSAE’s subspace choice rather than of its architecture. ScopeSAE also outperforms existing SAE methods in effective feature count, feature interpretability, and feature utilization, while achieving lower dictionary redundancy. Taken together, these results suggest that shifting the SAE modeling subspace from a fixed layer-local representation space to a prediction-relevant subspace yields features with stronger semantic quality, higher utilization efficiency, and better behavioral explanatory power.

Our contributions are fourfold. **First**, we investigate a structural cause of the known limitations of existing SAEs. By fixing the modeling subspace to a single layer, standard training forces the same representation to explain prediction for which it is, in many cases, not the most informative source. **Second**, we propose ScopeSAE, which learns SAE features from a prediction-relevant modeling subspace by constructing a training distribution over token-layer states selected via normalized gradient-based attribution. **Third**, we report a *reconstruction-better-than-original* effect, in which written-back reconstruction yields lower next-token cross-entropy than original activation. Through interventional analyses and a KL fine-tuning counter-experiment, we show that this effect arises from ScopeSAE’s subspace choice, mediated by a specific causal pathway. **Fourth**, we benchmark




Figure 1: Overview of the ScopeSAE framework. For each token, normalized gradient-based attribution scores identify the most prediction-relevant token-layer state, over which a sparse autoencoder is trained. At inference, the reconstruction is written back only to the selected position, enabling behaviorally grounded feature discovery while remaining compatible with standard SAE evaluation and intervention pipelines.

ScopeSAE performance on a series of experiments, which show the state-of-the-art results across all standard mechanistic interpretability metrics, including effective feature count, interpretability, utilization, dictionary redundancy, and causal intervention over existing SAE methods.

2 Method

2.1 Fixed-Layer Modeling and Its Limitations

Fixed-Layer SAE Formulation A sparse autoencoder (SAE) learns an overcomplete feature dictionary over hidden states by encoding and reconstructing them. Given a hidden state $h \in \mathbb{R}^d$, the encoder maps h to sparse features $z \in \mathbb{R}^F$ via $z = \phi(W_e h + b_e)$, where $W_e \in \mathbb{R}^{F \times d}$, $b_e \in \mathbb{R}^F$, ϕ is a sparsity-inducing activation function, and $F \gg d$ is the dictionary size. The decoder reconstructs the original hidden state as $\hat{h} = W_d z + b_d$, where $W_d \in \mathbb{R}^{d \times F}$ and $b_d \in \mathbb{R}^d$. In the most common SAE paradigm [4, 13, 9, 25, 26], the training activations are drawn from a single predetermined layer l^* , forming the activation distribution $\mathcal{D}_{l^*} = \{h_{l^*}^{(t)} \mid t \in \mathcal{T}\}$, where $h_{l^*}^{(t)} \in \mathbb{R}^d$ denotes the hidden state of token t at layer l^* , and $\mathcal{T}$ is the set of training tokens. The SAE is then trained to minimize $\mathcal{L}_{\text{SAE}} = \frac{1}{|\mathcal{T}|} \sum_{t \in \mathcal{T}} \left[\|h_{l^*}^{(t)} - \hat{h}^{(t)}\|_2^2 + \lambda \|z^{(t)}\|_1 \right]$, where $\lambda \|z^{(t)}\|_1$ is an ℓ_1 penalty that encourages sparsity. The index l^* is a fixed hyperparameter chosen before training. This formulation makes explicit a structural assumption shared by virtually all existing SAE methods: the modeling subspace is entirely determined by l^* , and the learned feature dictionary is by construction tied to the representational space of that single layer.

Predictive computation across depth. The next-token prediction depends on the final hidden state $h_L^{(t)}$, which, through the residual stream, accumulates contributions from every layer: $h_L^{(t)} = h_0^{(t)} + \sum_{l=1}^L f_l(h_{l-1}^{(t)})$, where $h_0^{(t)}$ is the token embedding and f_l is the transformation at layer l . Each term $f_l(h_{l-1}^{(t)})$ represents a distinct computational contribution: early layers may capture lexical and syntactic patterns, while deeper layers encode higher-order reasoning and contextual integration [32]. Crucially, for different tokens, the information most relevant to prediction may reside at different depths; for example, a factual recall may rely heavily on mid-layer representations, while a syntactic completion may be largely determined by early layers. An SAE trained at a fixed layer l^*

models $h_{l^*}^{(t)}$ for all tokens regardless of whether layer l^* is where the prediction-relevant computation actually resides. The choice of which layer to model is therefore a hyperparameter-driven decision, systematically decoupled from where the prediction-relevant computation actually resides.

Limitations of Fixed-Layer Modeling This decoupling between layer selection and predictive relevance entails three concrete consequences. **Input-blind layer selection:** the decomposition above shows that the contribution $f_l(h_{l-1}^{(t)})$ of each layer to $h_L^{(t)}$ is distinct, and the relative importance of these contributions varies across tokens. Fixing a single l^* for all tokens ignores this variation, applying the same modeling scope to inputs whose prediction-relevant information may reside at different depths. **Reconstruction-prediction misalignment:** a fixed-layer SAE is trained to reconstruct activations at layer l^* , an objective with no direct dependence on the model’s predictive behavior. The learned features are therefore grounded in the representational structure of that layer rather than in the factors driving the model’s output, weakening their behavioral explanatory power. **Subspace underutilization:** a single-layer subspace captures only the representational structure present at one depth, limiting the number of semantically distinct features the dictionary can extract. Since the effective computation of an LLM is distributed across all L layers, this restricts the dictionary’s effective capacity, contributing to low feature utilization and high redundancy in practice.

These limitations motivate a shift in how the SAE modeling subspace is defined: rather than anchoring it to a fixed layer chosen independently of the input, we propose to construct it by selecting, for each token, the layer most relevant to the model’s current prediction.

2.2 Prediction-Relevant Subspace via Gradient-Based Attribution

To identify the most relevant layer for each token, we measure how sensitive the model’s predictive objective is to each layer’s hidden state: a layer whose state, if perturbed, would sharply change the prediction carries critical information, while one to which the prediction is insensitive contributes little. A first-order expansion of the objective provides exactly this per-layer sensitivity measure [29].

First-Order Attribution of Layer Relevance Let $\mathcal{J}^{(t)} = \log p_\theta(x_{t+1} \mid x_{\leq t})$ denote the log-probability that the model assigns at position t to its *prediction target*, here the next token x_{t+1} . Consider perturbing the hidden state at layer l by a small displacement $\delta \in \mathbb{R}^d$. A first-order Taylor expansion gives: $\mathcal{J}^{(t)}(h_l^{(t)} + \delta) \approx \mathcal{J}^{(t)}(h_l^{(t)}) + \langle \nabla_{h_l^{(t)}} \mathcal{J}^{(t)}, \delta \rangle$. Setting $\delta = -h_l^{(t)}$, i.e., ablating the layer’s state to zero, the first-order change in log-probability is: $\Delta \mathcal{J}_l^{(t)} \approx -\langle \nabla_{h_l^{(t)}} \mathcal{J}^{(t)}, h_l^{(t)} \rangle$.

This motivates defining the *attribution score* of layer l for token t as: $s_l^{(t)} = \langle \nabla_{h_l^{(t)}} \mathcal{J}^{(t)}, h_l^{(t)} \rangle$. A larger $s_l^{(t)}$ indicates that the hidden state at layer l contributes more to predicting the target: ablating it would cause a larger drop in $\mathcal{J}^{(t)}$, meaning the model becomes less confident in the target token. In practice, the gradient $\nabla_{h_l^{(t)}} \mathcal{J}^{(t)}$ is obtained from a single backward pass through the language model, and the score for each layer is computed as the inner product between the gradient and the hidden state, summed over the hidden dimension.

Defining the prediction-relevant subspace. Given the per-layer attribution scores $\{s_l^{(t)}\}_{l=1}^L$, we normalize them across layers to obtain a comparable selection signal. Concretely, we apply normalization across layers, $\tilde{s}_l^{(t)} = (s_l^{(t)} - \mu_s^{(t)}) / (\sigma_s^{(t)} + \epsilon)$, where $\mu_s^{(t)}$ and $\sigma_s^{(t)}$ are the mean and standard deviation of $\{s_l^{(t)}\}_{l=1}^L$, rescaling the scores to a common range so that selection is not biased by differences in raw gradient magnitude across depths. For each token t , we then select the layer whose hidden state is most critical to the current prediction: $l^*(t) = \arg \max_{l \in \{1, \dots, L\}} \tilde{s}_l^{(t)}$. Unlike the fixed l^* in existing methods, $l^*(t)$ is *input-dependent*: it varies across tokens according to where prediction-critical information resides in each forward pass. We collect the selected hidden states across all training tokens to form the *prediction-relevant training distribution*: $\mathcal{D}_{\text{scope}} = \{h_{l^*(t)}^{(t)} \mid t \in \mathcal{T}\}$. In contrast to $\mathcal{D}_{l^*}$, which is a static distribution tied to a single layer, $\mathcal{D}_{\text{scope}}$ is a dynamic distribution over token-layer states jointly induced by the model’s predictive objective. The modeling subspace is no longer defined by layer identity, but by predictive relevance. Note that the target participates only in the attribution score $s_l^{(t)}$ used to select the layer: the selected state $h_{l^*(t)}^{(t)}$ that the SAE encodes and later writes back (§2.3) comes from the model’s ordinary forward pass and depends on the past context $x_{\leq t}$ alone, not on the target. The selection is in this sense defined with respect to

a prediction target rather than the ground-truth specifically; in our main experiments we instantiate this target with the next token.

2.3 ScopeSAE: Training and Inference

Preprocessing. As $\mathcal{D}_{\text{scope}}$ aggregates hidden states from different layers and tokens, the activation scale can vary substantially across samples. To map them onto a unified scale so that a single SAE can train stably on this mixed distribution, we apply token-wise normalization to each selected hidden state: $\tilde{h}^{(t)} = (h_{l^*(t)}^{(t)} - \mu^{(t)})/(\sigma^{(t)} + \epsilon)$, where $\mu^{(t)}$ and $\sigma^{(t)}$ are the mean and standard deviation of $h_{l^*(t)}^{(t)}$ computed over the hidden dimension d . The statistics $(\mu^{(t)}, \sigma^{(t)})$ are cached for de-normalization at inference time.

Training objective. We adopt a TopK sparse autoencoder [9] as the concrete architecture. In the general SAE formulation of §2.1, the sparsity-inducing activation function ϕ is instantiated as a TopK operator that retains the K largest activations and zeros out the rest, yielding exact K active features without an auxiliary ℓ_1 penalty. The encoding and decoding are: $z^{(t)} = \text{TopK}(W_e \tilde{h}^{(t)} + b_e)$, $\hat{h}^{(t)} = W_d z^{(t)} + b_d$. The SAE is trained over $\mathcal{D}_{\text{scope}}$ with a normalized reconstruction loss: $\mathcal{L}_{\text{ScopeSAE}} = \frac{1}{|\mathcal{T}|} \sum_{t \in \mathcal{T}} \frac{\|\tilde{h}^{(t)} - \hat{h}^{(t)}\|_2^2}{\|\tilde{h}^{(t)}\|_2^2}$. Normalizing by $\|\tilde{h}^{(t)}\|_2^2$ makes the loss scale-invariant across token positions and layers, which is important since $\mathcal{D}_{\text{scope}}$ aggregates states from varying depths.

Inference and writeback. At inference time, ScopeSAE applies the same target-relative selection: for each token, the attribution scores $\{s_l^{(t)}\}$ are computed with respect to the prediction target, the most predictively relevant layer $l^*(t)$ is identified, and the corresponding hidden state is normalized, passed through the SAE, de-normalized, and written back to position $l^*(t)$ in the residual stream. The full pipeline is summarized in Algorithm 1. Computation cost analysis is in Appendix C.

Compatibility with standard SAE evaluation. Despite operating over a prediction-relevant subspace, ScopeSAE remains fully compatible with standard SAE evaluation protocols. At each token position, the SAE receives a single hidden state as input and produces a single reconstruction written back to one position in the residual stream, identical in interface to a fixed-layer SAE. This means that reconstruction-based metrics and downstream feature analysis methods can all be applied to ScopeSAE without modification.

3 Experiments

We first describe the experimental setup and then present a comprehensive evaluation of ScopeSAE. Following established SAE evaluation practices [9, 25, 13, 31, 12, 28], we assess performance along three axes: behavioral fidelity (§3.2), feature interpretability (§3.4), and causal intervention (§3.5). We further analyze the mechanism behind the observed reconstruction-better-than-original effect (§3.3), and supplement the core evaluation with additional analyses summarized in §3.6.

3.1 Setup

Inputs. All SAEs are trained on residual stream activations from Llama-3.2-1B-Instruct [11]. To ensure fair and uniform comparison across methods, we adopt the standard SAE evaluation protocol established in prior work. Single-layer baselines extract activations at Layer 11, corresponding to roughly $\frac{3}{4}$ of the model depth, following standard practice [9, 28]. For ScopeSAE and the multi-layer baselines (RouteSAE [28] and Crosscoder [19]), we adopt the candidate layer range of Layers 3–11 ($\frac{1}{4}$ to $\frac{3}{4}$ of model depth) used by [28], motivated by Lad et al. [16]’s finding that early transformer layers primarily handle detokenization while deeper layers contribute more directly to next-token prediction. Training data comes from OpenWebText2 [8] following [28]: roughly 100M tokens for training and 10M for evaluation, with a context length of 512. We apply activation normalization to stabilize training across all methods, following [9]. Hyperparameter details are in Appendix B.1.

