

FROM CONCEPT ALIGNMENT TO CAUSAL GROUNDING: AN INTERVENTION TEST OF CHAIN-OF-THOUGHT FAITHFULNESS

Qianli Wang^{1,5} Yilong Wang^{1,†} Dennis Wei^{2,†} Jingyi Sun³
 Simon Ostermann^{5,6,7} Pepa Atanasova^{3,‡} Nils Feldhus^{4,5,‡}

¹Technische Universität Berlin ²IBM Research ³University of Copenhagen
⁴University of Groningen ⁵German Research Center for Artificial Intelligence (DFKI)
⁶Saarland Informatics Campus ⁷Centre for European Research in Trusted AI (CERTAIN)

Correspondence: qianli.wang@tu-berlin.de

[†]Equal contribution. [‡]These authors contributed equally to this work as joint last author.

ABSTRACT

Chain-of-thought (CoT) can sound plausible yet be unfaithful to the model’s underlying reasoning. Most prior work probes CoT faithfulness through input–output behavior or input attributions, leaving internal computation largely unexplored. We instead cast faithfulness as internal concept grounding: *Does a large language model’s (LLM) CoT reasoning engage the same internal concepts that support the LLM’s direct prediction, and do the shared concepts causally drive its answer?* Encoding a prediction pass and a CoT pass with a single shared sparse autoencoder (SAE), a reliable approximator of the latent concepts LLMs use, makes their internal concepts directly comparable. We introduce three correlational metrics of concept-level alignment and a causal metric, Δp , which ablates the shared concepts and measures the drop in answer probability. Across five LLMs and four datasets, concept alignment is generally high, as indicated by the correlational metrics; yet these only identify which concepts are shared, not how much they causally contribute. Δp fills this gap: causal faithfulness varies substantially with model depth, peaking at mid-to-late layers rather than the final ones, and model scale reshapes the layer-wise profile. Moreover, causally important shared concepts are not always verbalized in the CoT. These dissociations suggest that faithfulness cannot be reliably assessed from surface-level or representational correspondence alone; assessing it requires causal tests of whether the internal concepts underlying a CoT actually drive the model’s prediction.

1 INTRODUCTION

Chain-of-thought (CoT) is a promising technique that seemingly exposes an LLM’s decision-making on various critical tasks (Wei et al., 2022; Kojima et al., 2022), such as safety monitoring (Korbak et al., 2025). This apparent transparency, however, hinges on CoT faithfulness, i.e., the degree to which it accurately represents or contributes to the underlying reasoning of the model, rather than a plausible post-hoc rationalization (Jacovi & Goldberg, 2020; Agarwal et al., 2024). The distinction matters most in high-stakes settings, where unfaithful CoTs can lull users into over-relying on systems whose stated reasoning masks their true decision basis (Paul et al., 2024).

Prior work on measuring CoT faithfulness falls along complementary axes. One line of work intervenes on the reasoning chain, or the parametric knowledge that the reasoning invokes, and treats the CoT as faithful only if the answer changes accordingly (Lanham et al., 2023; Zaman & Srivastava, 2025; Tutek et al., 2025); another injects answer hints and treats the CoT as unfaithful if a hint is used silently without acknowledgment (Turpin et al., 2023; Chen et al., 2025; Chua & Evans, 2025). Both infer internal alignment only indirectly, from input–output behavior. Attribution-based




Figure 1: SAE-based CoT faithfulness evaluation pipeline. We apply a shared SAE to the residual stream, extracting sparse activations from prediction and CoT-derived prediction passes. The two activations are then compared using two complementary families of faithfulness indicators: correlational metrics (CC-SAE, Jaccard similarity, and Recall), which quantify the overlap and alignment between the prediction and CoT activations; and a causal metric, which ablates the shared features from the reconstructed prediction stream and measures the resulting change in answer probability.

methods (Parcalabescu & Frank, 2024; Admoni et al., 2026) assess the alignment between input attributions for prediction and those for CoT, but they capture only dependence on the input and miss what the model computes internally. These evaluations largely treat the model as a black box, *without testing which internal computations causally drive the answer*.

A more direct line pries open this black box: it uses mechanistic interpretability (MI) to inspect internal computation as signals for CoT faithfulness (Occhipinti et al., 2026; Shen et al., 2026). Closest to our work, Chen et al. (2026) train two SAEs and patch CoT features into a no-CoT run, reading the gain in answer probability to test whether reasoning concepts can steer a direct answer. However, these approaches do not examine (1) whether and to what extent an LLM’s CoT engages the same internal concepts that also support the LLM’s direct prediction, (2) whether these shared concepts causally drive the prediction, and (3) how these relations are distributed across model depth.

We fill this gap by casting CoT faithfulness as internal concept grounding: *does an LLM’s CoT engage the same internal concepts that also support the LLM’s direct prediction, and do the shared concepts causally drive its answer?* Concretely, we encode the answer-position residuals of a prediction pass and a CoT-derived prediction pass with a single shared SAE, making cross-pass concept overlap well-defined (Figure 1). On this shared concept space, we propose two notions of faithfulness: (i) **correlational metrics** (CC-SAE, Jaccard similarity, and prediction-grounded recall) quantifying alignment in concept usage, as a form of *self-consistency*¹ (Parcalabescu & Frank, 2024; Admoni et al., 2026); and (ii) a **causal metric** Δp that ablates the shared concepts $S_\cap$ from the prediction or CoT pass and measures the drop in answer probability, as a measure of *faithfulness*.

Applying this framework to five LLMs and four benchmarks yields six main findings. **First**, correlational metrics suggest decent concept alignment overall, which provides a solid basis for using Δp . Yet such metrics only identify which features are shared across passes. Precisely quantifying their contribution requires causal metrics like Δp . **Second**, we show that Δp captures genuine causal grounding rather than artifacts of ablation count, feature selection, or the SAE itself, through random and norm-matched controls and SAE-configuration ablations, respectively.

¹*Self-consistency* measures the alignment between a model’s contribution to answer prediction and its generated CoT (Parcalabescu & Frank, 2024), which is necessary but not sufficient for *faithfulness*.

Third, our notion of faithfulness is not a single number: it varies substantially across depth, generally peaking at mid-to-late rather than final layers, with model-specific dips, underscoring the need for layer-aware reporting. **Fourth**, architectures differ qualitatively rather than merely in magnitude: Qwen3 models exhibit *distributed* faithfulness that remains high across most layers, whereas Llama-3.1-8B shows *localized* faithfulness concentrated in a few layers. **Fifth**, within a model family, larger models are more faithful overall, and scaling reshapes the layer-wise profile, concentrating low-faithfulness computation into fewer, earlier layers while stabilizing alignment across the remaining depth. **Lastly**, a case study of layer-wise concept transition described in natural language reveals that causally important shared concepts are not always verbalized in the CoT.

2 RELATED WORK

CoT Faithfulness Evaluation. One line of work intervenes on input or CoT and examines whether the intervention is reflected in the answer (Lanham et al., 2023; Zaman & Srivastava, 2025; Ye et al., 2026). The second line of work injects answer hints and verifies whether they are verbalized in CoT (Turpin et al., 2023; Arcuschin et al., 2025; Chua & Evans, 2025). Furthermore, Tutek et al. (2025); Sun et al. (2026) unlearn knowledge used in CoT steps from model parameters and assess whether the answer changes after unlearning. Paul et al. (2024) apply causal mediation and counterfactual reasoning-chain interventions to test whether intermediate CoT steps influence final answers. Parcalabescu & Frank (2024) measure the alignment between input contribution to answer prediction and CoT reasoning. Our work is close to theirs but measures alignment and grounding in the SAE concept space (§3.2), which addresses limitations of input-level alignment: SAE concepts match the granularity of internal computation and are near-monosemantic where tokens are not.

Mechanistic Interpretability for CoT Faithfulness Evaluation. Recent advancements in MI enable the inspection of internal representations to assess faithfulness directly. Occhipinti et al. (2026) train a linear probe to distinguish faithful and unfaithful CoTs. Shen et al. (2026) apply circuit tracing and a supervised detector to measure CoT unfaithfulness as the discrepancy between sentence-level circuits and the displayed CoT. Chen et al. (2026), the closest to us, train *two separate* SAEs and patch CoT features into a non-CoT run, reading the causal effect from the gain in answer probability, which is a magnitude-based test of whether injected reasoning features can steer a direct answer. Our intervention runs in the opposite direction (Figure 6): with a *single shared* SAE that makes cross-pass concept overlap well-defined, we ablate the shared concepts $S_{\cap}$ between answer prediction and CoT (§3.5), yielding a test of whether the direct answer depends on, not merely is steerable by, the concepts that the CoT instates.

3 METHODOLOGY

3.1 PROBLEM SETUP

Inspired by Parcalabescu & Frank (2024), who propose a faithfulness metric that compares input attributions for the answer and for CoT, we instead take one step further into the model internals. Sparse autoencoders (SAEs) have been shown to recover latent concepts reliably (§3.2) that can influence LLM outputs when manipulated (Bricken et al., 2023; Templeton et al., 2026). We measure CoT faithfulness by checking whether *the LLM’s CoT-based prediction relies on the same internal representations (SAE concepts) used in a direct prediction*. Specifically, given an input x , we compare model behavior under two conditions: ① **Prediction Pass:** The model receives the input question x and is asked to predict the answer $\hat{y}$. We extract the residual stream activation at the answer token position, denoted $\mathbf{r}_{\text{pred}} \in \mathbb{R}^{d_{\text{model}}}$, where d_{model} denotes the model’s hidden dimension (Table 3). ② **CoT Pass:** Following prior work (Parcalabescu & Frank, 2024; Tutek et al., 2025), we adopt a two-step procedure: the model first generates a reasoning chain c based on input x (Kojima et al., 2022). We then form the sequence $(x; c)$, run a forward pass, and extract the residual stream activation $\mathbf{r}_{\text{cot}} \in \mathbb{R}^{d_{\text{model}}}$ at the answer position (§3.3), following Dura et al. (2026); Chen et al. (2026).

3.2 SPARSE AUTOENCODERS

To identify which internal representations are genuinely used by the model during prediction and during CoT, we apply sparse autoencoders (SAEs), which reconstruct residual-stream activations

through a sparse, overcomplete bottleneck (Cunningham et al., 2023; Zhang et al., 2026) (Figure 1). Given a residual stream vector $\mathbf{r} \in \mathbb{R}^{d_{\text{model}}}$, the SAE encoder produces a sparse activation vector $\mathbf{a}$:

$$\mathbf{a} = \sigma(\mathbf{W}_{\text{enc}}\mathbf{r} + \mathbf{b}_{\text{enc}}) \in \mathbb{R}^{d_{\text{sae}}} \quad (1)$$

where σ denotes the SAE activation function and d_{sae} represents the width of the SAE (typically $d_{\text{sae}} \gg d_{\text{model}}$; an ablation study of the SAE setup is performed in §5.4). The decoder reconstructs:

$$\hat{\mathbf{r}} = \mathbf{W}_{\text{dec}}\mathbf{a} + \mathbf{b}_{\text{dec}} \quad (2)$$

Each dimension of $\mathbf{a}$ ideally corresponds to a learned (SAE) concept; sparsity pushes these concepts toward monosemanticity rather than superposition (Bricken et al., 2023; Shu et al., 2025).

Why SAEs? SAE concepts instantiate the three desiderata (§1): encoding both passes with a single shared SAE makes concept usage directly *comparable*; sparsity promotes *monosemanticity*, so overlap is meaningful rather than an artifact of entangled directions; and individual features can be selectively ablated, making concepts *intervenable* for causal tests (§3.5).

3.3 ACTIVATION EXTRACTION

In both passes, we encode the residual at the answer position, ensuring that the same functional position is compared across the prediction and CoT passes:

$$\mathbf{a}_{\text{pred}} = \text{SAE}_{\text{enc}}(\mathbf{r}_{\text{pred}}^{(\text{ans})}) \in \mathbb{R}^{d_{\text{sae}}} \quad (3) \quad \mathbf{a}_{\text{cot}} = \text{SAE}_{\text{enc}}(\mathbf{r}_{\text{cot}}^{(\text{ans})}) \in \mathbb{R}^{d_{\text{sae}}} \quad (4)$$

Following Chen et al. (2026); Dura et al. (2026), we extract the residual $\mathbf{r}_{\text{cot}}^{(\text{ans})}$ at the final answer position rather than aggregating across CoT tokens, which would be naturally noisy: by the point of answer commitment, relevant concepts surfaced along the reasoning path are already aggregated into this position. This mirrors the prediction pass, so $\mathbf{a}_{\text{pred}}$ and $\mathbf{a}_{\text{cot}}$ are extracted at the same functional (answer-commitment) position. The intuition is that if the CoT is faithfully grounded in the model’s prediction process, it should engage answer-relevant concepts.²

3.4 CORRELATIONAL METRICS

Given SAE activation vectors $\mathbf{a}_{\text{pred}}$ and $\mathbf{a}_{\text{cot}}$, we first define three correlational metrics that quantify the overlap and alignment in concept usage between the two passes, as a form of concept-level *self-consistency* (Parcalabescu & Frank, 2024; Admoni et al., 2026):

CC-SAE. CC-SAE measures the strength-weighted alignment of activations in concept space:

$$\text{CC-SAE} = \mathbf{a}_{\text{pred}} \cdot \mathbf{a}_{\text{cot}} / \|\mathbf{a}_{\text{pred}}\| \cdot \|\mathbf{a}_{\text{cot}}\| \quad (5)$$

Jaccard Similarity. A set-based measure ignoring activation magnitudes, where $S_{\text{pred}} = \{i : \mathbf{a}_{\text{pred},i} > 0\}$ and $S_{\text{cot}} = \{i : \mathbf{a}_{\text{cot},i} > 0\}$ are the sets of active concept indices:

$$\text{Jaccard} = |S_{\text{pred}} \cap S_{\text{cot}}| / |S_{\text{pred}} \cup S_{\text{cot}}| \quad (6)$$

Prediction-Grounded Recall. A directional measure of how much of the prediction’s concept usage appears in the CoT:

$$\text{Recall} = \sum_{i \in S_{\text{pred}}} \mathbf{a}_{\text{cot},i} \cdot \mathbb{I}[\mathbf{a}_{\text{cot},i} > 0] / \sum_{i \in S_{\text{pred}}} \mathbf{a}_{\text{pred},i} \quad (7)$$

3.5 CAUSAL METRIC

The correlational metrics above measure whether the same concepts appear in both passes, but not whether those concepts causally drive the answer. To address this, we further introduce a causal intervention. Let $S_{\cap} = S_{\text{pred}} \cap S_{\text{cot}}$ be the set of shared active concepts. We ablate these concepts from the prediction pass³ by subtracting their contribution from the residual:

$$\mathbf{r}'_{\text{pred}} = \mathbf{r}_{\text{pred}} - \mathbf{W}_{\text{dec}}[\mathbf{a}_{\text{pred}} \odot \mathbf{m}_{S_{\cap}}] \quad (8)$$

²Following Tutek et al. (2025), we evaluate faithfulness in cases where CoT and no-CoT predictions of models agree. This avoids confounding the CoT-induced performance gain with the CoT faithfulness evaluation.

³Ablation results on the CoT pass are presented and discussed in §5.2.2 and Appendix F.2.




Figure 2: CoT faithfulness evaluation across layers for Qwen and Llama on LogiQA (Gemma results in Figure 10). **Left column:** Causal faithfulness (Δp) measures the probability drop when ablating shared features. **Middle column:** Correlational metrics (CC-SAE, Jaccard, Recall) between prediction and CoT activations. **Right column:** Spearman correlation heatmap between each correlational metric and Δp across layers. Shaded regions indicate 95% confidence intervals (left and middle columns). Additional results on the *remaining datasets* are provided in Appendix C.4.

where $\mathbf{m}_{S_{\cap}}$ is a binary mask that selects only the shared concepts. We then measure the change in answer probability, as a measure of *faithfulness*:

$$\Delta p = P(\hat{y} \mid \mathbf{r}_{\text{pred}}) - P(\hat{y} \mid \mathbf{r}'_{\text{pred}}) \quad (9)$$

$\Delta p \gg 0$ means that ablation substantially lowers answer probability, so these concepts *causally drive* the prediction, suggesting a more faithful CoT. Conversely, $\Delta p \approx 0$ means that ablating shared concepts leaves the answer intact: the CoT concepts are *not* causally responsible, so the model reaches the answer by other means, suggesting a highly unfaithful CoT and a potential concept-level signature of post-hoc rationalization (such examples with low Δp are presented in Appendix H).

4 EXPERIMENTAL SETUP

Models. We employ five open-source LLMs with varying sizes: Llama-3.1-8B (Grattafiori et al., 2024), Gemma-2- $\{2\text{B}, 9\text{B}\}$ (Team et al., 2024) and Qwen3- $\{1.7\text{B}, 8\text{B}\}$ (Yang et al., 2025), with off-the-shelf Llama-Scope SAEs (He et al., 2024), Gemma-Scope SAEs (Lieberum et al., 2024) and Qwen-Scope SAEs (Deng et al., 2026), respectively, trained on the residual streams.

Datasets. We adopt four diverse datasets, including one math and three multiple-choice question-answering (MCQA) datasets: GSM8K (Cobbe et al., 2021), LogiQA (Liu et al., 2021), OpenbookQA (Mihaylov et al., 2018), and ARC-Easy (Clark et al., 2018).

5 RESULTS

In this section, we analyze the correlational metrics and examine the relationship between correlational and causal metrics (§5.1), validate the causal metric Δp under controlled conditions (§5.2), analyze causal layer-wise faithfulness (§5.3), ablate the SAE configuration (§5.4), and characterize SAE concept transitions across layers through natural-language descriptions (§5.5).

5.1 CORRELATIONAL METRICS AND THEIR RELATIONSHIP TO THE CAUSAL METRIC

We first observe that the correlational metrics (§3.4) achieve reasonably high magnitudes (CC-SAE > Recall > Jaccard) (Figure 2; middle column). This indicates decent conceptual alignment between prediction and CoT passes, i.e., *self-consistency*, which provides a solid basis for using Δp (§5.3).




Figure 3: Three base causal ablation conditions and top- k upper bound comparison against $\Delta p(S_\cap)$ on the prediction pass across models and datasets. Full raw results are provided in Table 10. The **Condition D** results can be found in Appendix F.1.3.

We further explore the correlational metrics with respect to how and to what extent they correlate with the causal metric Δp (§3.5). Figure 2 (right column) illustrates that the correlational metrics align more closely with Δp in late layers (except for Llama-3.1-8B), while in early layers the correlation can even become negative. Additionally, Table 15 reveals that CC-SAE has a larger magnitude and correlates more strongly with Δp , outperforming Recall, followed by Jaccard. Overall, all three correlations are negligible to weak, with substantial variance across layers.

Table 16 provides more evidence of weak correlation by binarizing both CC-SAE (θ^*) and Δp (τ^*) using per-dataset median thresholds. The alignment-causal-effect association is weak (A : above-median CC-SAE; B : above-median Δp): $\bar{P}(A|B) = 0.53$, $\bar{P}(B|A) = 0.54$, with lift = 1.08. High alignment thus raises the probability of a high causal effect by only 8% over baseline: concept overlap identifies *which* features are shared, but high overlap (*self-consistency*) neither guarantees nor precludes a strong causal effect (*faithfulness*), consistent with Parcalabescu & Frank (2024). Taken together, *quantifying the contribution of overlapping features therefore requires causal metrics*.

5.2 VALIDATION OF THE CAUSAL ROLE OF INTERSECTION FEATURES

5.2.1 PREDICTION PASS

To validate the strength of the causal effect of intersection features $S_\cap$ with size k , we evaluate four base conditions for sanity checks (the Venn diagram of S_{pred} and S_{cot} is illustrated in Figure 7):

- **Condition A:** Sample k random features from $(S_{\text{pred}} \cup S_{\text{cot}}) \setminus S_\cap$. It aims to test whether intersection features influence the answer probability more than other features that are active in either pass.
- **Condition B:** Sample k random features from $S_{\text{pred}} \setminus S_\cap$. It aims to test whether the causal effect is specific to features reused by CoT.
- **Condition C:** Norm-matched sampling from $S_{\text{pred}} \setminus S_\cap$, controlling for the L_2 norm of the ablation vector (Appendix F.1.2). It aims to test whether the effects are not driven by ablation magnitude.
- **Condition D:** Ablate $S_\cap$ together with k additional random features from $S_{\text{pred}} \setminus S_\cap$. If $S_\cap$ captures the causally important features, adding prediction-only features should not noticeably increase Δp beyond ablating $S_\cap$ alone, which serves as a completeness check.

Additionally, we define and compute a **top- k upper bound** (magnitude-ranked baseline), i.e., ablating the k features in S_{pred} with the largest activation magnitude.

Results. Figure 3 and Table 10 show that the intersection ablation consistently yields close to the largest Δp , while random baselines A–C yield Δp close to 0. This reveals two findings. First, the causal effect is specific to intersection features rather than to arbitrary subsets of S_{pred} or features active in one pass alone. Second, causal importance is as driven by conceptual reuse and grounding as it is by activation magnitude: the top- k upper bound, which selects the k highest-magnitude features in S_{pred} , is only marginally above $\Delta p(S_\cap)$, so intersection features are nearly as causally potent as the best magnitude-selected subset of the same size, despite not being chosen for magnitude. For **Condition D**, Table 11 illustrates a negligible incremental effect across all model-dataset combinations, indicating that $S_\cap$ alone is complete for the causal effect. Features that are active only in the prediction pass but are not reused from CoT contribute minimally to the model’s final answer.

5.2.2 CoT PASS

To establish that $S_{\cap}$ reflects genuine *reasoning-relevant* computation, we further ablate these features on the **CoT pass** and examine their causal effect from two angles (Appendix F.2). First, we test cross-path consistency in terms of the causal effect of $S_{\cap}$. Second, analogous to the prediction pass (§5.2), we run four base conditions for sanity (§F.2.1).

Cross-pass Consistency. For cross-pass consistency, we compare the causal effect of $S_{\cap}$ across the two inference paths. Comparing $\Delta p(S_{\cap})$ across paths, we observe similarly high values on both prediction and CoT paths (Table 14). This rules out the interpretation that intersection features are shallow prediction heuristics that become irrelevant when explicit reasoning is available; rather, they remain causally central to both inference modes.

Control Conditions. Table 13 shows that **Conditions A–C** yield Δp substantially lower than $S_{\cap}$. The top- k upper bound remains marginally above $\Delta p(S_{\cap})$, confirming that intersection features, selected by conceptual reuse rather than magnitude, are nearly as causally potent as the best magnitude-selected subset. Furthermore, **Condition D** yields Δp nearly identical to ablating all of S_{cot} (Table 14), indicating that intersection features account for essentially all causal contributions of CoT-active features to the reasoning process. In other words, CoT-only features provide no additional effect on the CoT pass. All three findings mirror those observed in the prediction pass (§5.2.1).

5.3 LAYER-WISE CAUSAL CoT FAITHFULNESS

Prior faithfulness metrics (Lanham et al., 2023; Turpin et al., 2023; Ye et al., 2026) typically treat faithfulness as model-level, reporting a single number. In contrast, by operating in the internal representation space, our proposed metrics enable *a layer-wise faithfulness analysis and show that faithfulness varies notably across layers* (Figure 2 for LogiQA; results for OpenbookQA, ARC-Easy, and GSM8K are in Appendix C.4).⁴ Our experimental results further reveal three patterns.