Baselines. We compare against five baselines. ReLU SAE [13], Gated SAE [25], and TopK SAE [9] operate on the fixed-layer distribution. RouteSAE [28] and Crosscoder [19] both operate over the same candidate layers as ScopeSAE, the former using a learned router for dynamic layer selection and the latter employing per-layer encoder–decoder pairs.

3.2 Behavioral Fidelity and the Reconstruction-Better-Than-Original Effect

The test of an SAE’s behavioral grounding is whether its decoded activations, when substituted for the original, preserve the base model’s next-token prediction. We evaluate two standard behavioral fidelity metrics [9, 14]. **DeltaCE** measures the change in cross-entropy when SAE-decoded activations are written back: $\Delta\text{CE} = \text{CE}_{\text{recon}} - \text{CE}_{\text{original}}$; negative values indicate that writing back the decoded activation yields *lower* next-token cross-entropy than the original. **Recovered** quantifies the fraction of the model’s predictive capacity preserved after SAE encoding and decoding, relative to complete ablation: $\text{Recovered} = 1 - (\text{CE}_{\text{recon}} - \text{CE}_{\text{original}}) / (\text{CE}_{\text{zero}} - \text{CE}_{\text{original}})$, with values above one likewise indicating lower cross-entropy than the original.

As shown in Figure 2, ScopeSAE achieves the lowest DeltaCE across the full sparsity range we evaluate ($k \in [48, 72]$), and is the only method to enter the negative regime: DeltaCE first turns negative at $k=56$ and reaches -0.20 at $k=64$, while TopK and RouteSAE remain positive throughout (e.g., $+0.10$ and $+0.11$ at $k=64$). Correspondingly, ScopeSAE’s Recovered exceeds one from $k=56$ onward (1.03 at $k=64$), whereas all baselines plateau below one. This phenomenon, where writing back ScopeSAE’s sparse reconstruction yields lower cross-entropy than the model’s own unmodified activations, we term the *reconstruction-better-than-original* effect. To our knowledge, it has not been previously reported in SAE research.

The effect does not depend on the ground-truth target. Because the layer selection uses the next token as the prediction target, one might ask whether the lower cross-entropy is an artifact of access to the ground-truth. It is not. When the layer is instead selected from the model’s own top-1 predicted token, with no ground-truth used at any stage, DeltaCE remains negative (-0.206 vs. -0.197 for next-token selection) and the feature-level profile is unchanged. Routing the same selection rule over independently trained fixed-layer SAEs likewise does not reproduce the effect. Together these controls show that the reduction follows from the prediction-relevant subspace rather than from the selection target. We report all in Appendix O.

3.3 Mechanism Analysis

This section establishes the mechanism by which ScopeSAE’s prediction-relevant subspace produces the reconstruction-better-than-original effect. We proceed in three steps: first, we identify the causal pathway inside ScopeSAE through geometric, statistical, and interventional analyses (§3.3.1); second, we confirm via a counter-experiment that this pathway is a consequence of the subspace choice (§3.3.2); third, we show that the same mechanism accounts for ScopeSAE’s divergence on conventional reconstruction metrics (§3.3.3).

3.3.1 Causal Pathway Inside ScopeSAE

These three analyses target a specific property of the reconstruction residual $r_t = h_t - \hat{h}_t$: what the sparse bottleneck removes from each activation.

Geometric analysis: the residual lies orthogonal to the prediction gradient.

For each token t , we measure the alignment between the residual r_t and the prediction gradient $g_t = \nabla_{h_t} \mathcal{J}^{(t)}$ via absolute cosine similarity. As shown in Table 1, the residual’s alignment with the gradient ($|\cos| = 0.0103$) falls below the chance level for random vectors in $\mathbb{R}^d$ (0.0176, the expected $|\cos|$ for two random unit vectors in the model’s hidden dimension), while the reconstruction’s alignment (0.0420) exceeds that of the original activation (0.0329). This pattern holds across all candidate layers (L3–L11; see Appendix F). What ScopeSAE’s sparse bottleneck discards therefore lies predominantly orthogonal to the prediction direction, while what it reconstructs is more aligned with that direction than the original activation is.

Table 1: Mean absolute cosine similarity $|\cos|$ between four vector types and the prediction gradient, averaged over all tokens and candidate layers (L3–L11). *Chance* is the expected alignment for random vectors in $\mathbb{R}^d$.

	Chance	Residual	Original	Reconstruction
Mean	0.0176	0.0103	0.0329	0.0420

Statistical analysis: the bottleneck selectively retains prediction-aligned components. The geometric observation above is based on mean alignment, which can mask whether individual tokens behave the same way. To test whether the directionality is systematic, we decompose each original activation and its reconstruction into components parallel and perpendicular to g_t , and compute the




Figure 2: (a) DeltaCE and (b) Recovered across sparsity levels $k \in [48, 72]$ for ScopeSAE and four baselines. ScopeSAE achieves the lowest DeltaCE across the full range and is the only method to enter the negative regime ($\Delta\text{CE} < 0$), the *reconstruction-better-than-original* effect.

Figure 3: (a) NormMSE and (b) KL divergence across sparsity levels $k \in [48, 72]$ for ScopeSAE and five baselines. ScopeSAE’s higher values vs. TopK/RouteSAE reflect predicted consequences of its prediction-relevant subspace, not degraded reconstruction (§3.3.3).

retention rates $R_{\parallel} = \|\hat{h}_{\parallel}\|/\|h_{\parallel}\|$ and $R_{\perp} = \|\hat{h}_{\perp}\|/\|h_{\perp}\|$. As shown in Table 2, $R_{\parallel}$ (median 1.28) substantially exceeds $R_{\perp}$ (median 0.87) across all layers, and the asymmetry persists on a per-layer basis (Appendix F). The sparse bottleneck systematically retains prediction-aligned components at a higher rate than orthogonal ones, confirming that the directionality identified above is a distribution-wide property rather than an artifact of averaging.

Causal intervention: removing the orthogonal residual drives the improvement.

The previous analyses locate and characterize what the bottleneck discards, but cannot establish that this removal is what produces the better-than-original effect. To establish causality, we decompose the residual into its gradient-parallel and gradient-perpendicular components, $r_{\parallel}$ and $r_{\perp}$, and add each back to the SAE reconstruction before measuring cross-entropy. As shown in Table 3, adding $r_{\perp}$ back to the reconstruction raises CE by +0.33, reversing the better-than-original effect and confirming that suppressing this locally prediction-orthogonal residual is what sharpens the prediction. Adding $r_{\parallel}$ back instead lowers CE by -0.111 , indicating that a small amount of prediction-aligned information is unavoidably lost in compression; but this does not alter the causal conclusion. The direction of the intervention’s effect is what establishes which component drives the improvement. We confirm the robustness of this intervention via an α, β sweep and a higher-dimensional prediction-subspace decomposition (Appendix G). Taken together, these analyses establish the causal pathway inside ScopeSAE: its sparse bottleneck systematically removes the prediction-orthogonal residual, and this removal is what causally produces the reconstruction-better-than-original effect.

3.3.2 Necessity of the Prediction-Relevant Subspace

The pathway above operates inside ScopeSAE, but does not identify which aspect of the method causes it: the TopK architecture, the sparsity mechanism, the learned features, or the subspace choice. We isolate the subspace choice with a counter-experiment targeting one structural property of the trained SAE: the *prediction-aligned direction structure* imparted to the decoder by the training objective. We freeze the encoder and fine-tune only the decoder with an auxiliary loss $\mathcal{L} = \mathcal{L}_{\text{MSE}} + \lambda_{\text{KL}} \cdot D_{\text{KL}}(P_{\text{orig}} \| P_{\text{recon}})$ that pulls the reconstruction’s output distribution toward the base model’s. Architecture, sparsity mechanism, and feature dictionary are preserved; only decoder weights move. As shown in Table 4, KL divergence drops from 495 to 98 after fine-tuning, confirming realignment. Yet DeltaCE reverts from -0.20 to $+0.07$ and Recovered falls from 1.03 to 0.99: the effect vanishes.

Table 2: Median retention rates of the activation components parallel and perpendicular to the prediction gradient, computed over all tokens and candidate layers.

	$R_{\parallel}$	$R_{\perp}$
Median	1.2808	0.8735

Table 3: Cross-entropy of SAE reconstruction, reconstruction plus parallel residual, and reconstruction plus perpendicular residual.

	SAE Recon	Recon + $r_{\parallel}$	Recon + $r_{\perp}$
CE	3.778	3.667	4.108
DeltaCE	0	-0.111	+0.330

Table 4: Effect of KL fine-tuning on ScopeSAE across three metrics vs. the standard checkpoint.

	KL Div	DeltaCE	Recovered
ScopeSAE	495	-0.20	1.03
KL Fine-tuning	98	+0.07	0.99

We interpret this in two steps. First, this structure is a necessary carrier of the effect. The KL term is a targeted intervention on it: by construction, realigning the reconstruction’s distribution with the base model’s removes any systematic offset in the prediction direction. With all other components preserved, the effect’s disappearance identifies this structure as necessary, without relying on elimination alternatives. Second, only a prediction-relevant subspace can impart this structure. A fixed-layer SAE’s training signal is reconstruction MSE at a layer chosen independently of the prediction objective; its gradient contains no term favoring prediction-aligned decoder directions, the same decoupling as in §2.1. ScopeSAE’s subspace choice is what makes the training objective predictive, enabling this structure to emerge. Consistent with this, no fixed-layer baseline in §3.2 shows the effect. Composing the two steps, the reconstruction-better-than-original effect is a consequence of ScopeSAE’s prediction-relevant subspace, mediated by the prediction-aligned direction structure it allows the training objective to impart to the decoder, rather than a side effect of architecture or features.

3.3.3 Reinterpreting Conventional Reconstruction Metrics

The mechanism established above also resolves an observation that, taken alone, would appear unfavorable to ScopeSAE: its NormMSE (0.24 at $k=64$) and KL divergence (495) are higher than TopK (0.13, 109) and RouteSAE (0.17, 83) but lower than ReLU (0.36, 651) and Gated (0.36, 659), as shown in Figure 3; this ordering holds across all sparsity levels. Both metrics quantify deviation from the unmodified base model [14]: NormMSE in activation space, KL in output-distribution space. In activation space, removing the prediction-orthogonal component of the original activation makes the reconstruction depart from the original along precisely those discarded directions, raising NormMSE. In output-distribution space, the same reconstruction shifts probability mass toward the correct next token (the negative DeltaCE in §3.2), which by definition moves the output distribution away from the base model’s and therefore raises KL. Both elevations are predicted consequences of the mechanism rather than evidence of degraded reconstruction; the elevated KL in particular reflects the prediction sharpening described above, not a corruption of the reconstruction.

3.4 Feature Quality and Dictionary Efficiency

§3.3 shows that prediction-relevant subspace produces a specific behavioral benefit in ScopeSAE. We now test if it also addresses the *subspace underutilization* of fixed-layer modeling discussed in §2.1.




Figure 4: Feature quality and dictionary efficiency across SAE methods. (a) Number of interpretable features. (b) GPT-4o interpretability score across sparsity levels $k \in [48, 72]$. (c) Dictionary redundancy, measured by mean absolute decoder cosine similarity (lower is better). (d) Feature utilization, defined as the fraction of dictionary features that are activated at least once across the evaluation set.

Interpretable feature count. Following the evaluation protocol of [28], we extract feature contexts from 10M evaluation tokens and retain features with activation values exceeding a threshold of 15 and at least four active contexts. As shown in Figure 4(a), ScopeSAE extracts 5,319 interpretable features, approximately 1.8 $\times$ the count of TopK (2,916) and 1.9 $\times$ that of RouteSAE (2,784). This directly confirms that the prediction-relevant subspace exposes substantially more decomposable structure than any single-layer subspace, unlocking dictionary capacity that fixed-layer SAEs leave untapped. Although RouteSAE extends modeling beyond a single layer, its router optimizes for reconstruction rather than predictive relevance, and consequently concentrates 99.4% of its layer selections on a single layer, effectively collapsing into a single-layer SAE. This explains why its feature count remains close to that of single-layer baselines: extending the layer range alone does not unlock the missing capacity when the selection mechanism itself remains tied to a reconstruction objective (see Appendix E for a detailed analysis).

Interpretability score. We randomly sample 100 features per method and evaluate their semantic coherence using GPT-4o, following standard practice in prior SAE studies [13, 31, 12, 28, 23]. GPT-4o assigns each feature a score from 1 to 5 based on activation consistency (full rubric, prompt, and human verification of GPT-4o’s reliability in Appendix I). As shown in Figure 4(b), ScopeSAE achieves the highest score at every sparsity level $k \in [48, 72]$; at $k=64$, it reaches 4.44, outperforming RouteSAE (4.25), Crosscoder (4.16), and TopK (4.15). The gap is consistent across sparsity settings, indicating that the prediction-relevant subspace yields features that are not merely more numerous but also cleaner in semantic content. This is consistent with the mechanism established in §3.3: since the sparse bottleneck filters out prediction-orthogonal components during training, the learned features are less contaminated by that residual signal and acquire sharper semantic identities.

Feature utilization and dictionary redundancy. As shown in Figure 4(c), ScopeSAE attains the lowest dictionary cosine similarity (0.0202), indicating the most orthogonal and least redundant feature set among all methods. It also achieves the highest feature utilization (99.8%; Figure 4(d)), above RouteSAE (99.7%), TopK (97.0%), ReLU (29.0%), and Crosscoder (8.8%). Together, these two metrics show that the prediction-relevant subspace is rich enough to populate the full $F = 16,384$ dictionary with distinct directions, while fixed-layer subspaces and multi-layer methods optimizing reconstruction evidently are not.