Both peak and dip locations are model-specific, with peaks at mid-to-late rather than final layers. For all models except Gemma-2-2B, peak faithfulness often occurs at mid-to-late layers (e.g., on LogiQA, L24/28 for Qwen3-1.7B, L31/42 for Gemma-2-9B; Table 4) and not at the final layer. A plausible explanation is that the final layers are devoted to output formatting and next-token prediction rather than manipulating reasoning concepts. In contrast, faithfulness dips appear at model-specific layers: Qwen3 models exhibit early-mid to mid-layer dips, while Llama-3.1-8B exhibits high layer-to-layer variance with no clear trend (Figure 2; left column). Gemma-2 models are rather task-specific, with the two models showing opposite patterns: Gemma-2-9B drops at late layers on MCQA (Figure 4; Figure 8a–8c) but at early layers on GSM8k (Figure 8d), while Gemma-2-2B exhibits a weaker and less consistent version of the reverse.

Faithfulness manifests as distributed or localized patterns across architectures. Figure 2 and Figure 4 demonstrate that layer-wise Δp stability varies substantially across models: Llama-3.1-8B exhibits the highest layer-to-layer variance, with adjacent layers differing by up to $\Delta p = 0.3$ (Table 4); Qwen3 models are the most stable across datasets, and Gemma-2-2B shows moderate fluctuations of consistent magnitude across layers (Figure 9). Additionally, the contrast between Qwen3’s *distributed faithfulness*, which is consistently high across most layers with great stability, and Llama-3.1-8B’s *localized faithfulness*, concentrated at, e.g., L18–20 and L23–25 on LogiQA, suggests that *CoT faithfulness may differ qualitatively across architectures*,






Figure 4: Causal faithfulness (Δp) variation across normalized depth on LogiQA. Full results are presented in Figure 8.

⁴ S_{pred} and S_{cot} overlap only partially (33–75%), indicating that $S_{\cap}$ captures a selective subset of prediction-relevant features that are also activated during CoT reasoning (Appendix D.1), rather than all features active at the answer position. Additional sanity checks rule out the potential that these are merely answer-identity features (Appendix D.2) or ubiquitous features (Appendix D.3). Moreover, we find that $S_{\cap}$ is largely sufficient to recover the answer alone, compared to the random baseline (Appendix D.4).



(a) Aggregated Δp value and Pearson's r comparison (see raw values in Table 20).



(b) Layer-wise Δp , CC-SAE, Jaccard, and Recall variation on GSM8K (see full results in Figure 15).

Figure 5: SAE configuration ablation study on Gemma-2-2B with **increased width** and **increased sparsity**. Full results with other models across all datasets are provided in Appendix I.

not only in magnitude. Moreover, distributed faithfulness suggests a tight coupling between CoT and prediction throughout the model’s computation, potentially suggesting more robust reasoning that is less susceptible to perturbations at any single layer. In contrast, localized faithfulness implies that the alignment emerges only at specific computational stages, with the CoT reasoning proceeding largely independently elsewhere (Appendix C.5).

Scaling reshapes the layer-wise profile, not just its magnitude. Within the same model family, larger models generally exhibit greater faithfulness in magnitude (Δp) (Figures 2 and 4), aligned with Siegel et al. (2026); Wang et al. (2026) (Table 14 provides additional evidence: for example, Qwen3-8B achieves a normalized $\Delta p = \frac{\Delta p(S_{\cap})}{\Delta p(S_{\text{pred}})}$ of 0.96 on LogiQA, while Qwen3-1.7B reaches only 0.90). Beyond this uniform shift, scaling also reshapes the layer-wise profile. For instance, within the Qwen3 family, scaling from 1.7B to 8B shifts the faithfulness dip from a normalized depth⁵ of 44% (L12-14) to 23% (L6-8) on LogiQA (Figure 2; Table 4), suggesting that larger models may compress the early processing region. Within the Gemma-2 family, scaling from 2B to 9B additionally reduces the coefficient of variation across layers from 23% to 12%.

Implications for CoT faithfulness evaluation. The substantial layer-wise variation raises a practical question: which layer should practitioners use when reporting faithfulness? We recommend reporting either (1) the **peak-layer** Δp , which captures the model’s maximum potential for faithful reasoning, or (2) the **mean** Δp across layers, which reflects overall alignment but may underestimate faithfulness for models with localized patterns.⁶

5.4 SAE CONFIGURATION ABLATION STUDY

Motivation and Setup. SAEs introduce hyperparameters, i.e., *dictionary width* and *target sparsity*, that may confound downstream analyses. To verify that our findings genuinely reflect model behavior rather than SAE-specific artifacts, we conduct ablation studies on Gemma-2-2B (the results of other models are provided in Appendix I). We compare the baseline (width=16k, $\ell_0 = 105$) against two ablation conditions: increased width (16k $\rightarrow$ 65k) and increased sparsity ($\ell_0 = 105 \rightarrow 46$).

⁵Normalized depth is defined as the dip layer index divided by the total number of layers (Table 4).

⁶More actionable implications are further discussed in Appendix C.6. In addition, we provide a comparison between our proposed metrics and existing CoT faithfulness metrics in Appendix E. Meanwhile, we present post-hoc rationalization examples identified by Δp in Appendix H.

Table 1: Natural language descriptions of a subset of shared SAE concepts $S_\cap$. Full trajectory is presented in Table 24 and additional case studies for the remaining datasets are provided in Appendix J.

Gemma-2-9B | LogiQA

Q. Many elderly Beijings have never forgotten the unique skills they saw at temple fairs ... as the modernization of society accelerates, Chinese folk culture is facing unprecedented crisis ... all of which make the development of folk art face difficulties. Which of the following can be derived from this text?

A: Marketization is the way out for folk art
B: Folk culture needs rescue protection ★
C: Urban construction should highlight cultural characteristics
D: ... social status of folk art talents should be improved

Pred: B ✓ **CoT:** B ✓ **Gold:** B ★

CoT: "... the survival of folk art is facing unprecedented crisis ... the development of folk art is facing difficulties. So, the answer is B. The text does not say that marketization is the way out ... that urban construction should highlight cultural characteristics ... So, the answer is B."

L	ID ($S_\cap$)	Feature Description
L5	#4649	Specialized compound nouns ... social sciences, economics, and formal academic assessments.
L11	#10504	Pinyin transliterations of Chinese proper nouns ... names, locations, cultural terms.
L21	#7862	Formal academic or examination-style text ... Chinese-to-English translations in logical reasoning.
L41	#13986	Academic terminology ... cultural heritage, archaeological sites, and logical observations.

marks features overlapping with concepts verbalized in the CoT.

Results. Figure 5a shows a highly correlated Δp under different SAE configurations (Pearson $r > 0.8$). This indicates that the SAE choice does not substantially alter *which* examples are judged as faithful. Figure 5b further demonstrates broadly consistent layer-wise Δp profiles across configurations, with all three exhibiting noticeably similar trends in general. Some variation emerges at mid-to-late layers (L15-20), where the baseline configuration yields slightly higher Δp than the wider or sparser alternatives.⁷ However, the overall profile shape and peak locations remain aligned, suggesting that the layer-wise patterns reported in §5.3 are not artifacts of a particular SAE configuration. Moreover, CC-SAE, Jaccard, and Recall report moderate correlations across configurations, with Jaccard exhibiting the largest sensitivity to sparsity changes, which is expected given its direct dependence on feature set cardinality. Overall, these results indicate that *causal faithfulness captures properties of the model’s reasoning process rather than idiosyncrasies of the SAE*.

5.5 LAYER-TO-LAYER TRANSITIONS IN SHARED CONCEPTS

To interpret the shared concepts $S_\cap$ activated at each layer and analyze the feature transition across layers, we map the active $S_\cap$ to natural-language descriptions using an automated interpretability pipeline following Paulo et al. (2025). Specifically, we extract the top-activating shared concepts at every layer and retrieve their corresponding descriptions, allowing us to qualitatively trace how the model’s encoded concepts evolve across layers (further trajectories are provided in Appendix J).

Table 1 illustrates a Chinese folk art case study from LogiQA, mapping its shared concept ($S_\cap$) transitions and layer-wise descriptions extracted from Gemma-2-9B. $S_\cap$ exhibits a distinct trajectory across layers: early layers encode topic-specific concepts; middle and late layers shift toward more abstract concepts detached from specific entities or explicitly tied to evaluating candidate answers. Specifically, the trajectory reveals a cross-lingual transition: $S_\cap$ is initially anchored in Chinese topic-related entities; at a middle layer (L21), a Chinese-to-English translation feature emerges, suggesting a bridge from topic-specific Chinese concepts into English ones (Wendler et al., 2024).

Moreover, we observe a high Δp in layers associated with question content or answer evaluation; conversely, Δp drops noticeably for layers with topically irrelevant concepts, e.g., L16 and L25 (Table 24). Additionally, certain highly activated concepts across layers mirror content explicitly verbalized in the CoT, suggesting that these concepts indeed track specific reasoning content. Nevertheless, *not all concepts with high Δp are verbalized in the CoT. Conversely, some content of the CoT verbalization may not appear among the highly activated layer-wise concepts.*

⁷One possible explanation is that intermediate layers encode increasingly compositional features with potentially higher superposition density, making them more sensitive to SAE hyperparameter (Gurnee et al., 2026).

6 CONCLUSION

We recast CoT faithfulness as an internal concept grounding: whether an LLM’s CoT reasoning relies on the same SAE concepts that also support the LLM’s direct prediction, and whether the shared concepts causally drive its answer. A single shared SAE renders the cross-pass concept overlap well-defined, supporting correlational alignment metrics and a causal metric Δp , a necessity test validated against count, arbitrary feature selection, and SAE-choice confounds. Across five LLMs and four reasoning benchmarks, although concept alignment reveals substantial feature overlap between passes, it neither establishes nor negates a strong causal link; precisely determining a feature’s contribution requires causal metrics like Δp . Faithfulness generally peaks at mid-to-late depths, is distributed in some architectures yet localized in others, and, within model families, is reorganized rather than merely raised by scale, arguing for layer-aware reporting. Natural-language concept descriptions further show that causally important shared concepts may not be verbalized in the CoT. Overall, our findings suggest a future direction for faithfulness evaluations that also test whether the internal representations associated with a CoT causally support the model’s prediction.

AI USE STATEMENT

The authors used Claude (Code) to assist with language polishing (including grammar, clarity, and coherence), literature search, and code implementation support. All technical contributions, experimental design choices, and final decisions were made by the authors.

ACKNOWLEDGMENTS

We thank **Yifan Wang**, **Yihong Liu**, and **Ruta Binkyte** for their valuable feedback and constructive criticism, which helped improve the paper. Additionally, we thank **Pingjun Hong** and **Xu Shen** for providing CIE scores (Shen et al., 2026) to measure the correlation between our metrics and existing ones (Appendix E).

REFERENCES

- Sahar Admoni, Ofra Amir, Assaf Hallak, and Yftah Ziser. Aligning what LLMs do and say: Towards self-consistent explanations. In Maria Liakata, Viviane P. Moreira, Jiajun Zhang, and David Jurgens (eds.), *Findings of the Association for Computational Linguistics: ACL 2026*, pp. 987–1003, San Diego, California, United States, July 2026. Association for Computational Linguistics. ISBN 979-8-89176-395-1. doi: 10.18653/v1/2026.findings-acl.49. URL <https://aclanthology.org/2026.findings-acl.49/>.
- Chirag Agarwal, Sree Harsha Tanneru, and Himabindu Lakkaraju. Faithfulness vs. Plausibility: On the (Un)Reliability of Explanations from Large Language Models, 2024. URL <https://arxiv.org/abs/2402.04614>.
- Iván Arcuschin, Jett Janiak, Robert Krzyzanowski, Senthoooran Rajamanoharan, Neel Nanda, and Arthur Conmy. Chain-of-Thought Reasoning in the Wild Is Not Always Faithful. In *Workshop on Reasoning and Planning for Large Language Models*, 2025. URL <https://openreview.net/forum?id=L8094Whth0>.
- Trenton Bricken, Adly Templeton, Joshua Batson, Brian Chen, Adam Jermy, Tom Conerly, Nick Turner, Cem Anil, Carson Denison, Amanda Askell, Robert Lasenby, Yifan Wu, Shauna Kravec, Nicholas Schiefer, Tim Maxwell, Nicholas Joseph, Zac Hatfield-Dodds, Alex Tamkin, Karina Nguyen, Brayden McLean, Josiah E Burke, Tristan Hume, Shan Carter, Tom Henighan, and Christopher Olah. Towards monosemanticity: Decomposing language models with dictionary learning. *Transformer Circuits Thread*, 2023. URL <https://transformer-circuits.pub/2023/monosemantic-features/index.html>.
- Xi Chen, Aske Plaat, and Niki van Stein. How does chain of thought think? mechanistic interpretability of chain-of-thought reasoning with sparse autoencoding. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 40, pp. 30297–30305, 2026. doi: 10.1609/aaai.v40i36.40281.