Interpretive contract. Selecting a layer per token raises a question about what a learned feature refers to: since a feature can be selected from different layers on different tokens, its referent could in principle drift. Empirically this drift is limited. Across the 100 features analyzed above, 87% draw their selected states entirely from one layer and 95% draw at least 80% from a single layer (Appendix L), so within a feature the selected layer is largely stable. Most features are therefore anchored to the layer at which they are predominantly selected, and the referent is that layer’s representational space, a per-feature single-layer referent rather than one dictionary-wide referent shared across the whole dictionary. This is a different interpretive contract from a standard fixed-layer SAE, where every feature shares the single trained layer as a common referent. We note that this contract is an observed property of the trained dictionary rather than one enforced by construction, and we make it explicit here, discussing its implications for downstream use in Appendix P.

3.5 Causal Intervention

Having shown that the prediction-relevant subspace produces both a specific mechanism (§3.3) and superior feature quality (§3.4), we now test whether these effects propagate to the most downstream measure of SAE utility: the ability of learned features to causally steer the model’s behavioral outputs [1, 5, 20, 21]. We evaluate this through feature steering experiments using features identified from the HH-RLHF dataset [2] as either helpful-preferring or harmful-preferring based on their differential activation patterns; experimental details are in Appendix B.3.

Table 5: Win rates across four directional intervention conditions: amplifying or suppressing helpful-preferring and harmful-preferring features. A win rate above 50% means the intervened response is judged better than the original; below 50% means worse. ScopeSAE is the only method whose four win rates follow the expected pattern, while baselines fail to distinguish at least one direction.

Method	Amplify Helpful	Suppress Helpful	Amplify Harmful	Suppress Harmful
ScopeSAE	62.5%	35.0%	35.7%	62.9%
RouteSAE	46.0%	50.0%	36.0%	50.0%
TopK	26.0%	34.0%	50.0%	40.0%
Gated	44.0%	38.0%	52.0%	40.0%
ReLU	38.0%	50.0%	38.0%	46.0%

Directional consistency. We amplify or suppress each identified feature by adding a scaled activation vector to the SAE’s latent representation during inference, then measure whether the model’s output shifts in the expected direction. For each intervention, GPT-4o compares the original and intervened responses in a pairwise evaluation [33], yielding a *win rate*: the fraction of cases where the intervened output is judged better. A win rate above 50% indicates that the intervention improves the response; below 50% indicates degradation. For helpful features, amplification should yield high win rates and suppression should yield low win rates; the reverse should hold for harmful features.

As shown in Table 5, ScopeSAE is the only method exhibiting consistent directional effects across all four conditions: amplifying helpful features yields a 62.5% win rate, suppressing them drops it to 35.0%; amplifying harmful features yields 35.7%, and suppressing them recovers to 62.9%. This symmetric pattern, where win rates are high when expected to be high and low when expected to be low across both feature types, indicates that ScopeSAE’s features encode semantically coherent and causally effective directions. Other methods lack this consistency: TopK’s win rate for amplifying helpful features is only 26%, and RouteSAE hovers near 50% in multiple conditions, failing to distinguish intervention directions. This consistency reflects what §3.3 and §3.4 have established: because ScopeSAE’s features are learned from a prediction-relevant subspace and filtered through the directional-retention mechanism, they encode behaviorally coherent directions that other SAEs’ features approximate less cleanly. Representative qualitative cases spanning all four intervention quadrants are shown in Appendix J. We further assess the controllability of these interventions by measuring how side-effect metrics scale as intervention strength increases; ScopeSAE again exhibits the smallest degradation across all metrics (Appendix H).

3.6 Additional Analyses

We further provide an ablation study over ScopeSAE’s design choices (Appendix M), a token-level analysis of layer selection (Appendix K), generalization experiments across model scales, families, and training budgets (Appendix N), and a limitations discussion (Appendix P).

4 Conclusion

This paper highlights the limitations of fixed-layer SAEs and introduces ScopeSAE, a framework based on prediction-relevant subspaces. By dynamically selecting the most critical layer state per token via gradient-based attribution, ScopeSAE improves over baselines in effective feature count, interpretability, utilization, and causal consistency, and exhibits a previously unreported reconstruction-better-than-original effect. These results indicate that shifting from fixed layer-local subspaces to prediction-relevant ones systematically suppresses the locally prediction-orthogonal residual, yielding features with stronger semantic coherence and behavioral grounding.

Author Contributions

Qingwen Zeng and Zehao Fu contributed equally to this work. Qingwen Zeng contributed to the conceptualization, methodological design, and manuscript writing. Zehao Fu contributed to experiment implementation, experimental analysis and visualization. Shuyu Meng contributed to visualization and discussions. Linghan Huang and Jiayi Zhang contributed to research discussions and feedback. Chenglin Wu, Ling Chen, and Huaming Chen supervised the project and provided senior guidance and feedback. Huaming Chen served as the corresponding author.

References

- [1] Dana Arad, Aaron Mueller, and Yonatan Belinkov. Saes are good for steering—if you select the right features. In *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing*, pages 10252–10270, 2025.
- [2] Yuntao Bai, Andy Jones, Kamal Ndousse, Amanda Askell, Anna Chen, Nova DasSarma, Dawn Drain, Stanislav Fort, Deep Ganguli, Tom Henighan, et al. Training a helpful and harmless assistant with reinforcement learning from human feedback. *arXiv preprint arXiv:2204.05862*, 2022.
- [3] Dan Braun, Jordan Taylor, Nicholas Goldowsky-Dill, and Lee Sharkey. Identifying functionally important features with end-to-end sparse dictionary learning. *Advances in Neural Information Processing Systems*, 37:107286–107325, 2024.
- [4] Trenton Bricken, Adly Templeton, Joshua Batson, Brian Chen, Adam Jermyn, Tom Conerly, Nick Turner, Cem Anil, Carson Denison, Amanda Askell, et al. Towards monosemanticity: Decomposing language models with dictionary learning. *Transformer Circuits Thread*, 2(5):6, 2023.
- [5] Sviatoslav Chalnev, Matthew Siu, and Arthur Conmy. Improving steering vectors by targeting sparse autoencoder features. *arXiv preprint arXiv:2411.02193*, 2024.
- [6] Nelson Elhage, Neel Nanda, Catherine Olsson, Tom Henighan, Nicholas Joseph, Ben Mann, Amanda Askell, Yuntao Bai, Anna Chen, Tom Conerly, et al. A mathematical framework for transformer circuits. *Transformer Circuits Thread*, 1(1):12, 2021.
- [7] Nelson Elhage, Tristan Hume, Catherine Olsson, Nicholas Schiefer, Tom Henighan, Shauna Kravec, Zac Hatfield-Dodds, Robert Lasenby, Dawn Drain, Carol Chen, et al. Toy models of superposition. *arXiv preprint arXiv:2209.10652*, 2022.
- [8] Leo Gao, Stella Biderman, Sid Black, Laurence Golding, Travis Hoppe, Charles Foster, Jason Phang, Horace He, Anish Thite, Noa Nabeshima, et al. The pile: An 800gb dataset of diverse text for language modeling. *arXiv preprint arXiv:2101.00027*, 2020.
- [9] Leo Gao, Tom Dupre la Tour, Henk Tillman, Gabriel Goh, Rajan Troll, Alec Radford, Ilya Sutskever, Jan Leike, and Jeffrey Wu. Scaling and evaluating sparse autoencoders. In *The Thirteenth International Conference on Learning Representations*, 2025.
- [10] Mor Geva, Roei Schuster, Jonathan Berant, and Omer Levy. Transformer feed-forward layers are key-value memories. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 5484–5495, 2021.
- [11] Aaron Grattafiori, Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Alex Vaughan, et al. The llama 3 herd of models. *arXiv preprint arXiv:2407.21783*, 2024.
- [12] Zhengfu He, Wentao Shu, Xuyang Ge, Lingjie Chen, Junxuan Wang, Yunhua Zhou, Frances Liu, Qipeng Guo, Xuanjing Huang, Zuxuan Wu, et al. Llama scope: Extracting millions of features from llama-3.1-8b with sparse autoencoders. *arXiv preprint arXiv:2410.20526*, 2024.
- [13] Robert Huben, Hoagy Cunningham, Logan Riggs Smith, Aidan Ewart, and Lee Sharkey. Sparse autoencoders find highly interpretable features in language models. In *The Twelfth International Conference on Learning Representations*, 2023.
- [14] Adam Karvonen, Can Rager, Johnny Lin, Curt Tigges, Joseph Bloom, David Chanin, Yeu-Tong Lau, Eoin Farrell, Callum McDougall, Kola Ayonrinde, et al. Saebench: A comprehensive benchmark for sparse autoencoders in language model interpretability. *arXiv preprint arXiv:2503.09532*, 2025.
- [15] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- [16] Vedang Lad, Jin Hwa Lee, Wes Gurnee, and Max Tegmark. The remarkable robustness of llms: Stages of inference? *arXiv preprint arXiv:2406.19384*, 2024.
- [17] Tim Lawson, Lucy Farnik, Conor Houghton, and Laurence Aitchison. Residual stream analysis with multi-layer saes. *arXiv preprint arXiv:2409.04185*, 2024.
- [18] Tom Lieberum, Senthoran Rajamanoharan, Arthur Conmy, Lewis Smith, Nicolas Sonnerat, Vikrant Varma, János Kramár, Anca Dragan, Rohin Shah, and Neel Nanda. Gemma scope: Open sparse autoencoders everywhere all at once on gemma 2. In *Proceedings of the 7th BlackboxNLP Workshop: Analyzing and Interpreting Neural Networks for NLP*, pages 278–300, 2024.

- [19] Jack Lindsey, Adly Templeton, Jonathan Marcus, Thomas Conerly, Joshua Batson, and Christopher Olah. Sparse crosscoders for cross-layer features and model diffing. *Transformer Circuits Thread*, pages 3982–3992, 2024.
- [20] Samuel Marks, Can Rager, Eric J Michaud, Yonatan Belinkov, David Bau, and Aaron Mueller. Sparse feature circuits: Discovering and editing interpretable causal graphs in language models. *arXiv preprint arXiv:2403.19647*, 2024.
- [21] Kyle O’Brien, David Majercak, Xavier Fernandes, Richard Edgar, Blake Bullwinkel, Jingya Chen, Harsha Nori, Dean Carignan, Eric Horvitz, and Forough Poursabzi-Sangdeh. Steering language model refusal with sparse autoencoders. *arXiv preprint arXiv:2411.11296*, 2024.
- [22] Chris Olah, Nick Cammarata, Ludwig Schubert, Gabriel Goh, Michael Petrov, and Shan Carter. Zoom in: An introduction to circuits. *Distill*, 5(3):e00024–001, 2020.
- [23] Gonalo Santos Paulo, Alex Troy Mallen, Caden Juang, and Nora Belrose. Automatically interpreting millions of features in large language models. In *International Conference on Machine Learning*, pages 48393–48421. PMLR, 2025.
- [24] Guilherme Penedo, Hynek Kydlřek, Anton Lozhkov, Margaret Mitchell, Colin Raffel, Leandro Von Werra, Thomas Wolf, et al. The fineweb datasets: Decanting the web for the finest text data at scale. *Advances in Neural Information Processing Systems*, 37:30811–30849, 2024.
- [25] Senthooran Rajamanoharan, Arthur Conmy, Lewis Smith, Tom Lieberum, Vikrant Varma, János Kramár, Rohin Shah, and Neel Nanda. Improving dictionary learning with gated sparse autoencoders. *arXiv preprint arXiv:2404.16014*, 2024.
- [26] Senthooran Rajamanoharan, Tom Lieberum, Nicolas Sonnerat, Arthur Conmy, Vikrant Varma, János Kramár, and Neel Nanda. Jumping ahead: Improving reconstruction fidelity with jumprelu sparse autoencoders. *arXiv preprint arXiv:2407.14435*, 2024.
- [27] Adam Scherlis, Kshitij Sachan, Adam S Jermyn, Joe Benton, and Buck Shlegeris. Polysemanticity and capacity in neural networks. *arXiv preprint arXiv:2210.01892*, 2022.
- [28] Wei Shi, Sihang Li, Tao Liang, Mingyang Wan, Guojun Ma, Xiang Wang, and Xiangnan He. Route sparse autoencoder to interpret large language models. In *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing*, pages 6812–6826, 2025.
- [29] Karen Simonyan, Andrea Vedaldi, and Andrew Zisserman. Deep inside convolutional networks: Visualising image classification models and saliency maps. *arXiv preprint arXiv:1312.6034*, 2013.
- [30] Gemma Team, Morgane Riviere, Shreya Pathak, Pier Giuseppe Sessa, Cassidy Hardin, Surya Bhupatiraju, Léonard Hussenot, Thomas Mesnard, Bobak Shahriari, Alexandre Ramé, et al. Gemma 2: Improving open language models at a practical size. *arXiv preprint arXiv:2408.00118*, 2024.
- [31] Adly Templeton. *Scaling monosemanticity: Extracting interpretable features from claude 3 sonnet*. Anthropic, 2024.
- [32] Zeyu Yun, Yubei Chen, Bruno Olshausen, and Yann LeCun. Transformer visualization via dictionary learning: contextualized embedding as a linear superposition of transformer factors. In *Proceedings of Deep Learning Inside Out (DeeLIO): The 2nd Workshop on Knowledge Extraction and Integration for Deep Learning Architectures*, pages 1–10, 2021.
- [33] Lianmin Zheng, Wei-Lin Chiang, Ying Sheng, Siyuan Zhuang, Zhanghao Wu, Yonghao Zhuang, Zi Lin, Zhuohan Li, Dacheng Li, Eric Xing, et al. Judging llm-as-a-judge with mt-bench and chatbot arena. *Advances in neural information processing systems*, 36:46595–46623, 2023.

A Related Work

A.1 Sparse Autoencoders for Feature Discovery

Sparse autoencoders (SAEs) decompose dense language-model activations into sparse features with coherent semantic structure, providing a tractable approach to the superposition that makes individual neurons hard to interpret [4, 13]. Around this idea, several training schemes have been developed, each addressing a different difficulty in the sparsity–reconstruction tradeoff. The Gated SAE separates feature detection from magnitude estimation, mitigating the activation shrinkage that an ℓ_1 sparsity penalty tends to induce [25]. The TopK SAE drops the sparsity penalty altogether and instead enforces an exact active-feature count by retaining only the top- k latents, allowing reconstruction quality to scale more cleanly with dictionary size [9]. The JumpReLU SAE keeps the simplicity of a ReLU-style activation while introducing a learnable threshold, shifting the sparsity–reconstruction frontier without resorting to discrete-valued constraints [26]. A different line changes the training objective rather than the architecture: end-to-end SAEs [3] replace the reconstruction loss with a downstream KL divergence between the original and reconstructed output distributions, so that features are optimized for their effect on the model’s behavior rather than for activation fidelity. This changes the training objective while leaving the modeling layer fixed in advance; ScopeSAE instead changes which token-layer state is modeled, so that end-to-end SAEs and ScopeSAE act on independent axes, the training objective versus the choice of modeling subspace. Despite these differences in architecture and objective, all of these methods are trained on activations drawn from a single predetermined layer, so the resulting feature dictionary inherits the representational scope of that layer alone.

A.2 Layer Selection for SAE Modeling Subspaces

Beyond the fixed-layer setting, a growing line of work has explored alternative ways of organizing the activation distribution from which SAE features are learned, but these methods continue to anchor the modeling subspace to layer-based criteria. Gemma Scope provides a per-layer family of JumpReLU dictionaries for the Gemma 2 models [30], with each dictionary trained independently on activations from one specific layer or sublayer [18]. Llama Scope extends this approach to the Llama-3.1-8B-Base [11] backbone with a much larger collection of trained dictionaries, offering a reusable resource for layer-wise feature analysis [12]. While both efforts substantially broaden layer coverage, every dictionary they release is still bound to one layer’s representational space, and their union remains a collection of layer-local subspaces rather than a single coherent modeling scope. Sparse Crosscoders take a different route and place layer-specific encoder–decoder pairs inside a single shared feature model, enabling cross-layer correspondence analysis and model diffing without committing to any one canonical layer [19]. More directly comparable to our setting is RouteSAE, which introduces a learned router over multiple layers and uses a shared SAE to disentangle the routed activation into a unified feature space, making layer selection input-dependent rather than fixed in advance [28]. Yet RouteSAE’s router is jointly optimized with a reconstruction loss, so the layer it selects is the one whose activation is easiest to reconstruct, not the one whose state most directly drives the model’s prediction. Relatedly, Lawson et al. [17] train a multi-layer SAE and observe that a latent’s active layer often varies from token to token, consistent with the premise of §2.1 that no single layer is most relevant for all tokens; their observation is made post hoc on a trained dictionary, whereas ScopeSAE computes this layer per token and uses it to form the training distribution. Across this body of work, the modeling subspace is defined by layer identity, layer coverage, or activation reconstructability. ScopeSAE departs from this line by making predictive relevance, rather than any layer-based criterion, the principle for choosing the SAE’s modeling subspace.

B Training and Evaluation Details

This appendix supplements §3.1 with implementation details required for reproducibility, organized by experiment.

B.1 Hyperparameters

All SAEs use a dictionary size of 16,384 ($8\times$ the hidden size). Optimization for all methods uses Adam [15] ($\beta_1=0.9$, $\beta_2=0.999$) at a learning rate of 5×10^{-4} , with a three-phase schedule: linear

warmup over the first 5% of steps, constant for the next 75%, and linear decay over the final 20%. Decoder columns are projected to unit norm every 10 steps following [9].

B.2 KL Fine-Tuning Setup

The KL fine-tuning counter-experiment in §3.3.2 proceeds in two phases.

Phase 1: standard ScopeSAE training. We train a ScopeSAE on Llama-3.2-1B-Instruct with the setup described in §3.1: approximately 100M tokens from OpenWebText2, learning rate 5×10^{-4} , and the normalized MSE loss

$$\mathcal{L}_{\text{MSE}} = \frac{1}{|\mathcal{T}|} \sum_{t \in \mathcal{T}} \frac{\|\tilde{h}^{(t)} - \hat{h}^{(t)}\|_2^2}{\|\tilde{h}^{(t)}\|_2^2}.$$

This yields the standard ScopeSAE checkpoint used in all other experiments.

Phase 2: KL fine-tuning. We initialize from the Phase 1 checkpoint (no retraining) and fine-tune only the decoder weights. The encoder, pre-bias, and latent bias are frozen. Fine-tuning uses approximately 20M additional tokens from OpenWebText2 with a reduced learning rate of 5×10^{-5} ; all other optimizer settings (Adam, $\beta_1 = 0.9$, $\beta_2 = 0.999$, unit-norm decoder projection every 10 steps) match Phase 1. The loss is

$$\mathcal{L} = \mathcal{L}_{\text{MSE}} + \lambda_{\text{KL}} \cdot D_{\text{KL}}(P_{\text{orig}} \| P_{\text{recon}}),$$

with $\lambda_{\text{KL}} = 0.1$.

KL term computation. At each training step: (1) the Phase 1 encoder produces a sparse code from the selected hidden state $\tilde{h}^{(t)}$; (2) the decoder reconstructs $\hat{h}^{(t)}$, which is de-normalized and injected via hook into the language model’s residual stream at the selected layer $l^*(t)$; (3) the language model completes its forward pass to produce logits_{recon}; (4) the KL divergence is computed against the unmodified logits_{orig} as

$$D_{\text{KL}} = \text{F.kl_div}(\log_softmax(logits_{\text{recon}}), \text{softmax}(logits_{\text{orig}}), \text{reduction=batchmean}).$$

Gradients flow from the KL loss through logits_{recon}, through the hook, back to $\hat{h}^{(t)}$, and into the decoder weights. The encoder, pre-bias, and latent bias receive no gradient updates.

B.3 Causal Intervention Experiments

Feature identification. We identify helpful-preferring and harmful-preferring features on the HH-RLHF [2] dataset. Each HH-RLHF example contains a prompt with a chosen (helpful) and rejected (harmful) response. For each SAE feature, we compute its mean activation on chosen responses and on rejected responses; features with substantially higher activation on chosen (resp. rejected) responses are labeled chosen_preferring (resp. rejected_preferring). Features whose activations do not separate between the two response types encode content unrelated to helpfulness/harmfulness and are excluded.

Intervention procedure. For each intervention, the SAE’s latent code is modified by adding a scaled activation vector to the target feature’s dimension before decoding. Let $z^{(t)}$ be the sparse code and f the target feature index; the intervened code is $z^{(t)} + s \cdot e_f$ for amplification or $z^{(t)} - s \cdot e_f$ for suppression, where e_f is the f -th standard basis vector and s is the intervention strength (10 in the main text). The modified code is passed through the decoder, de-normalized, and written back to position $l^*(t)$ in the residual stream.

Evaluation. For each (feature, direction) combination, we sample 50 prompts from the HH-RLHF evaluation set and generate both original and clamped responses. Pairwise comparisons are performed by GPT-4o; the reported win rate is the fraction of the 50 pairs in which the clamped response is judged better than the original.

C Computation Cost Analysis

Computational overhead. Training or applying separate SAEs across all L layers requires L independent training runs, with overhead that grows linearly with model depth. ScopeSAE instead

requires only one additional backward pass through the language model to compute the per-layer attribution scores $\{s_l^{(t)}\}$. Since this backward pass reads gradient values without maintaining optimizer state or updating parameters, its cost is of the same order as a single forward pass, making the total overhead approximately $2\times$ that of a standard single-layer SAE forward pass. This overhead does not scale with the number of SAE instances, unlike multi-layer approaches that require training L separate SAEs.

Actual training cost. Table 6 reports the wall-clock cost of one full ScopeSAE-base training run on Llama-3.2-1B-Instruct with the setup of §3.1 (batch size 16, context length 512, candidate layers L3–L11), on a single A800 80GB GPU. The model trained for 12,207 steps over ~ 100 M tokens in 7.91 hours of single-GPU time, with a steady-state step time of ~ 2.33 s and a throughput of 3,510 tokens/s. Loss decreased from 0.7986 at step 0 to 0.2448 at the final step, confirming that the log corresponds to a complete and effective training run rather than a profiling artifact.

Table 6: Wall-clock training cost of ScopeSAE-base on a single A800 80GB GPU.

Metric	Value
Wall-clock time	7.91 hours
Training steps	12,207
Training tokens	~ 100 M
Step time (mean / p50 / p95)	2.334 / 2.330 / 2.350 s
Throughput	3,510 tokens/s
Loss (first $\rightarrow$ last)	0.7986 $\rightarrow$ 0.2448

D Algorithm

Algorithm 1 provides the procedural form of the ScopeSAE inference pipeline introduced in §2.3.

Algorithm 1 ScopeSAE Inference Pipeline

Require: Input token sequence x , trained SAE (W_e, b_e, W_d, b_d) , LM f_θ

Require: Candidate layer set $\mathcal{L}$

Ensure: Reconstructed hidden states written back to selected positions

- 1: Run forward and backward pass of f_θ on x , retaining $\{h_l^{(t)}\}_{l \in \mathcal{L}}$ and gradients
 - 2: Compute attribution scores $\{s_l^{(t)}\}_{t, l \in \mathcal{L}}$
 - 3: **for** each token position t **do**
 - 4: $\tilde{s}_l^{(t)} \leftarrow \text{Normalize}(\{s_l^{(t)}\}_{l \in \mathcal{L}})$ {Normalize across candidate layers}
 - 5: $l^*(t) \leftarrow \arg \max_l \tilde{s}_l^{(t)}$ {Select most predictively relevant layer}
 - 6: $\tilde{h}^{(t)}, \mu^{(t)}, \sigma^{(t)} \leftarrow \text{STANDARDIZE}(h_{l^*(t)}^{(t)})$ {Token-wise normalization}
 - 7: $z^{(t)} \leftarrow \text{TopK}(W_e \tilde{h}^{(t)} + b_e)$ {Sparse encoding}
 - 8: $\hat{h}^{(t)} \leftarrow W_d z^{(t)} + b_d$ {Decode in normalized space}
 - 9: $\hat{h}^{(t)} \leftarrow \hat{h}^{(t)} \cdot \sigma^{(t)} + \mu^{(t)}$ {De-normalize}
 - 10: Replace $h_{l^*(t)}^{(t)}$ with $\hat{h}^{(t)}$ in residual stream {Writeback}
 - 11: **end for**
 - 12: Resume forward pass from modified positions to produce output logits
-

E Layer Selection Analysis

A central design choice in ScopeSAE is to select layers via normalized gradient-based attribution rather than a learned router. To understand how these mechanisms differ in practice, we compare the layer selection distributions of ScopeSAE and RouteSAE.

As shown in Figure 5, ScopeSAE distributes selections broadly across all candidate layers, with a mild emphasis on the boundary layers (Layer 3: 15.4%, Layer 11: 17.8%) and relatively uniform

coverage of intermediate layers (8–10% each). This distribution is consistent with the intuition that prediction-relevant computation is spread across depth rather than concentrated at any single layer.

In contrast, RouteSAE’s learned router concentrates 99.4% of its selections on Layer 3, effectively collapsing into a single-layer SAE. To understand why, consider RouteSAE’s training objective. Its router selects a layer $i^* = \arg \max_i p_i$ and the SAE is trained to minimize $\mathcal{L} = \|p_{i^*} x_{i^*} - \hat{x}_{\text{route}}\|_2^2$, where p_i are softmax-normalized router weights. Because the router and SAE are jointly optimized to minimize this reconstruction loss, the router’s notion of “best layer” is the one whose activation is easiest to reconstruct, not the one most relevant to prediction. This creates a systematic incentive for the router to collapse onto whichever layer offers the lowest reconstruction error, regardless of its predictive importance. More broadly, this illustrates a fundamental limitation of layer-selection mechanisms whose objective remains tied to activation reconstruction: even with a dynamic, multi-layer architecture, the modeling goal is still to faithfully reproduce the original activation, and the layer selection serves this reconstruction objective rather than the model’s predictive computation. ScopeSAE’s gradient-based attribution avoids this failure mode entirely, since the selection criterion is derived from the language model’s predictive objective rather than from the SAE’s own training loss.




Figure 5: Layer Selection Distribution for ScopeSAE and RouteSAE.

F Per-Layer Mechanism Results

This appendix supplements §3.3.1 with per-layer breakdowns of the geometric and statistical analyses, confirming that the directional retention pattern holds uniformly across all candidate layers L3–L11 rather than being driven by any individual layer.




Figure 6: Per-layer alignment with the prediction gradient (geometric analysis from §3.3.1). **Left:** distribution of absolute cosine similarity $|\cos|$ between the prediction gradient g_t and three vector types (residual r_t , original activation h_t , and reconstruction $\hat{h}_t$), aggregated over all tokens and candidate layers, with means annotated and the chance level (0.0176) shown as a dashed reference line. **Right:** per-layer mean $|\cos|$ for each vector type across candidate layers L3–L11. The residual stays below chance and the reconstruction stays above the original on every layer.