-
- Yanda Chen, Joe Benton, Ansh Radhakrishnan, Jonathan Uesato, Carson Denison, John Schulman, Arushi Somani, Peter Hase, Misha Wagner, Fabien Roger, Vlad Mikulik, Samuel R. Bowman, Jan Leike, Jared Kaplan, and Ethan Perez. Reasoning Models Don’t Always Say What They Think, 2025. URL <https://arxiv.org/abs/2505.05410>.
- James Chua and Owain Evans. Are DeepSeek R1 And Other Reasoning Models More Faithful?, 2025. URL <https://arxiv.org/abs/2501.08156>.
- Peter Clark, Isaac Cowhey, Oren Etzioni, Tushar Khot, Ashish Sabharwal, Carissa Schoenick, and Oyvind Tafjord. Think you have solved question answering? try arc, the ai2 reasoning challenge, 2018. URL <https://arxiv.org/abs/1803.05457>.
- Karl Cobbe, Vineet Kosaraju, Mohammad Bavarian, Mark Chen, Heewoo Jun, Lukasz Kaiser, Matthias Plappert, Jerry Tworek, Jacob Hilton, Reiichiro Nakano, Christopher Hesse, and John Schulman. Training verifiers to solve math word problems, 2021. URL <https://arxiv.org/abs/2110.14168>.
- Hoagy Cunningham, Aidan Ewart, Logan Riggs, Robert Huben, and Lee Sharkey. Sparse autoencoders find highly interpretable features in language models, 2023. URL <https://arxiv.org/abs/2309.08600>.
- Boyi Deng, Xu Wang, Yaoning Wang, Yu Wan, Yubo Ma, Baosong Yang, Haoran Wei, Jialong Tang, Huan Lin, Ruize Gao, Tianhao Li, Qian Cao, Xuancheng Ren, Xiaodong Deng, An Yang, Fei Huang, Dayiheng Liu, and Jingren Zhou. Qwen-scope: Turning sparse features into development tools for large language models, 2026. URL <https://arxiv.org/abs/2605.11887>.
- Murat Dura, Serkan Öztürk, and Selma Tekir. Mechanistic interpretability of chain-of-thought reasoning via sequential activation patching, 2026. URL <https://arxiv.org/abs/2608.22332>.
- 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. URL <https://openreview.net/forum?id=tcsZt9ZNKD>.
- Aaron Grattafiori, Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Alex Vaughan, Amy Yang, Angela Fan, Anirudh Goyal, Anthony Hartshorn, Aobo Yang, Archi Mitra, Archie Sravankumar, Artem Korenev, Arthur Hinsvark, Arun Rao, Aston Zhang, Aurelien Rodriguez, Austen Gregerson, Ava Spataru, Baptiste Roziere, Bethany Biron, Binh Tang, Bobbie Chern, Charlotte Caucheteux, Chaya Nayak, Chloe Bi, Chris Marra, Chris McConnell, Christian Keller, Christophe Touret, Chunyang Wu, Corinne Wong, Cristian Canton Ferrer, Cyrus Nikolaidis, Damien Allonsius, Daniel Song, Danielle Pintz, Danny Livshits, Danny Wyatt, David Esiobu, Dhruv Choudhary, Dhruv Mahajan, Diego Garcia-Olano, Diego Perino, Dieuwke Hupkes, Egor Lakomkin, Ehab AlBadawy, Elina Lobanova, Emily Dinan, Eric Michael Smith, Filip Radenovic, Francisco Guzmán, Frank Zhang, Gabriel Synnaeve, Gabrielle Lee, Georgia Lewis Anderson, Govind Thattai, Graeme Nail, Gregoire Mialon, Guan Pang, Guillem Cucurell, Hailey Nguyen, Hannah Korevaar, Hu Xu, Hugo Touvron, Iliyan Zarov, Imanol Arrieta Ibarra, Isabel Kloumann, Ishan Misra, Ivan Evtimov, Jack Zhang, Jade Copet, Jaewon Lee, Jan Geffert, Jana Vranes, Jason Park, Jay Mahadeokar, Jeet Shah, Jelmer van der Linde, Jennifer Billock, Jenny Hong, Jenya Lee, Jeremy Fu, Jianfeng Chi, Jianyu Huang, Jiawen Liu, Jie Wang, Jiecao Yu, Joanna Bitton, Joe Spisak, Jongsoo Park, Joseph Rocca, Joshua Johnstun, Joshua Saxe, Junteng Jia, Kalyan Vasuden Alwala, Karthik Prasad, Kartikeya Upasani, Kate Plawiak, Ke Li, Kenneth Heafield, Kevin Stone, Khalid El-Arini, Krithika Iyer, Kshitiz Malik, Kuenley Chiu, Kunal Bhalla, Kushal Lakhotia, Lauren Rantala-Young, Laurens van der Maaten, Lawrence Chen, Liang Tan, Liz Jenkins, Louis Martin, Lovish Madaan, Lubo Malo, Lukas Blecher, Lukas Landzaat, Luke de Oliveira, Madeline Muzzi, Mahesh Pasupuleti, Mannat Singh, Manohar Paluri, Marcin Kardas, Maria Tsimpoukelli, Mathew Oldham, Mathieu Rita, Maya Pavlova, Melanie Kambadur, Mike Lewis, Min Si, Mitesh Kumar Singh, Mona Hassan, Naman Goyal, Narjes Torabi, Nikolay Bashlykov, Nikolay Bogoychev, Niladri Chatterji, Ning Zhang, Olivier Duchenne, Onur Çelebi, Patrick Alrassy, Pengchuan Zhang, Pengwei Li, Petar Vasic, Peter Weng, Prajjwal Bhargava, Pratik Dubal, Praveen Krishnan,

Punit Singh Koura, Puxin Xu, Qing He, Qingxiao Dong, Ragavan Srinivasan, Raj Ganapathy, Ramon Calderer, Ricardo Silveira Cabral, Robert Stojnic, Roberta Raileanu, Rohan Maheswari, Rohit Girdhar, Rohit Patel, Romain Sauvestre, Ronnie Polidoro, Roshan Sumbaly, Ross Taylor, Ruan Silva, Rui Hou, Rui Wang, Saghar Hosseini, Sahana Chennabasappa, Sanjay Singh, Sean Bell, Seohyun Sonia Kim, Sergey Edunov, Shaoliang Nie, Sharan Narang, Sharath Rapparthi, Sheng Shen, Shengye Wan, Shruti Bhosale, Shun Zhang, Simon Vandenhende, Soumya Batra, Spencer Whitman, Sten Sootla, Stephane Collot, Suchin Gururangan, Sydney Borodinsky, Tamar Herman, Tara Fowler, Tarek Sheasha, Thomas Georgiou, Thomas Scialom, Tobias Speckbacher, Todor Mihaylov, Tong Xiao, Ujjwal Karn, Vedanuj Goswami, Vibhor Gupta, Vignesh Ramanathan, Viktor Kerkez, Vincent Gonguet, Virginie Do, Vish Vogeti, Vitor Albiero, Vladan Petrovic, Weiwei Chu, Wenhan Xiong, Wenyin Fu, Whitney Meers, Xavier Martinet, Xiaodong Wang, Xiaofang Wang, Xiaoqing Ellen Tan, Xide Xia, Xinfeng Xie, Xuchao Jia, Xuwei Wang, Yaelle Goldschlag, Yashesh Gaur, Yasmine Babaei, Yi Wen, Yiwen Song, Yuchen Zhang, Yue Li, Yuning Mao, Zacharie DelPierre Coudert, Zheng Yan, Zhengxing Chen, Zoe Papakipos, Aaditya Singh, Aayushi Srivastava, Abha Jain, Adam Kelsey, Adam Shajnfeld, Adithya Gangidi, Adolfo Victoria, Ahuva Goldstand, Ajay Menon, Ajay Sharma, Alex Boesenberg, Alexei Baevski, Allie Feinstein, Amanda Kallet, Amit Sangani, Amos Teo, Anam Yunus, Andrei Lupu, Andres Alvarado, Andrew Caples, Andrew Gu, Andrew Ho, Andrew Poulton, Andrew Ryan, Ankit Ramchandani, Annie Dong, Annie Franco, Anuj Goyal, Aparajita Saraf, Arkabandhu Chowdhury, Ashley Gabriel, Ashwin Bharambe, Assaf Eisenman, Azadeh Yazdan, Beau James, Ben Maurer, Benjamin Leonhardi, Bernie Huang, Beth Loyd, Beto De Paola, Bhargavi Paranjape, Bing Liu, Bo Wu, Boyu Ni, Braden Hancock, Bram Wasti, Brandon Spence, Brani Stojkovic, Brian Gamido, Britt Montalvo, Carl Parker, Carly Burton, Catalina Mejia, Ce Liu, Changhan Wang, Changkyu Kim, Chao Zhou, Chester Hu, Ching-Hsiang Chu, Chris Cai, Chris Tindal, Christoph Feichtenhofer, Cynthia Gao, Damon Civin, Dana Beaty, Daniel Kreymer, Daniel Li, David Adkins, David Xu, Davide Testuggine, Delia David, Devi Parikh, Diana Liskovich, Didem Foss, Dingkan Wang, Duc Le, Dustin Holland, Edward Dowling, Eissa Jamil, Elaine Montgomery, Eleonora Presani, Emily Hahn, Emily Wood, Eric-Tuan Le, Erik Brinkman, Esteban Arcaute, Evan Dunbar, Evan Smothers, Fei Sun, Felix Kreuk, Feng Tian, Filippos Kokkinos, Firat Ozgenel, Francesco Caggioni, Frank Kanayet, Frank Seide, Gabriela Medina Florez, Gabriella Schwarz, Gada Badeer, Georgia Swee, Gil Halpern, Grant Herman, Grigory Sizov, Guangyi, Zhang, Guna Lakshminarayanan, Hakan Inan, Hamid Shojanazeri, Han Zou, Hannah Wang, Hanwen Zha, Haroun Habeeb, Harrison Rudolph, Helen Suk, Henry Aspegren, Hunter Goldman, Hongyuan Zhan, Ibrahim Damlaj, Igor Molybog, Igor Tufanov, Ilias Leontiadis, Irina-Elena Veliche, Itai Gat, Jake Weissman, James Geboski, James Kohli, Janice Lam, Japhet Asher, Jean-Baptiste Gaya, Jeff Marcus, Jeff Tang, Jennifer Chan, Jenny Zhen, Jeremy Reizenstein, Jeremy Teboul, Jessica Zhong, Jian Jin, Jingyi Yang, Joe Cummings, Jon Carvill, Jon Shepard, Jonathan McPhie, Jonathan Torres, Josh Ginsburg, Junjie Wang, Kai Wu, Kam Hou U, Karan Saxena, Kartikay Khandelwal, Katayoun Zand, Kathy Matosich, Kaushik Veeraraghavan, Kelly Michelena, Keqian Li, Kiran Jagadeesh, Kun Huang, Kunal Chawla, Kyle Huang, Lailin Chen, Lakshya Garg, Lavender A, Leandro Silva, Lee Bell, Lei Zhang, Liangpeng Guo, Licheng Yu, Liron Moshkovich, Luca Wehrstedt, Madian Khabsa, Manav Avalani, Manish Bhatt, Martynas Mankus, Matan Hasson, Matthew Lennie, Matthias Reso, Maxim Groshev, Maxim Naumov, Maya Lathi, Meghan Keneally, Miao Liu, Michael L. Seltzer, Michal Valko, Michelle Restrepo, Mihir Patel, Mik Vyatskov, Mikayel Samvelyan, Mike Clark, Mike Macey, Mike Wang, Miquel Jubert Hermoso, Mo Metanat, Mohammad Rastegari, Munish Bansal, Nandhini Santhanam, Natascha Parks, Natasha White, Navyata Bawa, Nayan Singhal, Nick Egebo, Nicolas Usunier, Nikhil Mehta, Nikolay Pavlovich Laptev, Ning Dong, Norman Cheng, Oleg Chernoguz, Olivia Hart, Omkar Salpekar, Ozlem Kalinli, Parkin Kent, Parth Parekh, Paul Saab, Pavan Balaji, Pedro Rittner, Philip Bontrager, Pierre Roux, Piotr Dollar, Polina Zvyagina, Prashant Ratanchandani, Pritish Yuvraj, Qian Liang, Rachad Alao, Rachel Rodriguez, Rafi Ayub, Raghotham Murthy, Raghu Nayani, Rahul Mitra, Rangaprabhu Parthasarathy, Raymond Li, Rebekkah Hogan, Robin Battey, Rocky Wang, Russ Howes, Ruty Rinott, Sachin Mehta, Sachin Siby, Sai Jayesh Bondu, Samyak Datta, Sara Chugh, Sara Hunt, Sargun Dhillon, Sasha Sidorov, Satadru Pan, Saurabh Mahajan, Saurabh Verma, Seiji Yamamoto, Sharadh Ramaswamy, Shaun Lindsay, Shaun Lindsay, Sheng Feng, Shenghao Lin, Shengxin Cindy Zha, Shishir Patil, Shiva Shankar, Shuqiang Zhang, Shuqiang Zhang, Sinong Wang, Sneha Agarwal, Soji Sajuyigbe, Soumith Chintala, Stephanie Max, Stephen Chen, Steve Kehoe, Steve Satterfield, Sudarshan Govindaprasad, Sumit Gupta, Summer Deng, Sungmin Cho, Sunny Virk, Suraj Subramanian, Sy Choudhury, Sydney Goldman, Tal Remez, Tamar Glaser, Tamara Best, Thilo