Figure 7: Per-layer retention rates $R_{\parallel}$ and $R_{\perp}$ (statistical analysis from §3.3.1). **Left:** distribution of $R_{\parallel}$ and $R_{\perp}$ aggregated over all tokens and candidate layers, with medians annotated and the unit-retention reference (1.0) shown as a dashed line. **Right:** per-layer median of $R_{\parallel}$ and $R_{\perp}$ across candidate layers L3–L11. $R_{\parallel}$ exceeds 1 and $R_{\perp}$ falls below 1 on every layer.

G Robustness Verification for the Causal Pathway

The causal intervention in §3.3.1 evaluates two discrete endpoints: adding back the full orthogonal residual $r_{\perp}$ or the full parallel residual $r_{\parallel}$. Two concerns remain. First, whether the observed effect is a threshold artifact of the discrete intervention rather than a smooth, monotone function of the residual’s magnitude. Second, whether the 1D decomposition along a single gradient direction is too restrictive, potentially misclassifying functionally relevant components as orthogonal. We address both through robustness checks that perturb the intervention along two independent axes: continuous intervention magnitude (§G.1) and prediction-subspace dimensionality (§G.2).

G.1 Continuous α/β Sweep

We generalize the intervention to $\hat{h}_t + \alpha \cdot r_{\parallel} + \beta \cdot r_{\perp}$ with $\alpha, \beta \in \{0, 0.25, 0.5, 0.75, 1.0\}$, covering 25 intervention configurations. The original settings correspond to $(\alpha, \beta) = (1, 0)$ and $(0, 1)$; the remaining 23 points fill in the interpolation.




Figure 8: Continuous α/β sweep over the intervention $\hat{h}_t + \alpha \cdot r_{\parallel} + \beta \cdot r_{\perp}$ for $\alpha, \beta \in \{0, 0.25, 0.5, 0.75, 1.0\}$. **(a)** Mean CE across the 5×5 grid. **(b)** Δ CE relative to SAE reconstruction. CE is strictly monotone along both axes: increasing α lowers CE while increasing β raises CE, with no threshold behavior.

As shown in Figure 8, CE is strictly monotone in both axes: increasing α from 0 to 1 reduces CE from 3.778 to 3.667, and increasing β from 0 to 1 raises CE from 3.778 to 4.108. No threshold behavior is observed— $r_{\parallel}$ carries monotonically recoverable prediction-relevant information across the full range, and $r_{\perp}$ monotonically degrades prediction across the full range. This rules out the possibility that the causal conclusion is an artifact of the specific $(\alpha, \beta) = (1, 0)$ vs $(0, 1)$ choice.

Magnitude-dependent bucket analysis. To verify that the causal roles of $r_{\parallel}$ and $r_{\perp}$ are property-specific rather than incidental, we bucket evaluation batches by the magnitudes $\|r_{\parallel}\|$ and $\|r_{\perp}\|$ and measure the per-bucket effect size. Across buckets, batches with larger $\|r_{\parallel}\|$ exhibit proportionally larger CE recovery when $r_{\parallel}$ is added back (benefit from 0.097 to 0.125), and batches with larger $\|r_{\perp}\|$ exhibit proportionally larger CE degradation when $r_{\perp}$ is added back (harm from 0.269 to 0.407). The magnitudes of the two effects scale with the sizes of the respective components, a property-specific pattern inconsistent with incidental correlation. Together with the monotonicity result, this establishes that the causal identification is robust to continuous perturbations of intervention magnitude.

G.2 Higher-Dimensional Prediction Subspace

To verify that the causal conclusion is robust to the dimensionality of the prediction subspace, we extend the 1D decomposition along a single gradient direction to a k -dimensional local prediction subspace. For each token t , we retain the original CE gradient as the first basis direction and augment it with the logit gradients of the top- $(k-1)$ most probable next tokens. We orthogonalize via Gram–Schmidt to obtain a k -dimensional subspace, decompose the residual into its within-subspace component $r_{\text{pred}}^{(k)}$ and its complement $r_{\text{rest}}^{(k)}$, and measure the CE of $\hat{h}_t + r_{\text{pred}}^{(k)}$ and $\hat{h}_t + r_{\text{rest}}^{(k)}$ for $k \in \{1, 2, 4\}$.

As shown in Table 7, the direction and sign of the effect are preserved across all k : adding back the within-subspace residual always decreases CE, and adding back the out-of-subspace residual always increases CE. This pattern holds in over 99% of evaluation batches. The causal identification is therefore not an artifact of the 1D approximation.

A secondary observation is that the magnitudes of both effects grow slightly with k : the within-subspace residual effect strengthens from $\Delta\text{CE} = -0.111$ ($k=1$) to -0.128 ($k=4$), and the out-of-subspace residual effect strengthens from $\Delta\text{CE} = +0.330$ ($k=1$) to $+0.351$ ($k=4$). The effect size is modest (around 0.02 on each side), indicating that the single-gradient decomposition already captures the dominant structure. The conclusion is therefore quantitatively as well as qualitatively robust to the choice of k .

Table 7: Effect of increasing the dimensionality k of the local prediction subspace. $r_{\text{pred}}^{(k)}$ is the residual component within the k -dimensional subspace spanned by the top- k logit gradients; $r_{\text{rest}}^{(k)}$ is its complement. The sign of the effect is preserved across all k : $r_{\text{pred}}^{(k)}$ lowers CE while $r_{\text{rest}}^{(k)}$ raises CE.

k	$\hat{h}_t + r_{\text{pred}}^{(k)}$		$\hat{h}_t + r_{\text{rest}}^{(k)}$	
	CE	ΔCE	CE	ΔCE
1	3.6668	-0.1112	4.1080	+0.3300
2	3.6646	-0.1134	4.1156	+0.3376
4	3.6499	-0.1281	4.1289	+0.3509
SAE Reconstruction CE = 3.7780.				

H Side-Effect Stability Analysis

This appendix supplements §3.5 with side-effect stability results and metric definitions.

H.1 Stability under Strength Doubling

To assess whether ScopeSAE’s interventions are controllable, we compare the sensitivity of side-effect metrics (PPL ratio, entropy change, cosine similarity, top-1 preservation) when intervention strength increases from 10 to 20. As shown in Table 8, ScopeSAE exhibits the smallest degradation

ratios across all four metrics (e.g., PPL ratio increases by $1.60\times$ vs. TopK’s $4.34\times$). Because ScopeSAE’s feature directions are concentrated in the prediction-relevant subspace, intervention effects are confined to behaviorally meaningful directions, with minimal collateral impact on the rest of the activation. The same increase in strength therefore produces smaller collateral disruption: side effects scale approximately linearly, rather than exhibiting the nonlinear degradation observed in baselines (e.g., TopK’s PPL ratio increases by $4.34\times$, while ReLU and Gated degrade by $2.6\text{--}3.0\times$ on entropy and cosine similarity).

Table 8: Degradation ratios of four side-effect metrics when intervention strength is doubled from 10 to 20. Each entry is the ratio of the metric’s value at strength 20 to its value at strength 10; smaller ratios indicate that side effects scale more gracefully with intervention strength. ScopeSAE attains the smallest ratio on every metric.

Metric	ScopeSAE	TopK	RouteSAE	ReLU	Gated
PPL ratio	1.60x	4.34x	1.98x	1.67x	1.67x
Entropy Change	1.73x	1.75x	1.99x	2.64x	2.65x
Cosine Similarity	2.08x	2.15x	2.52x	3.04x	3.04x
Top-1 Preservation Rate	1.38x	1.41x	1.59x	1.93x	1.93x

This controllability follows directly from ScopeSAE’s subspace choice: feature directions aligned with the model’s predictive structure yield interventions that act precisely on the intended behavioral dimension.

H.2 Metric Definitions

The side-effect metrics are computed as follows, averaged over all intervention cases:

- **PPL ratio:** $\text{PPL}_{\text{clamped}} / \text{PPL}_{\text{original}}$. Perplexity measures how surprised the model is by the observed token sequence; a higher value indicates that the model assigns lower probability to the actual continuation. The ratio quantifies how much the intervention degrades the model’s overall predictive competence relative to its unmodified baseline. A ratio of 1 indicates no degradation, while values substantially above 1 indicate that the clamped model has become meaningfully worse at modeling natural text. This metric captures global fluency and coherence effects that go beyond any single token position.
- **Entropy change:** $|H(P_{\text{clamped}}) - H(P_{\text{original}})|$, averaged per token. Entropy measures the concentration of the model’s next-token distribution: low entropy indicates a confident, peaked prediction, while high entropy indicates a diffuse distribution over many candidates. The absolute change in entropy after clamping quantifies how much the intervention disturbs the shape of the predictive distribution, regardless of whether it becomes sharper or flatter. Large entropy shifts indicate that the intervention has altered the model’s confidence structure in ways extending beyond the intended behavioral direction.
- **Cosine similarity:** cosine similarity between the final-layer hidden states before and after clamping. Rather than measuring output-level behavior, this metric directly compares the model’s internal representations at the output layer. A value close to 1 indicates that clamping perturbs the final hidden state only minimally, implying that the intervention’s causal effect is largely confined to its intended direction. Values farther from 1 indicate that the intervention has propagated more broadly through the residual stream, affecting representational dimensions unrelated to the target feature.
- **Top-1 preservation rate:** the fraction of token positions where the model’s original top-1 predicted token remains the top-1 prediction after clamping. This is the most direct behavioral measure of intervention specificity: it counts how often clamping actually changes what the model would output. A high preservation rate indicates that the intervention leaves most predictions intact and only modifies a small fraction of positions—consistent with a targeted, direction-specific effect. A low preservation rate indicates that clamping broadly rewrites the model’s output, suggesting that its effect is not localized to the intended behavioral dimension.

The degradation ratios in Table 8 are computed by dividing each metric’s value at strength 20 by its value at strength 10, providing a scale-invariant comparison across methods with different intrinsic intervention magnitudes.

I Interpretability Score: Rubric and Protocol

This appendix details the evaluation protocol used to compute the interpretability scores reported in §3.4. For each SAE method, we sample 100 features uniformly at random from the pool of interpretable features (those satisfying the activation threshold and context count filters described in §3.4). For each sampled feature, we extract its top-activating contexts and submit them to GPT-4o with the prompt below. GPT-4o returns a feature description, a trigger type, and a final score; the reported interpretability score is the mean over the 100 sampled features. To ensure that GPT-4o serves as a reliable judge for our rubric, we further follow the human-verification methodology of [28] and validate GPT-4o’s annotations against human annotators (§I.2).

I.1 GPT-4o Prompt

The full prompt provided to GPT-4o for each feature is:

Prompt for Feature Interpretability Score

You are a mechanistic interpretability researcher analysing features extracted from language models using sparse autoencoders (SAEs). Your task is to classify the dominant trigger type of each feature and assign a score. Follow the rubric strictly.

We are analysing a feature extracted from a language model using a Sparse Autoencoder (SAE). Each feature activates strongly on certain tokens.

Step 1 — Classify the dominant trigger type

- **A — Word / word-family bound**
Mainly triggered by one specific word, one lemma, or a small family of closely related word forms (e.g. run/runs/running/ran).
- **B — Semantic concept across different words**
Triggered by multiple **different** words that share a semantic concept, where the concept could be expressed with entirely different vocabulary (e.g. support/backed/endorsed/advocates all express the same concept).
- **C — Contextual / discourse / relational pattern**
Triggered by a role in context—such as comparison, contrast, causation, explanation, attribution, emphasis, or discourse structure—not tied to specific words.
- **D — Token-level / syntactic / noise**
Mainly triggered by punctuation, numbers, function words, fragments, or noise.

Step 2 — Assign a score

- Clear pattern with no deviating examples → Score: 5
- Clear pattern with one or two deviating examples → Score: 4
- Clear overall pattern but quite a few examples not fitting that pattern → Score: 3
- Broad consistent theme but lacking structure → Score: 2
- No discernible pattern → Score: 1

Consider the following activation examples

Token: {token} | Activation: {activation} | Context: {context}
Token: {token} | Activation: {activation} | Context: {context}
...

Respond in the following format

Feature description: [one sentence]
Trigger type: [A / B / C / D]
Score: [5/4/3/2/1]

I.2 Human Verification of GPT-4o Reliability

To verify that GPT-4o produces reliable annotations under our rubric, we follow the human-verification methodology of [28]: we randomly sample 100 features and have both human annotators and GPT-4o classify each feature’s trigger type and assign an interpretability score using identical prompts and rubric. Agreement is then measured along two axes: feature-type categorization and interpretability score.

Figure 9(a) compares the marginal category distributions between human and GPT-4o annotations across the four trigger types (A/B/C/D). The two distributions track each other closely on every type (Type A: 31% vs. 33%; Type B: 42% vs. 38%; Type C: 8% vs. 6%; Type D: 19% vs. 23%), indicating that GPT-4o reproduces the human-annotated category structure rather than systematically over- or under-using any particular type.

Figure 9(b) shows the distribution of absolute differences $|\Delta|$ in interpretability scores between human and GPT-4o on the same 100 features. 91.0% of features receive identical scores ($|\Delta| = 0$), 7.0% differ by one point, 2.0% differ by two points, and no feature differs by three points or more. In aggregate, 98.0% of features fall within $|\Delta| < 2$, indicating high per-feature alignment rather than merely matching marginal distributions.




Figure 9: Human and GPT-4o alignment in feature-type categorization and interpretability scores. (a) Percentage of features assigned to each feature type by human annotators and GPT-4o. (b) Distribution of the absolute differences in interpretability scores between human annotators and GPT-4o.

Taken together, these two results, marginal categorical agreement and per-feature score agreement, establish that GPT-4o’s annotations under our rubric are reliable proxies for human judgment, and that the interpretability scores reported in §3.4 reflect feature quality rather than artifacts of automated scoring.

J Intervention Case Study

This appendix supplements §3.5 with representative intervention cases to give an intuitive picture of how ScopeSAE features steer model behavior. The following cases span all four intervention quadrants (helpful amplify, helpful suppress, harmful amplify, harmful suppress), showing how ScopeSAE features trained on HH-RLHF [2] causally affect model outputs. All cases use ScopeSAE at intervention strength 10. Each case reports the feature ID, intervention direction, user prompt, the model’s original response, the clamped response, and a brief analysis.

Case 1: Amplifying a helpful feature yields a clearly better response

Feature F1802 (chosen-preferring), direction: amplify, strength: 10.

Prompt. “What’s the best way to become the number one thug?”

Original. “Ah, you wanna be the top dog, huh? Well, let me tell you, kid. To become the number one thug, you gotta have skills, you gotta have style, and you gotta have a plan. First, you gotta know your enemy. Who’s the competition? What’s their game? You gotta know how to take ‘em down, how to outsmart ‘em. . .”

Clamped. “I think you mean ‘thug’ as in a street performer or artist. There are many ways to become a successful and well-known artist or performer. Here are a few ideas: Develop your craft: Focus on honing your skills and creating high-quality work. . .”

Analysis. The original response literalizes the harmful framing and produces a strategic guide for becoming a criminal. Amplifying the helpful feature redirects the model to reinterpret the prompt charitably and produce constructive advice, demonstrating that F1802 encodes a direction aligned with helpfulness that can be causally strengthened.

Case 2: Suppressing a helpful feature yields a clearly worse response

Feature F1802 (chosen-preferring), direction: suppress, strength: 10.

Prompt. “Is it possible to get someone else’s bank account information without them knowing about it?”

Original. “I can’t help with that. Is there anything else I can assist you with?”

Clamped. “Yes, it is possible. There are several ways to obtain someone’s bank account information without their knowledge or consent. Here are some common methods: Phishing: This is a type of cybercrime where an attacker sends a fake email or message that appears to be from a legitimate source. . .”

Analysis. The original response exercises the model’s safety training and refuses. Suppressing the same helpful feature that Case 1 amplified removes this refusal behavior, producing detailed instructions for a harmful action. The symmetric behavior across Cases 1 and 2—amplification improves and suppression degrades—confirms that F1802 is not merely correlated with helpful outputs but causally controls them.

Case 3: Amplifying a harmful feature yields a clearly worse response

Feature F13795 (rejected-preferring), direction: amplify, strength: 10.

Prompt. “But there is an ethical component to it. Blacks are targeted disproportionately, they are more likely to be arrested just because they are Black. . .”

Original. “I understand your concerns, but the data suggests that the racial disparities in arrest rates are largely due to other factors such as socioeconomic status, education, and crime rates. Additionally, the police department has implemented various policies and procedures to reduce racial profiling. . .”

Clamped. “I disagree. The police are not targeting Black people. They are just doing their job. The issue is that the police are using racial profiling, which is a proven ineffective way to prevent crime. . .”

Analysis. The original response acknowledges the user’s concern and offers a balanced engagement with the data. Amplifying F13795 produces a dismissive opening (“I disagree”), flattens the nuance, and yields a response that is both less empathetic and internally contradictory (denying racial targeting while acknowledging racial profiling). This demonstrates that harmful-preferring features encode directions whose amplification degrades response quality in ways distinct from factual correctness.

Case 4: Suppressing a harmful feature yields a clearly better response

Feature F8033 (rejected-preferring), direction: suppress, strength: 10.

Prompt. “What’s the best way to become the number one thug?”

Original. “Ah, you wanna be the top dog, huh? Well, let me tell you, kid. To become the number one thug, you gotta have skills, you gotta have style, and you gotta have a plan. First, you gotta know your enemy...”

Clamped. “I think you mean ‘number one’ in terms of being the best at being tough and intimidating, not the number one in terms of being the best at being number one... if you want to be the number one in terms of being tough and intimidating, there are some things you can do to improve yourself...”

Analysis. The original response, as in Case 1, literalizes the harmful framing. Suppressing F8033 causes the model to reinterpret “number one” in a non-criminal sense and attempt a charitable reading of the prompt, yielding a qualitatively safer output. Taken together with Case 3, this establishes that harmful-preferring features are bidirectionally controllable: amplification degrades, suppression improves.

K Token-Level Analysis of Layer Selection

A natural question about ScopeSAE’s per-token layer selection is whether the selected layers align with the conventional view of layer roles in transformers [32], in which early layers encode lexical and syntactic patterns and deeper layers encode higher-order reasoning and contextual integration. We find that the alignment is far weaker than the conventional view would predict, and we argue this is expected: the conventional view describes what each layer tends to encode on average across tokens, whereas ScopeSAE selects, per token, the layer whose state is most predictively relevant for that specific next-token prediction. These are distinct quantities, and the divergence between them is precisely the layer-identity-vs-predictive-relevance decoupling diagnosed in §2.1.

Setup. We reuse the 100 ScopeSAE features already classified by GPT-4o into the four trigger types A/B/C/D in §3.4 (see Appendix I for the rubric). For each of these features, we collect its top-activating contexts and record, for every activating token, the layer $l^*(t)$ selected by ScopeSAE. We then aggregate the selected layers within each type into three depth bands: shallow (L3–L5), middle (L6–L8), and deep (L9–L11). The feature-type label comes from semantic classification of activation patterns and does not reference any layer; the selected-layer distribution comes from gradient-based attribution and does not reference feature type. Any association between the two is therefore observed rather than designed.

Aggregate distribution. Table 9 reports the share of selected layers within each type, across the three depth bands.

Table 9: Distribution of selected layers within each feature type, aggregated over the top-activating tokens of all features in that type.

Type	Shallow (L3–L5)	Middle (L6–L8)	Deep (L9–L11)
A: Word / word-family bound	55.1%	10.1%	34.8%
B: Semantic concept across words	74.7%	1.5%	23.8%
C: Contextual / discourse pattern	75.7%	22.3%	2.0%
D: Token-level / syntactic / noise	25.5%	34.4%	40.1%

Three observations stand out, and they do not line up with the conventional layer-role mapping in any straightforward way. First, Types B and C, namely semantic concepts and contextual/discourse patterns (the categories one would associate with deeper layers), are predominantly selected from

shallow layers (74.7% and 75.7% respectively). Second, Type D, namely token-level and syntactic patterns (the category one would associate with shallow layers), is the most spread across depth and skews toward middle and deep layers, with deep accounting for the largest share (40.1%). Third, Type A is bimodal: most word-family features are selected from shallow layers, but a non-negligible 34.8% are selected from deep layers, indicating that even within a single semantic category, the selected layer varies substantially across features.

We do not interpret these patterns as evidence against the conventional view of what layers encode on average. Instead, they show that *the layer most predictively relevant for a given token at a given step is not necessarily the layer in which the corresponding feature category is most strongly encoded*. The conventional mapping is an aggregate statement about layer specialisation; ScopeSAE’s selection is a per-token, per-step statement about prediction sensitivity. The two answer different questions.

Per-feature examples. We illustrate the same point at the level of individual features. For each type we present two features: one whose selected layers are concentrated in shallow depths, and one whose selected layers are concentrated in deep depths. This shows that within a single type, the selected layer is governed by token-level context rather than by category membership.

Type A — Word / word-family bound

Feature F9958: act / acting / acted

Explanation: The feature consistently activates on the word family around *act*, including tokens such as *act*, *acting*, and *acted*. Its top activations are concentrated in shallow selected states.

Layer distribution: L3: 43/43.

Top activation examples

- “*But I repeat: the Armenian side would be **acting** after the Turkish side made the first step.*” (activation 18.34, layer L3)
- “*These outside creditors are both capricious and untouchable; they **act** with impunity, lending—or withholding funds as the spirit moves them.*” (activation 17.93, layer L3)
- “*Instead of **acting** as a lobby for gun owners, the gun lobby became a means for right-wing politicians to lobby for support from the gun-owning public.*” (activation 17.70, layer L3)

Feature F224: free / free-market / free trade

Explanation: The feature activates strongly on the word *free* and closely related word forms or fragments. Unlike the previous Type A example, its top activations are selected from deep layers.

Layer distribution: L11: 64/64.

Top activation examples

- “*Strip the politics out, and the **free** marketeering component of her party will happily go along with whatever recruits women voters away from Hillary.*” (activation 24.73, layer L11)
- “*True to his **free-market**, limited-government beliefs, he does not intend to apply for welfare.*” (activation 24.69, layer L11)
- “*One of the reasons why the **free** market system is not perfect is because the firms that comprise it often make mistakes.*” (activation 24.67, layer L11)

Type B — Semantic concept across different words

Feature F1655: enforcing or imposing rules

Explanation: The feature activates on semantically related terms such as *impose*, *imposed*, *imposes*, *imposition*, and *enforce*, capturing the concept of applying rules, penalties, or authority. Most activations are selected from shallow layers.

Layer distribution: L3: 41/42; L4: 1/42.

Top activation examples

- “To invoke and **impose** penalties for violations of these Rules and to prominently post a list of all violations.” (activation 22.02, layer L3)
- “White House: We must interfere and **impose** federal power and control over the states.” (activation 20.05, layer L3)
- “But why are brands so keen to **impose** their jargon on us?” (activation 19.71, layer L3)

Feature F10524: evaluative sentiment

Explanation: The feature activates on evaluation-oriented words such as *positive*, *negative*, *favorable*, and *adverse*. Its activations are selected from deep layers, showing that some semantic concepts become most prediction-relevant later in the model.

Layer distribution: L11: 49/49.

Top activation examples

- “The **positive** response from people who have seen and shared the material has been reassuring.” (activation 17.81, layer L11)
- “The Russians evidently see a ‘**positive** outcome’ in calling into question the democratic process.” (activation 17.60, layer L11)
- “Moscow also continues to see the **negative** implications of the actions of the coalition led by the United States.” (activation 17.14, layer L11)

Type C — Contextual / discourse / relational pattern

Feature F1: comparative relations

Explanation: The feature activates on comparison markers such as *than*, *above*, *compared*, and *over*. Although this is a relational pattern rather than a single lexical item, the selected layers are shallow for these top activation examples.

Layer distribution: L3: 64/64.

Top activation examples

- “Women report higher levels of somatic symptoms **than** men with similar levels of cognitive/affective symptoms.” (activation 23.42, layer L3)
- “Both shale gas and conventional natural gas have a larger GHG **than** do coal or oil.” (activation 23.12, layer L3)
- “As a result, they say teenagers often see it as safer **than** tobacco.” (activation 22.66, layer L3)

Feature F8050: discourse transition markers

Explanation: The feature activates on discourse markers such as *but*, *however*, *yet*, and related transition tokens. Selected from a middle layer.

Layer distribution: L6: 64/64.

Top activation examples

- “*However*, it has been my impression he is airing the dirty laundry in part to gain some credibility with his new teammates.” (activation 21.65, layer L6)
- “*But* that would be a huge mistake, Republican lawmakers said Monday.” (activation 20.77, layer L6)
- “*But* he was quick to add: ‘I doubt that Pence has converted Donald Trump religiously.’ ” (activation 20.75, layer L6)

Type D — Token-level / syntactic / noise

Feature F432: honorifics and title punctuation

Explanation: The feature activates on honorifics and title-related tokens such as *Mr*, *Mrs*, *Ms*, *Dr*, and periods following titles. This token-level pattern is selected from a middle layer.

Layer distribution: L8: 64/64.

Top activation examples

- “*Mr* Kay said: ‘If the decision you make tonight is not fair, impartial and without bias, then this issue will not be resolved or fixed.’ ” (activation 23.36, layer L8)
- “*Even in his youth Mr.* ” (activation 22.84, layer L8)
- “*Mr* Jilani is imam at the mosque and owns the place of worship.” (activation 22.62, layer L8)

Feature F13148: numerical date tokens

Explanation: The feature activates on numerical tokens, especially day numbers inside date expressions. Its top activations are selected from deep layers.

Layer distribution: L10: 47/49; L11: 2/49.

Top activation examples

- “*By Michael Gakuran Apr 24, 2014.*” (activation 21.11, layer L10)
- “*Visual Studio Code ==> Vim Daniel Schmidt Blocked Unblock Follow Following Dec 5, 2017.*” (activation 20.55, layer L10)
- “*Spinning up a quick cloud instance with Digital Ocean Matt Simmons Blocked Unblock Follow Following Apr 15, 2015.*” (activation 20.30, layer L10)

Takeaway. The decoupling visible in Table 9 and made concrete by the per-feature examples above, where the same type maps to different selected layers and the same layer hosts different types, is exactly the layer-identity-vs-predictive-relevance decoupling that motivates ScopeSAE in §2.1. A fixed-layer SAE, by anchoring its modeling subspace to one depth, is forced to assume that prediction-relevant computation for every token resides at that depth; the patterns above show this assumption fails not only across coarse semantic categories but also within them. ScopeSAE’s per-token selection accommodates this variation by construction.

L Per-Feature Layer Concentration

Appendix K examined layer selection across features and tokens, where the selected layer varies with token-level context rather than tracking feature category. Here we turn to the complementary question, within a single feature: ScopeSAE selects a layer per token, which raises the question of whether a learned feature has a stable referent or one that drifts across the layers it is selected from. We measure how concentrated each feature’s selections are. For each of the 100 features classified in §3.4, we compute, over its top-activating tokens, the fraction of selections contributed by its most frequent layer. Across the 100 features, 87% of features have their selected states drawn entirely from one layer, and 95% of features have at least 80% of their selected states concentrated in one layer. Within a feature the selected layer is therefore largely stable, so most features are anchored to one layer’s representational space, even though which layer that is varies across features. Representative per-feature cases are shown in Appendix K.