-
- Koehler, Thomas Robinson, Tianhe Li, Tianjun Zhang, Tim Matthews, Timothy Chou, Tzook Shaked, Varun Vontimitta, Victoria Ajayi, Victoria Montanez, Vijai Mohan, Vinay Satish Kumar, Vishal Mangla, Vlad Ionescu, Vlad Poenaru, Vlad Tiberiu Mihailescu, Vladimir Ivanov, Wei Li, Wenchen Wang, Wenwen Jiang, Wes Bouaziz, Will Constable, Xiaocheng Tang, Xiaojian Wu, Xiaolan Wang, Xilun Wu, Xinbo Gao, Yaniv Kleinman, Yanjun Chen, Ye Hu, Ye Jia, Ye Qi, Yenda Li, Yilin Zhang, Ying Zhang, Yossi Adi, Youngjin Nam, Yu, Wang, Yu Zhao, Yuchen Hao, Yundi Qian, Yunlu Li, Yuzi He, Zach Rait, Zachary DeVito, Zef Rosnbrick, Zhaoduo Wen, Zhenyu Yang, Zhiwei Zhao, and Zhiyu Ma. The llama 3 herd of models, 2024. URL <https://arxiv.org/abs/2407.21783>.
- Yoav Gur-Arieh, Ana Marasović, and Mor Geva. Faithfulness metrics don’t measure faithfulness: A meta-evaluation with ground truth, 2026. URL <https://arxiv.org/abs/2605.25052>.
- Wes Gurnee, Nicholas Sofroniew, Adam Pearce, Mateusz Piotrowski, Isaac Kauvar, Runjin Chen, Anna Soligo, Paul Bogdan, Euan Ong, Rowan Wang, Ben Thompson, David Abrahams, Subhash Kantamneni, Emmanuel Ameisen, Joshua Batson, and Jack Lindsey. Verbalizable representations form a global workspace in language models, 2026. URL <https://arxiv.org/abs/2607.15495>.
- Zhengfu He, Wentao Shu, Xuyang Ge, Lingjie Chen, Junxuan Wang, Yunhua Zhou, Frances Liu, Qipeng Guo, Xuanjing Huang, Zuxuan Wu, Yu-Gang Jiang, and Xipeng Qiu. Llama scope: Extracting millions of features from llama-3.1-8b with sparse autoencoders, 2024. URL <https://arxiv.org/abs/2410.20526>.
- Alon Jacovi and Yoav Goldberg. Towards faithfully interpretable NLP systems: How should we define and evaluate faithfulness? In Dan Jurafsky, Joyce Chai, Natalie Schluter, and Joel Tetreault (eds.), *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pp. 4198–4205, Online, July 2020. Association for Computational Linguistics. doi: 10.18653/v1/2020.acl-main.386. URL <https://aclanthology.org/2020.acl-main.386/>.
- Takeshi Kojima, Shixiang Shane Gu, Machel Reid, Yutaka Matsuo, and Yusuke Iwasawa. Large Language Models are Zero-Shot Reasoners. In Alice H. Oh, Alekh Agarwal, Danielle Belgrave, and Kyunghyun Cho (eds.), *Advances in Neural Information Processing Systems*, 2022. URL https://proceedings.neurips.cc/paper_files/paper/2022/hash/8bb0d291acd4acf06ef112099c16f326-Abstract-Conference.html.
- Tomek Korbak, Mikita Balesni, Elizabeth Barnes, Yoshua Bengio, Joe Benton, Joseph Bloom, Mark Chen, Alan Cooney, Allan Dafoe, Anca Dragan, Scott Emmons, Owain Evans, David Farhi, Ryan Greenblatt, Dan Hendrycks, Marius Hobbhahn, Evan Hubinger, Geoffrey Irving, Erik Jenner, Daniel Kokotajlo, Victoria Krakovna, Shane Legg, David Lindner, David Luan, Aleksander Małdry, Julian Michael, Neel Nanda, Dave Orr, Jakub Pachocki, Ethan Perez, Mary Phuong, Fabien Roger, Joshua Saxe, Buck Shlegeris, Martín Soto, Eric Steinberger, Jasmine Wang, Wojciech Zaremba, Bowen Baker, Rohin Shah, and Vlad Mikulik. Chain of thought monitorability: A new and fragile opportunity for ai safety, 2025. URL <https://arxiv.org/abs/2507.11473>.
- Tamera Lanham, Anna Chen, Ansh Radhakrishnan, Benoit Steiner, Carson Denison, Danny Hernandez, Dustin Li, Esin Durmus, Evan Hubinger, Jackson Kernion, et al. Measuring Faithfulness in Chain-of-Thought Reasoning. *arXiv preprint arXiv:2307.13702*, 2023. URL <https://www.anthropic.com/research/measuring-faithfulness-in-chain-of-thought-reasoning>.
- Tom Lieberum, Senthooran Rajamanoharan, Arthur Conmy, Lewis Smith, Nicolas Sonnerat, Vikrant Varma, Janos Kramar, Anca Dragan, Rohin Shah, and Neel Nanda. Gemma scope: Open sparse autoencoders everywhere all at once on gemma 2. In Yonatan Belinkov, Najoung Kim, Jaap Jumelet, Hosein Mohebbi, Aaron Mueller, and Hanjie Chen (eds.), *Proceedings of the 7th BlackboxNLP Workshop: Analyzing and Interpreting Neural Networks for NLP*, pp. 278–300, Miami, Florida, US, November 2024. Association for Computational Linguistics. doi: 10.18653/v1/2024.blackboxnlp-1.19. URL <https://aclanthology.org/2024.blackboxnlp-1.19/>.
- Jian Liu, Leyang Cui, Hanmeng Liu, Dandan Huang, Yile Wang, and Yue Zhang. LogiQA: A Challenge Dataset for Machine Reading Comprehension with Logical Reasoning. In *Proceedings of the Twenty-Ninth International Joint Conference on Artificial Intelligence, IJCAI’20*,

-
- pp. 3622–3628. International Joint Conferences on Artificial Intelligence Organization, 2021. ISBN 9780999241165. doi: 10.24963/ijcai.2020/501. URL <https://doi.org/10.24963/ijcai.2020/501>.
- Uzay Macar, Paul C. Bogdan, Senthooan Rajamanoharan, and Neel Nanda. Thought branches: Interpreting LLM reasoning requires resampling. In *The Fourteenth International Conference on Learning Representations*, 2026. URL <https://openreview.net/forum?id=bVsAuIOvJ5>.
- Todor Mihaylov, Peter Clark, Tushar Khot, and Ashish Sabharwal. Can a suit of armor conduct electricity? a new dataset for open book question answering. In Ellen Riloff, David Chiang, Julia Hockenmaier, and Jun’ichi Tsujii (eds.), *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pp. 2381–2391, Brussels, Belgium, October–November 2018. Association for Computational Linguistics. doi: 10.18653/v1/D18-1260. URL <https://aclanthology.org/D18-1260/>.
- Giovanni Maria Occhipinti, Alessandro Abate, and Nandi Schoots. Probing and steering chain-of-thought unfaithfulness in language models. In *Third Workshop on Test-Time Updates (Main Track)*, 2026. URL <https://openreview.net/forum?id=LocRunEIxK>.
- Letitia Parcalabescu and Anette Frank. On Measuring Faithfulness or Self-Consistency of Natural Language Explanations. In Lun-Wei Ku, Andre Martins, and Vivek Srikumar (eds.), *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 6048–6089, Bangkok, Thailand, August 2024. Association for Computational Linguistics. doi: 10.18653/v1/2024.acl-long.329. URL <https://aclanthology.org/2024.acl-long.329/>.
- Debjit Paul, Robert West, Antoine Bosselut, and Boi Faltings. Making reasoning matter: Measuring and improving faithfulness of chain-of-thought reasoning. In Yaser Al-Onaizan, Mohit Bansal, and Yun-Nung Chen (eds.), *Findings of the Association for Computational Linguistics: EMNLP 2024*, pp. 15012–15032, Miami, Florida, USA, November 2024. Association for Computational Linguistics. doi: 10.18653/v1/2024.findings-emnlp.882. URL <https://aclanthology.org/2024.findings-emnlp.882/>.
- Gonçalo Santos Paulo, Alex Troy Mallen, Caden Juang, and Nora Belrose. Automatically interpreting millions of features in large language models. In Aarti Singh, Maryam Fazel, Daniel Hsu, Simon Lacoste-Julien, Felix Berkenkamp, Tegan Maharaj, Kiri Wagstaff, and Jerry Zhu (eds.), *Proceedings of the 42nd International Conference on Machine Learning*, volume 267 of *Proceedings of Machine Learning Research*, pp. 48393–48421. PMLR, 13–19 Jul 2025. URL <https://proceedings.mlr.press/v267/paulo25a.html>.
- Senthooan Rajamanoharan, Tom Lieberum, Nicolas Sonnerat, Arthur Conmy, Vikrant Varma, János Kramár, and Neel Nanda. Jumping ahead: Improving reconstruction fidelity with jumprelu sparse autoencoders, 2024. URL <https://arxiv.org/abs/2407.14435>.
- Xu Shen, Zhen Tan, Song Wang, Pingjun Hong, Rui Miao, Xin Wang, and Tianlong Chen. Detecting unfaithful chain-of-thought via circuit-guided internal-external discrepancy, 2026. URL <https://arxiv.org/abs/2605.25603>.
- Dong Shu, Xuansheng Wu, Haiyan Zhao, Daking Rai, Ziyu Yao, Ninghao Liu, and Mengnan Du. A survey on sparse autoencoders: Interpreting the internal mechanisms of large language models. In Christos Christodoulopoulos, Tanmoy Chakraborty, Carolyn Rose, and Violet Peng (eds.), *Findings of the Association for Computational Linguistics: EMNLP 2025*, pp. 1690–1712, Suzhou, China, November 2025. Association for Computational Linguistics. ISBN 979-8-89176-335-7. doi: 10.18653/v1/2025.findings-emnlp.89. URL <https://aclanthology.org/2025.findings-emnlp.89/>.
- Noah Y. Siegel, Nicolas Heess, Maria Perez-Ortiz, and Oana-Maria Camburu. Verbosity tradeoffs and the impact of scale on the faithfulness of LLM self-explanations. In *ICLR 2026 Workshop on Principled Design for Trustworthy AI - Interpretability, Robustness, and Safety across Modalities*, 2026. URL <https://openreview.net/forum?id=02pU6ZBiic>.