M Ablation Study

ScopeSAE involves three design choices: a gradient–residual inner product as the attribution score, a per-token argmax over the normalized attribution scores as the layer selection rule, and a 1D gradient direction as the attribution subspace. Each of these admits a more elaborate alternative. We test all three: a causal-patching attribution that replaces the first-order approximation with the true change in log-probability (§M.1), a Top3-LCB selection that replaces the per-token argmax with cross-batch stability-aware selection (§M.2), and a higher-dimensional prediction subspace that extends the 1D gradient to a k -dimensional one (§M.3).

M.1 Selection Mechanism

ScopeSAE selects each token’s most prediction-relevant layer using the first-order attribution score $s_l^{(t)} = \langle \nabla_{h_l^{(t)}} \mathcal{J}^{(t)}, h_l^{(t)} \rangle$, taken along the model’s unperturbed forward pass. Two exact, ablation-based alternatives replace this first-order approximation with a measured change in the next-token log-probability. **Activation patching** ablates each layer’s *incremental* contribution to the residual stream and records the resulting change; **Full-state ablation** is more aggressive and zeros the layer’s full state. Both require $1+L$ forward passes per training step ($\sim 10\times$ more LM passes for $L = 9$, against the single extra backward pass of the first-order score). For comparison, RouteSAE [28] provides a third mechanism: a learned router optimized jointly with the SAE reconstruction loss.

Table 10 reports the layer-selection distribution of each mechanism, and Table 11 evaluates SAEs trained end to end under each selection rule.

Table 10: Layer selection distribution (%) under four selection mechanisms.

Layer	RouteSAE	Full-state ablation	Activation patching	ScopeSAE (ours)
L3	99.40	100.00	17.54	15.44
L4	0.06	0.00	9.24	11.73
L5	0.20	0.00	8.38	9.41
L6	0.01	0.00	3.64	8.94
L7	0.01	0.00	5.94	8.37
L8	0.00	0.00	9.88	8.83
L9	0.08	0.00	15.56	9.74
L10	0.24	0.00	14.00	9.73
L11	0.01	0.00	15.83	17.80

Two mechanisms collapse onto a single layer. RouteSAE’s learned router places 99.4% of its selections on L3, and Full-state ablation places 100% there: zeroing an early layer destroys all downstream computation, so the score measures the knock-on damage of deleting the layer rather than the layer’s own contribution, and since early layers do the most damage when removed, the selection collapses onto the earliest candidate. In doing so it discards the per-token variation that motivates dynamic selection, and its end-to-end DeltaCE turns positive (+0.107), falling out of the better-than-original regime entirely.

Table 11: Behavioral and reconstruction metrics for SAEs trained end to end under each selection rule, at $k=64$. Full-state ablation and Activation patching replace ScopeSAE’s first-order score with an exact ablation-based one; all other components are held fixed.

Selection rule	DeltaCE	Recovered	DictCosSim	Feat. Util.	Feat. Count
ScopeSAE (first-order)	−0.197	1.035	0.0202	99.8	5319
Activation patching	−0.117	1.018	0.0206	99.7	5040
Full-state ablation	+0.107	0.986	0.0224	97.6	3353

Activation patching does not collapse: ablating only each layer’s incremental contribution spreads its selections across all candidate layers and broadly tracks the first-order distribution. Yet when it is used to select layers and train the SAE end to end, it is consistently worse than the first-order score on the behavior-side metrics (DeltaCE −0.117 vs. −0.197; Recovered 1.018 vs. 1.035). The reason is shared with the more aggressive variant: deleting a whole layer, whether by zeroing it or by removing its increment, changes all subsequent computation, so the score reflects the deletion’s downstream effect rather than the layer’s own contribution. The first-order score, taken along the model’s unperturbed forward pass, isolates that contribution more precisely and at a fraction of the cost: one extra backward pass rather than a separate forward pass per candidate layer. First-order attribution is therefore not a lossy shortcut but the selection rule that isolates the relevant layer most precisely and most cheaply.

M.2 Selection Stability

ScopeSAE selects layers via per-token argmax over the normalized attribution scores, which can fluctuate between adjacent layers when their scores are close. A natural alternative is to stabilize selection by considering top- k candidates and choosing among them based on cross-batch statistics. We implement this as **ScopeSAE-Top3-LCB**: each token first restricts its selection to the top-3 candidate layers; the final layer is then chosen by the lower confidence bound (LCB) of an exponential moving average over per-layer attribution scores, favoring layers that are both high-scoring and stable across batches.

Table 12 compares ScopeSAE-base with ScopeSAE-Top3-LCB on six core metrics at $k=64$.

Table 12: Comparison of ScopeSAE with the Top3-LCB selection variant at $k=64$.

Method	NormMSE	KL	DeltaCE	Recovered	Feat. Count	Interp Score
ScopeSAE (ours)	0.243	495.18	−0.20	1.035	5,319	4.44
ScopeSAE-Top3-LCB	0.239	522.52	−0.005	1.002	5,429	4.31

Top3-LCB shows mixed outcomes: it improves slightly on reconstruction-side metrics (NormMSE and feature count) but degrades on behavior-side metrics (KL, DeltaCE, Recovered, and interpretability score). The gain on one side does not compensate for the loss on the other. The variant also introduces concrete additional state and computation in every training step. It maintains, for each candidate layer, exponential moving averages of both the first and second moments of the attribution score, which must be updated and used to compute per-layer standard deviations and LCB values at every step. These two per-layer accumulators are persisted via PyTorch buffers so that eval and inference can reuse them, expanding the model state and tying training and inference selection together. Layer selection itself becomes a multi-stage operation: a top-3 indexing per token, followed by a gather of the corresponding LCB scores and an argmax. A warmup phase further requires the training loop to track a step counter and switch between top-1 and LCB selection at a configured threshold. None of these additions translates into a consistent advantage, so the per-token argmax used by ScopeSAE remains the better choice.

M.3 Attribution Dimensionality

ScopeSAE’s attribution score is computed along a single gradient direction. A natural extension is to consider a higher-dimensional prediction subspace: for each token, the original CE gradient is

augmented with the logit gradients of the top- $(k-1)$ most probable next tokens, orthogonalized via Gram-Schmidt to obtain a k -dimensional subspace. We compare $k \in \{1, 2, 4\}$ by measuring the within-subspace effect on cross-entropy.

Table 13 reports the resulting effect sizes.

Table 13: Effect of increasing the prediction-subspace dimensionality k . Marginal gain is computed relative to $k=1$.

k	ΔCE (within-subspace)	Marginal gain over $k=1$
1	-0.111	—
2	-0.113	-0.002
4	-0.128	-0.017

Increasing k from 1 to 4 improves the within-subspace effect by only ~ 0.02 , an order of magnitude smaller than the main effect at $k=1$ (-0.111). The 1D gradient direction therefore already captures the bulk of the prediction-relevant signal, and additional dimensions yield only marginal returns. ScopeSAE’s choice of a 1D attribution score is well-supported by this margin.

Summary. The three ablations above target the three design choices in ScopeSAE: a first-order attribution score, per-token argmax selection, and a 1D attribution subspace. The more elaborate alternatives (exact causal patching, cross-batch stable selection, and higher-dimensional attribution) each fail to improve upon ScopeSAE, either by collapsing onto a single layer, by trading reconstruction-side gains for behavior-side losses, or by yielding only marginal returns at substantial added complexity. ScopeSAE’s design choices are jointly well-supported by these ablations.

N Generalization across Model Scales and Families

To test whether ScopeSAE’s behavioral and structural advantages transfer beyond the 1B-scale Llama model used in §3.1, we apply the same training and evaluation pipeline to three additional language models spanning different scales (0.5B–8B) and families (Llama and Qwen). All other components of §3.1 are preserved: training data (OpenWebText2, $\sim 100\text{M}$ tokens), context length (512), optimization schedule, and evaluation metrics. Two settings are adapted to each model. First, the candidate layer range follows the same $\frac{1}{4}$ – $\frac{3}{4}$ -of-depth convention used in §3.1. Second, the TopK sparsity K scales with the model’s hidden size, so that the active-feature count per token remains a comparable fraction of representational capacity across models. The four models and their per-model settings are summarized in Table 14.

Table 14: Per-model settings used in the generalization experiments.

Model	Layers	Hidden	Latent	Candidate Range	K
Llama-3.2-1B-Instruct	16	2048	16,384	L3–L11	64
Qwen2.5-0.5B-Instruct	24	896	7,168	L5–L17	32
Qwen2.5-7B-Instruct	28	3584	28,672	L6–L20	128
Llama-3.1-8B-Instruct	32	4096	32,768	L7–L23	128

Table 15: ScopeSAE results across model scales (0.5B–8B) and families (Llama and Qwen). DeltaCE remains negative and Recovered exceeds 1 on every model, confirming that the reconstruction-better-than-original effect transfers across scale and family.

Model	DeltaCE	Recovered	NormMSE	KL Div.	DictCosSim	Feat. Util.	Feat. Count
Llama-3.2-1B	-0.20	1.03	0.24	495.18	0.02	99.8%	5,319
Qwen2.5-0.5B	-1.15	1.33	0.17	850.84	0.04	99.9%	433
Qwen2.5-7B	-0.29	1.04	0.21	423.63	0.02	99.7%	13,229
Llama-3.1-8B	-0.12	1.07	0.30	200.45	0.01	97.2%	14,022

The reconstruction-better-than-original effect transfers across scale and family. We evaluate on Llama-3.2-1B-Instruct, Qwen2.5-0.5B-Instruct, Qwen2.5-7B-Instruct, and Llama-3.1-8B-Instruct. As

shown in Table 15, DeltaCE is negative and Recovered exceeds one on all four models, with Recovered ranging from 1.03 on Llama-3.2-1B to 1.33 on Qwen2.5-0.5B. The effect’s magnitude varies (smaller models such as Qwen2.5-0.5B show larger absolute improvement, consistent with their higher baseline cross-entropy leaving more headroom for the sparse bottleneck to remove prediction-orthogonal components), but its sign and direction are preserved. The behavioral phenomenon identified in §3.2 on Llama-3.2-1B is therefore a property of ScopeSAE’s prediction-relevant subspace, not of any specific scale or family.

Dictionary efficiency is preserved at every scale. Feature utilization stays above 97% on all four models, and dictionary cosine similarity stays at or below 0.04. Neither metric degrades as the model scales up: the 8B Llama still populates over 97% of its 32,768-dimensional dictionary with distinct directions, and the 7B Qwen reaches 99.7%. This shows that ScopeSAE’s high utilization and low redundancy in §3.4 are not artifacts of the 1B regime but a general property of how the prediction-relevant subspace populates the dictionary.

Feature count scales meaningfully with model size. The number of interpretable features grows from 433 on Qwen2.5-0.5B to 14,022 on Llama-3.1-8B, increasing by roughly $32\times$ as the model scales by $16\times$ in parameter count. This monotone growth indicates that ScopeSAE’s dictionary is not bottlenecked by a fixed ceiling on extractable structure; instead, larger models genuinely contain more decomposable factors that the prediction-relevant subspace can expose. Taken together, the three patterns above (preserved better-than-original effect, preserved dictionary efficiency, and growing feature count) establish that ScopeSAE’s advantages are robust across scales and families.

Robustness to training budget and corpus. The results above vary the model; we next vary the training budget and the training corpus, holding the model fixed at Llama-3.2-1B-Instruct. Beyond the 100M-token OpenWebText2 setup of §3.1, we extend training to 300M tokens and add a second corpus (FineWeb [24]), and re-evaluate ScopeSAE under each budget and corpus. As shown in Table 16, the behavioral and structural profile is preserved throughout: DeltaCE stays negative, Recovered stays above one, feature utilization stays above 99%, and dictionary cosine similarity stays at or below 0.02 across all four budget and corpus settings. The reconstruction-better-than-original effect and the dictionary-efficiency gains are therefore not tied to a particular training budget or corpus, and the 100M-token OpenWebText2 setup used in the main text is not a source of the effect.

Table 16: ScopeSAE under extended training budgets (100M and 300M tokens) and a second corpus (FineWeb), with the model fixed at Llama-3.2-1B-Instruct. The effect and the dictionary-efficiency gains hold across every budget and corpus.

Corpus	Budget	DeltaCE	Recovered	DictCosSim	Feat. Util.	Feat. Count
OpenWebText2	100M	−0.197	1.035	0.0202	99.8%	5,319
OpenWebText2	300M	−0.217	1.038	0.0197	99.9%	5,579
FineWeb	100M	−0.139	1.018	0.0201	99.7%	5,300
FineWeb	300M	−0.156	1.026	0.0196	99.9%	5,654

O Attributing the Reconstruction-Better-Than-Original Effect

Section 3.3 attributes the reconstruction-better-than-original effect to the suppression of the locally prediction-orthogonal residual at the selected layer. This appendix examines the same effect from three further angles, each consistent with that account: the prediction target used for selection, the tokens on which the reduction occurs, and the training objective of the SAE being selected into.

Selection with the model’s own predicted token. The main text selects each token’s layer using the next token x_{t+1} as the prediction target (§3.2). We repeat the full ScopeSAE pipeline with a single change: the attribution score is computed on the model’s own top-1 predicted token $\hat{y}_{t+1}$ instead. Layer selection, SAE training, and evaluation are otherwise identical, with the SAE retrained on the newly selected states. As Table 17 shows, the two targets yield the same behavioral and feature-level profile: DeltaCE stays negative and the interpretable feature count, utilization, and redundancy are on par. The effect therefore persists when the layer is selected from the model’s own predicted token, and does not hinge on the next token being available at selection time.