Jingyi Sun, Qianli Wang, Pepa Atanasova, Nils Feldhus, and Isabelle Augenstein. Investigating the interplay between contextual and parametric chain-of-thought faithfulness under optimization, 2026. URL <https://arxiv.org/abs/2605.24960>.

Gemma Team, Morgane Riviere, Shreya Pathak, Pier Giuseppe Sessa, Cassidy Hardin, Surya Bhu-patiraju, Léonard Hussenot, Thomas Mesnard, Bobak Shahriari, Alexandre Ramé, Johan Fer-ret, Peter Liu, Pouya Tafti, Abe Friesen, Michelle Casbon, Sabela Ramos, Ravin Kumar, Char-line Le Lan, Sammy Jerome, Anton Tsitsulin, Nino Vieillard, Piotr Stanczyk, Sertan Girgin, Nikola Momchev, Matt Hoffman, Shantanu Thakoor, Jean-Bastien Grill, Behnam Neyshabur, Olivier Bachem, Alanna Walton, Aliaksei Severyn, Alicia Parrish, Aliya Ahmad, Allen Hutchi-son, Alvin Abdagic, Amanda Carl, Amy Shen, Andy Brock, Andy Coenen, Anthony Laforge, Antonia Paterson, Ben Bastian, Bilal Piot, Bo Wu, Brandon Royal, Charlie Chen, Chintu Kumar, Chris Perry, Chris Welty, Christopher A. Choquette-Choo, Danila Sinopalnikov, David Wein-berger, Dimple Vijaykumar, Dominika Rogozińska, Dustin Herbison, Elisa Bandy, Emma Wang, Eric Noland, Erica Moreira, Evan Senter, Evgenii Eltyshev, Francesco Visin, Gabriel Rasskin, Gary Wei, Glenn Cameron, Gus Martins, Hadi Hashemi, Hanna Klimczak-Plucińska, Harleen Batra, Harsh Dhand, Ivan Nardini, Jacinda Mein, Jack Zhou, James Svensson, Jeff Stanway, Jetha Chan, Jin Peng Zhou, Joana Carrasqueira, Joana Iljazi, Jocelyn Becker, Joe Fernandez, Joost van Amersfoort, Josh Gordon, Josh Lipschultz, Josh Newlan, Ju yeong Ji, Kareem Mohamed, Kar-tikeya Badola, Kat Black, Katie Millican, Keelin McDonell, Kelvin Nguyen, Kiranbir Sodhia, Kish Greene, Lars Lowe Sjoesund, Lauren Usui, Laurent Sifre, Lena Heuermann, Leticia Lago, Lilly McNealus, Livio Baldini Soares, Logan Kilpatrick, Lucas Dixon, Luciano Martins, Machel Reid, Manvinder Singh, Mark Iverson, Martin Görner, Mat Velloso, Mateo Wirth, Matt Davidow, Matt Miller, Matthew Rahtz, Matthew Watson, Meg Risdal, Mehran Kazemi, Michael Moyni-han, Ming Zhang, Minsuk Kahng, Minwoo Park, Mofi Rahman, Mohit Khatwani, Natalie Dao, Nenshad Bardoliwalla, Nesh Devanathan, Neta Dumai, Nilay Chauhan, Oscar Wahltinez, Pankil Botarda, Parker Barnes, Paul Barham, Paul Michel, Pengchong Jin, Petko Georgiev, Phil Cullit-ton, Pradeep Kuppala, Ramona Comanescu, Ramona Merhej, Reena Jana, Reza Ardeshtir Rokni, Rishabh Agarwal, Ryan Mullins, Samaneh Saadat, Sara Mc Carthy, Sarah Cogan, Sarah Perrin, Sébastien M. R. Arnold, Sebastian Krause, Shengyang Dai, Shruti Garg, Shruti Sheth, Sue Ron-strom, Susan Chan, Timothy Jordan, Ting Yu, Tom Eccles, Tom Hennigan, Tomas Kocisky, Tulsee Doshi, Vihan Jain, Vikas Yadav, Vilobh Meshram, Vishal Dharmadhikari, Warren Barkley, Wei Wei, Wenming Ye, Woohyun Han, Woosuk Kwon, Xiang Xu, Zhe Shen, Zhitao Gong, Zichuan Wei, Victor Cotruta, Phoebe Kirk, Anand Rao, Minh Giang, Ludovic Peran, Tris Warkentin, Eli Collins, Joelle Barral, Zoubin Ghahramani, Raia Hadsell, D. Sculley, Jeanine Banks, Anca Dra-gan, Slav Petrov, Oriol Vinyals, Jeff Dean, Demis Hassabis, Koray Kavukcuoglu, Clement Fara-bet, Elena Buchatskaya, Sebastian Borgeaud, Noah Fiedel, Armand Joulin, Kathleen Kenealy, Robert Dadashi, and Alek Andreev. Gemma 2: Improving open language models at a practical size, 2024. URL <https://arxiv.org/abs/2408.00118>.

Adly Templeton, Tom Conerly, Jonathan Marcus, Jack Lindsey, Trenton Bricken, Brian Chen, Adam Pearce, Craig Citro, Emmanuel Ameisen, Andy Jones, Hoagy Cunningham, Nicholas L Turner, Callum McDougall, Monte MacDiarmid, Alex Tamkin, Esin Durmus, Tristan Hume, Francesco Mosconi, C. Daniel Freeman, Theodore R. Sumers, Edward Rees, Joshua Batson, Adam Jermy, Shan Carter, Chris Olah, and Tom Henighan. Scaling monosemanticity: Extracting interpretable features from claude 3 sonnet, 2026. URL <https://arxiv.org/abs/2605.29358>.

Miles Turpin, Julian Michael, Ethan Perez, and Samuel R. Bowman. Language Mod-els Don’t Always Say What They Think: Unfaithful Explanations in Chain-of-Thought Prompting. In *Thirty-seventh Conference on Neural Information Processing Systems*, 2023. URL https://proceedings.neurips.cc/paper_files/paper/2023/hash/ed3fea9033a80feal376299fa7863f4a-Abstract-Conference.html.

Martin Tutek, Fateme Hashemi Chaleshtori, Ana Marasovic, and Yonatan Belinkov. Measuring Chain-of-Thought Faithfulness by Unlearning Reasoning Steps. In Christos Christodoulopou-los, Tanmoy Chakraborty, Carolyn Rose, and Violet Peng (eds.), *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing*, pp. 9935–9960, Suzhou, China, November 2025. Association for Computational Linguistics. ISBN 979-8-89176-332-6. doi: 10.18653/v1/2025.emnlp-main.504. URL <https://aclanthology.org/2025.emnlp-main.504/>.

-
- Qianli Wang, Nils Feldhus, Pepa Atanasova, Fedor Splitt, Simon Ostermann, Sebastian Möller, and Vera Schmitt. Can large language models still explain themselves? investigating the impact of quantization on self-explanations, 2026. URL <https://arxiv.org/abs/2601.00282>.
- Dennis Wei, Yannis Belkhir, Erik Miehl, and Radu Marinescu. Local causal attribution of chain-of-thought reasoning, 2026. URL <https://arxiv.org/abs/2606.21821>.
- Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, brian ichter, Fei Xia, Ed H. Chi, Quoc V Le, and Denny Zhou. Chain-of-Thought Prompting Elicits Reasoning in Large Language Models. In Alice H. Oh, Alekh Agarwal, Danielle Belgrave, and Kyunghyun Cho (eds.), *Advances in Neural Information Processing Systems*, 2022. URL https://proceedings.neurips.cc/paper_files/paper/2022/hash/9d5609613524ecf4f15af0f7b31abca4-Abstract-Conference.html.
- Chris Wendler, Veniamin Veselovsky, Giovanni Monea, and Robert West. Do llamas work in English? on the latent language of multilingual transformers. In Lun-Wei Ku, Andre Martins, and Vivek Srikumar (eds.), *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 15366–15394, Bangkok, Thailand, August 2024. Association for Computational Linguistics. doi: 10.18653/v1/2024.acl-long.820. URL <https://aclanthology.org/2024.acl-long.820/>.
- An Yang, Anfeng Li, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chang Gao, Chengen Huang, Chenxu Lv, Chujie Zheng, Dayiheng Liu, Fan Zhou, Fei Huang, Feng Hu, Hao Ge, Haoran Wei, Huan Lin, Jialong Tang, Jian Yang, Jianhong Tu, Jianwei Zhang, Jianxin Yang, Jiayi Yang, Jing Zhou, Jingren Zhou, Junyang Lin, Kai Dang, Keqin Bao, Kexin Yang, Le Yu, Lianghao Deng, Mei Li, Mingfeng Xue, Mingze Li, Pei Zhang, Peng Wang, Qin Zhu, Rui Men, Ruize Gao, Shixuan Liu, Shuang Luo, Tianhao Li, Tianyi Tang, Wenbiao Yin, Xingzhang Ren, Xinyu Wang, Xinyu Zhang, Xuancheng Ren, Yang Fan, Yang Su, Yichang Zhang, Yinger Zhang, Yu Wan, Yuqiong Liu, Zekun Wang, Zeyu Cui, Zhenru Zhang, Zhipeng Zhou, and Zihan Qiu. Qwen3 technical report, 2025. URL <https://arxiv.org/abs/2505.09388>.
- Donald Ye, Max Loffgren, Om Kotadia, Linus Wong, and Jonas Rohweder. Mechanistic evidence for faithfulness decay in chain-of-thought reasoning. In *Workshop on Latent & Implicit Thinking – Going Beyond CoT Reasoning*, 2026. URL <https://openreview.net/forum?id=wVj7dB7waI>.
- Kerem Zaman and Shashank Srivastava. A Causal Lens for Evaluating Faithfulness Metrics. In Christos Christodoulopoulos, Tanmoy Chakraborty, Carolyn Rose, and Violet Peng (eds.), *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing*, pp. 29413–29437, Suzhou, China, November 2025. Association for Computational Linguistics. ISBN 979-8-89176-332-6. doi: 10.18653/v1/2025.emnlp-main.1496. URL <https://aclanthology.org/2025.emnlp-main.1496/>.
- Hengyuan Zhang, Zhihao Zhang, Ercong Nie, Mingyang Wang, Zunhai Su, Yiwei Wang, Qianli Wang, Shuzhou Yuan, Xufeng Duan, Qibo Xue, Zeping Yu, Chenming Shang, Xiao Liang, Jing Xiong, Hui Shen, Chaofan Tao, Zhengwu Liu, Senjie Jin, Zhiheng Xi, Dongdong Zhang, Sophia Ananiadou, Tao Gui, Ruobing Xie, Hayden Kwok-Hay So, Hinrich Schuetze, Xuanjing Huang, Qi Zhang, and Ngai Wong. Locate, steer, and improve: A practical survey of actionable mechanistic interpretability in large language models. In Maria Liakata, Viviane P. Moreira, Jiajun Zhang, and David Jurgens (eds.), *Findings of the Association for Computational Linguistics: ACL 2026*, pp. 10317–10362, San Diego, California, United States, July 2026. Association for Computational Linguistics. ISBN 979-8-89176-395-1. URL <https://aclanthology.org/2026.findings-acl.502/>.

A PRELIMINARY

A.1 COMPARISON WITH PRIOR WORK

Table 2 illustrates the CoT faithfulness evaluation metric comparison. The evaluation dimensions are listed as follows:

- **Granularity:** The level at which the CoT is treated;
- **Intervention:** Whether the CoT is adapted or modified;
- **Causal:** Whether the method employs a causal framework (e.g., activation patching);
- **Inference Time:** The time complexity of the given metric;
- **Polysemanticity:** Whether the metric can isolate decoupled features;
- **Training:** Whether additional training is required for the *tested* model;
- **Per-Instance:** Whether the faithfulness evaluation is aggregate or per-instance.