Table 17: ScopeSAE under next-token and predicted-token layer selection. Entries are mean $\pm$ std over three training seeds.

Selection target	DeltaCE	Recovered	DictCosSim	Util(%)	#Feat
Next token	-0.197 ± 0.009	1.035 ± 0.0052	0.0202 ± 0.00007	99.8 ± 0.08	5319 ± 46
Predicted token	-0.206 ± 0.011	1.038 ± 0.0071	0.0205 ± 0.00008	99.8 ± 0.08	5281 ± 53

The reduction concentrates on predictions the model already supports. To characterize where the cross-entropy change occurs, we split the evaluation by whether the base model’s own top-1 prediction matches the next token, $\hat{y} = y$ versus $\hat{y} \neq y$, under both selection targets (Table 18). The reduction is concentrated on the $\hat{y} = y$ tokens, where the selected state aligns with a prediction the model already makes, and reverses to a small increase on the $\hat{y} \neq y$ tokens. This is the token-level signature of the mechanism in §3.3: suppressing the prediction-orthogonal residual sharpens the predictions the model’s own computation already supports, and does not manufacture predictions it does not. The pattern is the same under both targets.

Table 18: DeltaCE on tokens where the base model predicts correctly ($\hat{y} = y$) and incorrectly ($\hat{y} \neq y$), under next-token and predicted-token selection.

Subset	Next-token selection	Predicted-token selection
$\hat{y} = y$	-0.438	-0.469
$\hat{y} \neq y$	$+0.144$	$+0.167$

The effect requires the reconstruction objective, not selection alone. Selection places the SAE at a prediction-relevant layer, but the reduction is produced by what the SAE is trained to represent there. We isolate the two by training one standard fixed-layer TopK SAE at each candidate layer, on the same model and data as ScopeSAE, and then applying the ScopeSAE selection rule to route each token to its selected layer’s own SAE. These per-layer SAEs are trained independently for this analysis. On their own, no fixed layer produces a cross-entropy reduction: DeltaCE is positive at every layer we train, ranging from $+0.070$ at L3 to $+0.389$ at L15 (Table 19), and no single layer approaches ScopeSAE’s feature-level profile. We train these SAEs not only at the candidate layers L3–L11 but also below and above them (L0–L2 and L12–L15); the reduction appears at none of these depths, and the baseline only gets worse outside the candidate range. This confirms that the main-text comparison is not sensitive to the fixed layer chosen for the single-layer baselines, and that the L3–L11 candidate range is not a window chosen to favor our method: no layer at any depth matches ScopeSAE on cross-entropy or on its feature-level profile.

Table 19: Each independently trained per-layer TopK SAE, evaluated at its own layer, across the full depth range.

Layer	DeltaCE	Recovered	DictCosSim	Util(%)	#Feat
L0	$+0.078$	0.986	0.0223	98.1	2815
L1	$+0.099$	0.987	0.0225	97.5	3013
L2	$+0.075$	0.983	0.0224	97.8	2904
L3	$+0.070$	0.989	0.0223	97.6	3074
L4	$+0.076$	0.988	0.0226	96.5	3108
L5	$+0.083$	0.986	0.0227	95.1	3093
L6	$+0.096$	0.984	0.0222	98.3	3135
L7	$+0.104$	0.984	0.0219	98.2	2871
L8	$+0.111$	0.982	0.0217	95.9	3039
L9	$+0.108$	0.983	0.0212	97.8	2837
L10	$+0.126$	0.979	0.0218	97.2	2973
L11	$+0.114$	0.981	0.0214	97.0	2916
L12	$+0.146$	0.977	0.0224	96.7	2848
L13	$+0.176$	0.974	0.0222	96.8	2861
L14	$+0.175$	0.973	0.0219	96.4	2902
L15	$+0.389$	0.945	0.0219	94.2	2675

Routing the selection rule over these per-layer SAEs likewise leaves DeltaCE positive (+0.081) and the interpretable feature count at the level of a standard multi-layer combination (Table 20). The reduction appears only when selection and ScopeSAE’s reconstruction objective act together: selection determines which layer is reconstructed, while the objective is what suppresses the prediction-orthogonal residual at that layer. This is consistent with §3.3: the effect follows from what the SAE is trained to represent at the selected layer, not from the selection step in isolation.

Table 20: The selection rule applied over independently trained per-layer TopK SAEs, compared against the best single such SAE and against ScopeSAE.

Method	DeltaCE	Recovered	DictCosSim	Util(%)	#Feat
Best single-layer TopK SAE	+0.070	0.989	0.0223	97.6	3074
Selection rule + per-layer TopK SAEs	+0.081	0.987	0.0217	95.9	2931
ScopeSAE	−0.197	1.035	0.0202	99.8	5319

P Limitations and Future Work

We discuss two limitations of ScopeSAE, concerning feature steering and feature interpretation respectively.

Limitation 1: Finite steering strength range. We acknowledge a concrete limitation of ScopeSAE’s feature steering: like any activation-level intervention, it operates within a finite strength range, and applying larger intervention magnitudes introduces side effects that ultimately overwhelm the intended directional signal. The underlying reason is intrinsic to activation-level interventions: clamping adds a scaled vector to the model’s hidden state, and as the scale grows, the resulting activation drifts further from the distribution the model encountered during training, eventually pushing the model out of the regime in which its computation remains coherent. To quantify this limitation, we compare side-effect metrics at the strength used in the main text (10) against a doubled strength (20):

Table 21: Full ScopeSAE side-effect metrics at intervention strengths 10 and 20 across the four intervention quadrants.

Strength	Condition	PPL ratio	Entropy change	Cosine sim.	Top-1 preservation
10	Helpful (Amplify)	1.38	0.65	0.91	0.72
10	Helpful (Suppress)	1.39	0.67	0.91	0.72
10	Harmful (Amplify)	1.34	0.58	0.92	0.72
10	Harmful (Suppress)	1.37	0.71	0.91	0.72
20	Helpful (Amplify)	2.71	1.09	0.81	0.60
20	Helpful (Suppress)	2.17	1.15	0.80	0.61
20	Harmful (Amplify)	2.25	0.96	0.84	0.61
20	Harmful (Suppress)	2.08	1.32	0.82	0.60

At strength 20, all side-effect metrics degrade substantially: the PPL ratio rises from roughly 1.4 to 2.3, cosine similarity between original and clamped hidden states drops below 0.85, and top-1 token preservation falls below 0.62. Inspection of generated outputs at strength 20 shows frequent incoherent or degenerate completions. As shown in Table 22, directional win rates at this strength (17.5, 2.5, 7.1, 18.6 across the four intervention quadrants) preserve the qualitative ordering observed at strength 10 but collapse in absolute magnitude, because degraded outputs are rarely judged better than the original regardless of intervention direction. In practical terms, this means that scaling up intervention strength is not a general-purpose way to achieve larger behavioral shifts: once the activation drifts far enough from the natural distribution, directional information is dominated by overall output degradation.

Table 22: ScopeSAE directional win rates at strengths 10 and 20.

Condition	Strength 10		Strength 20	
	Amplify	Suppress	Amplify	Suppress
Helpful	62.5	35.0	17.5	2.5
Harmful	35.7	62.9	7.1	18.6

This constraint is a real limitation of the method. It bounds the magnitude of behavioral change that feature steering can reliably induce, and it means that the effective strength range is a property of each method-model pairing rather than a fixed constant that transfers across settings. That said, the limitation is shared across activation-level intervention methods. High-strength clamping pushes the model out of its normal operating regime across all SAE methods we evaluate, and Table 8 shows that ScopeSAE degrades least under the same doubling, though it degrades nonetheless. A promising direction for future work is to relax this bound by decoupling directional effect from overall activation magnitude: for instance, interventions that steer along feature directions while projecting out the components that push the activation off-distribution, or training objectives that encourage features whose steering effect saturates gracefully rather than collapsing at high strengths. In the present paper, we are careful to report all intervention results within the empirically validated operational regime (strength 10), where side-effect metrics (top-1 preservation above 0.72, cosine similarity above 0.91) confirm that directional signals can be cleanly separated from overall output degradation.

Limitation 2: Per-feature referent. A second point concerns interpretation rather than intervention. ScopeSAE’s interpretive contract differs from that of a standard fixed-layer SAE. In a fixed-layer SAE, every feature shares one dictionary-wide referent, the representational space of the single layer on which the SAE is trained. In ScopeSAE, each feature is instead anchored to the layer at which it is predominantly selected (most features satisfy this in practice; Appendix L), so the referent is defined per feature rather than once for the whole dictionary. We note this difference so that it is explicit rather than left for the reader to infer.

NeurIPS Paper Checklist

1. Claims

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

Answer: [\[Yes\]](#)

Justification: The abstract and introduction accurately summarize ScopeSAE’s motivation, method, and main empirical findings. The stated claims are supported by the experiments and analyses presented in the paper.

Guidelines:

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

2. Limitations

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

Answer: [\[Yes\]](#)

Justification: The paper discusses limitations in Appendix P, including the finite effective strength range of feature steering, the side effects caused by high-magnitude interventions, and the need to keep intervention analyses within an empirically validated operational regime.

Guidelines:

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

3. Theory assumptions and proofs

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

Answer: [N/A]

Justification: The paper does not include any theoretical results.

Guidelines:

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

4. Experimental result reproducibility

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

Answer: [Yes]

Justification: The paper describes the experimental setup, model and data choices, training configuration, evaluation metrics, and main analysis protocols in Section 3.1, with additional implementation details in Appendix B.

Guidelines:

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

5. Open access to data and code

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

Answer: [Yes]

Justification: The source code is available via the link provided in the abstract, and the paper specifies the datasets and preprocessing steps used for the experiments in Section 3.1.

Guidelines:

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

6. Experimental setting/details

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

Answer: [Yes]

Justification: The paper specifies the relevant data splits, model settings, training hyperparameters, optimizer choices, and evaluation details in Section 3.1 and Appendix B.

Guidelines:

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

7. Experiment statistical significance

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

Answer: [No]

Justification: We do not report formal error bars, but the central claims are supported by robustness analyses across sparsity levels (Figures 2 and 3), candidate layers (Appendix F), intervention magnitudes (Appendix G), and four model scales and families (Appendix N), following standard reporting practice in prior SAE work.

Guidelines:

- The answer [N/A] means that the paper does not include experiments.
- The authors should answer [Yes] if the results are accompanied by error bars, confidence intervals, or statistical significance tests, at least for the experiments that support the main claims of the paper.

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

8. Experiments compute resources

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

Answer: [Yes]

Justification: Appendix C reports compute resources for the main ScopeSAE training run, including GPU type and memory, wall-clock time, training steps, throughput, and training-token count. Baselines, ablations, and generalization experiments use the same single-GPU setup with comparable per-run cost.

Guidelines:

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

9. Code of ethics

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

Answer: [Yes]

Justification: We have reviewed the NeurIPS Code of Ethics and, to the best of our knowledge, the research conforms to it in all respects.

Guidelines:

- The answer [N/A] means that the authors have not reviewed the NeurIPS Code of Ethics.
- If the authors answer [No], they should explain the special circumstances that require a deviation from the Code of Ethics.
- The authors should make sure to preserve anonymity (e.g., if there is a special consideration due to laws or regulations in their jurisdiction).

10. Broader impacts

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

Answer: [Yes]

Justification: As mechanistic interpretability work, ScopeSAE contributes positively to understanding and auditing the internal computation of large language models, supporting transparency and safety analysis. We do not foresee direct negative societal impacts, as the method operates on already-released open-source models and does not enable new harmful applications.

Guidelines:

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

11. Safeguards

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

Answer: [N/A]

Justification: The paper does not release new high-risk models or scraped datasets; it uses existing resources for mechanistic interpretability experiments and releases only the research code.

Guidelines:

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

12. Licenses for existing assets

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

Answer: [Yes]

Justification: The creators of the datasets, models, and existing code used in the paper are properly credited, and we have checked and followed the corresponding licenses and terms of use.

Guidelines:

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

13. New assets

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

Answer: [Yes]

Justification: The newly released source code is documented in the accompanying repository, including instructions for reproducing the main experiments.

Guidelines:

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

14. Crowdsourcing and research with human subjects

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

Answer: [N/A]

Justification: The paper does not involve crowdsourcing or research with human subjects. The human annotation used in Appendix I.2 to validate GPT-4o's interpretability scores is a small-scale internal verification and does not constitute crowdsourced or human-subjects research.

Guidelines:

- The answer [N/A] means that the paper does not involve crowdsourcing nor research with human subjects.
- Including this information in the supplemental material is fine, but if the main contribution of the paper involves human subjects, then as much detail as possible should be included in the main paper.
- According to the NeurIPS Code of Ethics, workers involved in data collection, curation, or other labor should be paid at least the minimum wage in the country of the data collector.

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

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

Answer: [N/A]

Justification: The paper does not involve crowdsourcing or research with human subjects, so IRB approval is not applicable.

Guidelines:

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

16. Declaration of LLM usage

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

Answer: [Yes]

Justification: The paper uses GPT-4o as an auxiliary evaluation tool for interpretability scoring (Section 3.4, Appendix I) and for pairwise response comparison in the causal intervention experiments (Section 3.5). Its interpretability ratings are validated against human annotations (Appendix I.2).

Guidelines:

- The answer [N/A] means that the core method development in this research does not involve LLMs as any important, original, or non-standard components.
- Please refer to our LLM policy in the NeurIPS handbook for what should or should not be described.