For our proposed causal metric Δp at a given layer ℓ , an intervention can be applied, and its inference time is $O(1)$, as it only requires a constant number of forward passes, which is notably more efficient than most prior work. Moreover, it is training-free when SAEs are available, can mitigate polysemanticity, and can be used to measure the CoT faithfulness of each individual instance.

Table 2: CoT faithfulness evaluation metric comparison. Intervention: Whether the method perturbs the CoT or the input. n : number of input tokens; L : number of CoT steps; m : the number of traced tokens per step; P : resamples per step.

Method	Granularity	Intervention	Causal	Inference Time	Polysemanticity	Training	Per-Instance
Bias Feature (Turpin et al., 2023)	CoT	✓	✗	$O(1)$	✗	✗	✓
CoT Corruption (Lanham et al., 2023)	CoT	✓	✗	$O(1)$	✗	✗	✓
NLDD (Ye et al., 2026)	Step	✓	✗	$O(L)$	✗	✗	✓
FUR (Tutek et al., 2025)	Step	✓	✗	$O(L)$	✗	✓	✓
Thought Branches (Macar et al., 2026)	Step	✓	✓	$O(L \cdot R)$	✗	✗	✓
AttriCoT (Wei et al., 2026)	Step	✓	✓	$O(L)$	✗	✗	✓
CIE-Scorer (Shen et al., 2026)	Circuit	✓	✓	$O(L \cdot m)$	✗	✓	✓
CC-SHAP (Parcalabescu & Frank, 2024)	Input Token	✗	✗	$O(2^n)$	✗	✗	✓
CC-LIME (Admoni et al., 2026)	Input Token	✗	✗	$O(n)$	✗	✗	✓
SAE Patching (Chen et al., 2026)	SAE Concept	✓	✓	$O(1)$	✓	✓	✗
Our	SAE Concept	✓	✓	$O(1)$	✓	✗	✓

A.2 SEPARATE VS. SHARED SAEs




Figure 6: Comparison between Chen et al. (2026) and our metrics.

Figure 6 illustrates the differences between Chen et al. (2026) and our metrics. First, Chen et al. (2026) train and deploy two distinct SAEs on the prediction and CoT passes, respectively, which makes the two SAEs’ activations not directly comparable. In contrast, we use a single SAE for both passes, which enables the measurement of concept alignment and causal grounding. Furthermore, they patch a_{cot} into the prediction pass and measure the changes in answer log-probability. In comparison, we ablate the shared SAE features that are proven to causally drive the answer (§5.2). Additionally, we have validated that patching only the shared SAE features is sufficient to largely reproduce the original answer (Appendix D.4).

A.3 VENN DIAGRAM OF SAE FEATURES




Figure 7: Venn diagram of SAE features.

Figure 7 displays an overview of the SAE features. In general, the features can be categorized into three classes: $S_{\text{pred-only}}$, $S_{\text{cot-only}}$, and $S_{\cap}$, where $S_{\text{pred}} = S_{\text{pred-only}} \cup S_{\cap}$ and $S_{\text{cot}} = S_{\text{cot-only}} \cup S_{\cap}$. This is enabled by employing a shared SAE, which allows us to assess the shared concepts that drive the model toward answers faithfully across the prediction and CoT passes. Appendix D.1 further presents the SAE concept composition analysis.

B EXPERIMENTAL DETAIL

B.1 MODELS & SAEs

Table 3 illustrates the hidden state dimension, model link, and residual-stream SAE suite of each model.

Table 3: Hidden state dimension, model link, and residual-stream SAE suite of each model.

Model	d_{model}	Link	SAE Scope
Llama-3.1-8B (Grattafiori et al., 2024)	4096	meta-llama/Llama-3.1-8B	Llama-Scope (He et al., 2024)
Gemma-2-2B (Team et al., 2024)	2304	google/gemma-2-2b	Gemma-Scope (Lieberum et al., 2024)
Gemma-2-9B (Team et al., 2024)	3584	google/gemma-2-9b	Gemma-Scope (Lieberum et al., 2024)
Qwen3-1.7B (Yang et al., 2025)	2048	Qwen/Qwen3-1.7B	Qwen-Scope (Deng et al., 2026)
Qwen3-8B (Yang et al., 2025)	4096	Qwen/Qwen3-8B	Qwen-Scope (Deng et al., 2026)

We note that there is one shared SAE *between passes at a given layer*, rather than a single SAE shared across layers.⁸ We extract activations at the output of each transformer layer (post-attention, post-MLP, post-residual connection), corresponding to `resid_post` in TransformerLens⁹ terminology. Formally, for layer ℓ :

$$\mathbf{r}^{(\ell)} = \text{LayerNorm}(\mathbf{h}^{(\ell-1)} + \text{Attn}^{(\ell)} + \text{MLP}^{(\ell)}) \quad (10)$$

⁸Since the residual stream’s distribution and semantic abstraction level change qualitatively with layer depth, a single shared SAE cannot converge to a coherent sparse mapping.

⁹<https://github.com/TransformerLensOrg/TransformerLens>




Figure 8: Causal faithfulness (Δp) variation across normalized model depth.

B.2 DECODING CONFIGURATION

We evaluate CoT faithfulness using greedy decoding (`temperature=0`, `max_new_tokens=2048`) for CoT generation. The reasoning chain is extracted by splitting at the “Answer:” delimiter.

B.3 INFERENCE TIME

All experiments were conducted on a single NVIDIA RTX A6000 GPU. Under this setup, the average computational cost for a single Δp was approximately 15 seconds.

C CoT FAITHFULNESS EVALUATION

C.1 FAITHFULNESS METRIC DISCUSSION

We emphasize that our metric operationalizes faithfulness as **causal cross-path feature reuse**: whether the CoT-conditioned answer pass shares causally important features with the direct-answer pass. Our contribution is methodological, enabling internal comparison where prior work relied on behavioral proxies, rather than claiming to resolve the faithfulness question definitively.

C.2 CAUSAL FAITHFULNESS VARIATION

Figure 8 displays Δp variation across relative model depth, defined as the fraction of the current layer over the total model layers. Table 4 further reports the peak layer at which each model’s internal representations exhibit maximal faithfulness, alongside the corresponding normalized depth. Across most models, faithfulness peaks in the latter half of the network, with normalized depths typically exceeding 70–90%, suggesting that decisive, faithful representations emerge predominantly in deeper layers. In contrast, Gemma-2-2B is a notable outlier, with its faithfulness peak occurring




Figure 9: Layer stability across models.

at or near the earliest layers, implying that its most faithful representations are formed very early in the network rather than progressively refined through depth. This divergence highlights substantial variation in how faithfulness develops across layer depth, depending on model architecture and scale.

C.3 LAYER STABILITY

Figure 9 shows the layer stability across models. Gemma-2-9B and Qwen3 models exhibit generally high stability, whereas Llama3.1-8B demonstrates noticeable variance across different layers, suggesting localized faithfulness.

C.4 CoT FAITHFULNESS EVALUATION

Figure 10, 11, 12, and 13 show the faithfulness evaluation with three correlational metrics and Δp on LogiQA, OpenbookQA, ARC-Easy, and GSM8K, respectively.

C.5 LAYER-WISE FAITHFULNESS DYNAMICS

A natural question arising from our layer-wise analysis is whether models exhibit “repair”, i.e., whether faithfulness that degrades in earlier layers can be recovered in subsequent layers.




Figure 10: CoT faithfulness evaluation across layers for different models on LogiQA. **Left column:** Causal faithfulness (Δp) measures the probability drop when ablating shared features. **Middle column:** Correlational metrics (CC-SAE, Jaccard, Recall) between prediction and CoT activations. **Right column:** Spearman correlation heatmap between each correlational metric and Δp across layers. Shaded regions indicate 95% confidence intervals (left and middle columns). Additional results on the *remaining datasets* are provided in Appendix C.4.

C.5.1 SETUP

We track the evolution of the causal effect (Δp) across layers for each instance, using the first layer as a baseline. Specifically, for each instance i and layer ℓ , we compute the growth ratio:

$$r_i^{(\ell)} = \frac{\Delta p_i^{(\ell)}}{\Delta p_i^{(0)}} \quad (11)$$

We then report the median growth ratio across instances at each layer, along with the interquartile range to quantify variability. A ratio greater than 1 indicates increased faithfulness relative to the first layer, while a ratio less than 1 indicates decreased faithfulness.

C.5.2 RESULTS

Figure 14 presents the layer-wise faithfulness dynamics for five models across four datasets. We observe notable differences across model families.

Gemma models exhibit U-shaped repair patterns. Both Gemma-2-2B and Gemma-2-9B show a characteristic pattern where faithfulness drops in the early-to-middle layers before recovering toward the final layers. For Gemma-2-2B, the median Δp ratio drops to approximately 0.4–0.5 around layers 5–6 across all datasets, then gradually recovers to ≈ 1.0 by the final layer. This U-shaped trajectory suggests that intermediate layers may perform transformations that temporarily reduce alignment between CoT features and answer-relevant computations, but subsequent layers repair this misalignment.



Figure 11: CoT faithfulness evaluation across layers for different models on OpenbookQA. **Left column:** Causal faithfulness (Δp) measures the probability drop when ablating shared features. **Middle column:** Correlational metrics (CC-SAE, Jaccard, Recall) between prediction and CoT activations. **Right column:** Heatmap showing Spearman correlation between each correlational metric and Δp across layers. Shaded regions indicate 95% confidence intervals.

Qwen models maintain stable faithfulness across layers. In contrast, both Qwen3-1.7B and Qwen3-8B exhibit remarkably stable Δp ratios close to 1.0 throughout all layers. This suggests that whatever causal relationship exists between CoTs and predictions is established early and preserved throughout the forward pass, with minimal perturbation or correction in later layers.

Llama-3.1-8B shows high layer-wise variability. Llama-3.1-8B exhibits substantial oscillations in the Δp ratio across layers, with no clear monotonic trend, aligned with Figure 9. This high variability suggests that the SAE features at different layers may capture qualitatively different aspects of the reasoning process.

C.6 IMPLICATIONS

Beyond reporting, the layer-wise profile also carries practical guidance. As peak layers concentrate the causally important shared concepts, probing, steering, and CoT monitoring are best deployed at mid-to-late peak layers rather than at the final layers. This suggests that readouts anchored to last-layer representations may underestimate faithfulness. The distributed-localized faithfulness contrast further yields a testable prediction: models with localized faithfulness should be disproportionately brittle to pruning, quantization, or editing at their few alignment-bearing layers, while distributed models should degrade more gracefully.




Figure 12: CoT faithfulness evaluation across layers for different models on ARC-Easy. **Left column:** Causal faithfulness (Δp) measures the probability drop when ablating shared features. **Middle column:** Correlational metrics (CC-SAE, Jaccard, Recall) between prediction and CoT activations. **Right column:** Heatmap showing Spearman correlation between each correlational metric and Δp across layers. Shaded regions indicate 95% confidence intervals.

D SANITY CHECK

D.1 SAE CONCEPT COMPOSITION

Table 5 reports the feature set composition across models. The intersection $S_{\cap}$ comprises 33–75% of S_{pred} , leaving substantial pred-only ($|S_{\text{pred}} \setminus S_{\cap}|$) and CoT-only ($|S_{\text{cot}} \setminus S_{\cap}|$) subsets. For models with top- k SAEs (Qwen3 models), $|S_{\text{pred}}| = |S_{\text{cot}}| = 100$ is set by construction, yet overlap remains around 50%. The partial overlap indicates that *the two passes activate distinct feature sets despite predicting the same answer, ruling out trivial explanations where $S_{\text{pred}} \approx S_{\text{cot}}$.*

D.2 ANSWER IDENTITY

To rule out the possibility that shared features merely encode the answer token (e.g., A/B/C/D) rather than reflecting genuine reasoning on MCQA datasets, we identify answer-identity features by encoding each letter in isolation.¹⁰ Specifically,

1. **Isolated-token feature extraction.** For each answer letter $\ell \in \{A, B, C, D\}$, we tokenize ℓ in isolation, take the model’s residual stream at the target layer and final token position, and SAE-encode it. The set of features with nonzero activation, F_{ℓ} , is the answer-identity

¹⁰GSM8K is intentionally excluded from this sanity check, as it does not use fixed options.




Figure 13: CoT faithfulness evaluation across layers for different models on GSM8K. **Left column:** Causal faithfulness (Δp) measures the probability drop when ablating shared features. **Middle column:** Correlational metrics (CC-SAE, Jaccard, Recall) between prediction and CoT activations. **Right column:** Heatmap showing Spearman correlation between each correlational metric and Δp across layers. Shaded regions indicate 95% confidence intervals.


feature set for letter ℓ ; this is computed once per dataset–model pair and reused across items.

2. **Per-item prediction pass.** For each item, we build the prediction-only prompt (no CoT), take the model’s greedy answer $y = \arg \max_{\ell \in \{A, B, C, D\}} \text{logit}(\ell)$ at the final position, and extract S_{pred} . We use S_{pred} rather than $S_{\cap}$ as the ablated set in this control, since it isolates the identity confound from the prediction pass directly, independent of whether the CoT also activates the same identity features.
3. **Targeted ablation.** We compute the causal-faithfulness effect (§3.5) twice: once ablating the full S_{pred} (Δp_{full}), and once ablating $S_{\text{pred}} \setminus F_y$, i.e., S_{pred} with the letter’s own identity features removed ($\Delta p_{\text{no-pred}}$), i.e., we ablate F_{ℓ} – all four letters’ feature sets.
4. **Aggregation.** We average Δp_{full} and $\Delta p_{\text{no-pred}}$ across items per dataset–model pair and report retention as the ratio of means, $\overline{\Delta p_{\text{no-pred}}} / \overline{\Delta p_{\text{full}}}$.


Results. Table 6 shows that after removing these features from $S_{\cap}$, the causal effect remains nearly unchanged (mean retention = 0.997 ± 0.007). Furthermore, the number of option-related features contained in $S_{\cap}$ is very limited. Together, these observations indicate that Δp reflects *reasoning-related features, rather than superficial answer encoding*.




(a) Gemma-2-2B




(b) Gemma-2-9B



(c) Qwen3-1.7B



(d) Qwen3-8B



(e) Llama-3.1-8B

Figure 14: Layer-wise faithfulness dynamics across datasets.

Table 4: Faithfulness peak layer, corresponding Δp values, relative depth of the peak layer, and averaged Δp values across models.

	Model	Peak Layer	Peak Δp	Rel. Dep.	Avg Δp
LogiQA	Llama-3.1-8B	25	0.350	93%	0.145 ± 0.099
	Qwen3-1.7B	24	0.566	89%	0.512 ± 0.049
	Qwen3-8B	34	0.713	97%	0.642 ± 0.067
	Gemma-2-2B	1	0.269	4%	0.197 ± 0.045
	Gemma-2-9B	31	0.412	76%	0.356 ± 0.042
OpenbookQA	Llama-3.1-8B	25	0.585	93%	0.275 ± 0.292
	Qwen3-1.7B	19	0.694	70%	0.589 ± 0.458
	Qwen3-8B	27	0.872	77%	0.798 ± 0.339
	Gemma-2-2B	1	0.379	4%	0.282 ± 0.243
	Gemma-2-9B	32	0.679	78%	0.597 ± 0.280
ARC-Easy	Llama-3.1-8B	24	0.805	77%	0.371 ± 0.336
	Qwen3-1.7B	22	0.894	81%	0.795 ± 0.377
	Qwen3-8B	30	0.976	86%	0.924 ± 0.214
	Gemma-2-2B	0	0.615	0%	0.479 ± 0.298
	Gemma-2-9B	29	0.896	71%	0.810 ± 0.220
GSM8K	Llama-3.1-8B	5	0.062	16%	0.037 ± 0.089
	Qwen3-1.7B	6	0.187	22%	0.160 ± 0.243
	Qwen3-8B	19	0.197	54%	0.169 ± 0.287
	Gemma-2-2B	1	0.183	4%	0.124 ± 0.152
	Gemma-2-9B	1	0.276	2%	0.123 ± 0.184

Table 5: Feature set statistics across models and datasets (averaged across layers). $S_\cap$ comprises a fraction of S_{pred} , with substantial pred-only and CoT-only features.

Dataset	Model	$ S_{\text{pred}} $	$ S_{\text{cot}} $	$ S_\cap $	Pred-only	CoT-only	$ S_\cap / S_{\text{pred}} $
LogiQA	Llama-3.1-8B	17.6	18.6	13.4	4.2	5.2	75.3%
	Qwen3-1.7B	100.0	100.0	54.9	45.1	45.1	54.9%
	Qwen3-8B	100.0	100.0	61.0	39.0	39.0	61.0%
	Gemma-2-2B	68.2	66.8	40.6	27.6	26.2	62.2%
	Gemma-2-9B	93.1	93.9	56.7	36.4	37.2	63.9%
OpenbookQA	Llama-3.1-8B	18.3	18.9	12.7	5.6	6.3	68.3%
	Qwen3-1.7B	100.0	100.0	48.8	51.2	51.2	48.8%
	Qwen3-8B	100.0	100.0	49.5	50.5	50.5	49.5%
	Gemma-2-2B	74.6	69.1	37.3	37.3	31.8	53.8%
	Gemma-2-9B	90.7	93.3	51.9	38.8	41.4	61.0%
ARC-Easy	Llama-3.1-8B	18.7	19.5	13.5	5.2	6.0	70.7%
	Qwen3-1.7B	100.0	100.0	48.1	51.9	51.9	48.1%
	Qwen3-8B	100.0	100.0	51.1	48.9	48.9	51.1%
	Gemma-2-2B	73.1	67.7	38.2	34.8	29.5	56.2%
	Gemma-2-9B	85.3	88.7	52.8	32.5	35.9	65.8%
GSM8K	Llama-3.1-8B	18.6	18.7	10.2	8.4	8.5	50.4%
	Qwen3-1.7B	100.0	100.0	36.4	63.6	63.6	36.4%
	Qwen3-8B	100.0	100.0	33.3	66.7	66.7	33.3%
	Gemma-2-2B	56.8	55.3	25.8	30.9	29.4	50.1%
	Gemma-2-9B	72.1	71.1	24.9	47.2	46.2	40.4%

D.3 ROBUSTNESS TO UBIQUITOUS FEATURES

$S_\cap$ may partly consist of ubiquitous features, i.e., generic multiple-choice artifacts that recur across nearly all items regardless of content, rather than item-specific reasoning concepts. If such features drive a substantial share of Δp , the causal effect would reflect superficial task formatting rather than

Table 6: Answer-identity ablation results. $\delta_X = \Delta p_{\text{full}} - \Delta p_{\text{no-}X}$ is the drop in causal effect after removing option X ’s answer-identity features; $\#X$ is the number of option- X features in $S_{\cap}$ that are responsible for outputting the answer corresponding to one of the predefined option letters. Avg. retention ≈ 1 indicates the causal effect is driven by reasoning features, not answer-encoding.

Dataset	Model	Δp_{full}	δ_A	δ_B	δ_C	δ_D	Avg. Ret.	#A	#B	#C	#D	$ S_{\cap} $
ARC-Easy	Gemma-2-2B	0.6812	0.0000	0.0000	0.0000	0.0000	1.0000	8.8	7.3	6.7	7.4	89
	Gemma-2-9B	0.9126	0.0045	0.0000	0.0000	-0.0000	0.9985	4.0	2.4	2.3	4.4	91
	Llama-3.1-8B	0.8513	0.0001	0.0052	0.0030	0.0036	0.9963	4.9	2.1	5.7	4.0	37
	Qwen3-1.7B	0.9898	-0.0000	-0.0000	-0.0000	-0.0000	1.0000	4.5	5.6	5.6	5.4	100
	Qwen3-8B	0.9879	0.0000	0.0000	0.0000	0.0000	1.0000	3.2	4.6	4.7	4.6	100
LogiQA	Gemma-2-2B	0.3349	0.0000	0.0000	0.0000	0.0000	1.0000	8.7	6.3	5.2	5.5	84
	Gemma-2-9B	0.5236	0.0052	0.0000	0.0000	0.0000	0.9980	3.3	2.3	3.1	5.0	92
	Llama-3.1-8B	0.5177	0.0004	0.0140	0.0122	0.0116	0.9809	3.9	3.1	3.5	3.7	38
	Qwen3-1.7B	0.9129	-0.0000	-0.0000	-0.0000	-0.0000	1.0000	4.4	5.5	5.5	5.2	100
	Qwen3-8B	0.9053	0.0000	0.0000	0.0000	0.0000	1.0000	3.8	6.0	6.3	5.7	100
OpenBookQA	Gemma-2-2B	0.5020	0.0000	0.0000	0.0000	0.0000	1.0000	8.5	6.9	5.9	6.6	90
	Gemma-2-9B	0.7656	0.0223	0.0000	0.0000	0.0000	0.9923	2.9	2.2	2.4	3.8	83
	Llama-3.1-8B	0.7337	0.0002	0.0092	0.0050	0.0060	0.9924	4.8	2.1	5.7	4.1	38
	Qwen3-1.7B	0.9576	-0.0000	-0.0000	-0.0000	-0.0000	1.0000	4.3	5.2	5.2	5.1	100
	Qwen3-8B	0.9771	0.0000	0.0000	0.0000	0.0000	1.0000	3.2	4.5	4.6	4.6	100

genuine CoT-driven reasoning. For each dataset–model pair, we identify ubiquitous features and quantify their contribution to Δp in two passes:

1. **Pass 1: frequency collection.** For every item i , we extract $S_{\cap}^{(i)} = S_{\text{pred}}^{(i)} \cap S_{\text{cot}}^{(i)}$ (§3.5) and accumulate a corpus-level frequency count $\text{count}(f) = |\{i : f \in S_{\cap}^{(i)}\}|$ for every feature f observed in any item’s intersection set.
2. **Thresholding.** Given N items with a valid intersection set, we define the ubiquitous set $U_{\tau} = \{f : \text{count}(f)/N > \tau\}$ for a chosen frequency threshold $\tau \in (0, 1)$.
3. **Pass 2: ablation comparison.** For every item, we recompute the causal-faithfulness intervention (§3.5) three times, ablating (i) the full intersection $S_{\cap}^{(i)}$, (ii) the intersection with ubiquitous features removed, $S_{\cap}^{(i)} \setminus U_{\tau}$, and (iii) the ubiquitous features alone, $S_{\cap}^{(i)} \cap U_{\tau}$, yielding Δp_{full} , $\Delta p_{\text{no-gen}}$, and Δp_{gen} per item.
4. **Aggregation.** We average each quantity across items within a dataset–model pair and report **retention** as

$$\text{Retention} = \frac{\Delta p_{\text{no-gen}}}{\Delta p_{\text{full}}},$$

Retention ≈ 1 indicates the causal effect is largely content-specific, while retention $\ll 1$ indicates it is substantially driven by generic, ubiquitous features.

Results. Table 7 reports Δp_{full} , $\Delta p_{\text{no-gen}}$, Δp_{gen} , and retention for each dataset–model pair. Qwen3-1.7B and Qwen3-8B retain 1.00 of the causal effect after removing ubiquitous features across all three datasets, indicating their $S_{\cap}$ -driven causal effect is content-specific. Gemma-2-2B similarly retains 0.88–1.28. Llama-3.1-8B and Gemma-2-9B show more moderate retention. Averaged across all 15 dataset–model pairs, 87% of the causal effect measured by Δp survives removal of ubiquitous features, indicating that *the causal effect of $S_{\cap}$ predominantly reflects content-specific concept reuse rather than generic MCQA artifacts*.

D.4 SUFFICIENCY TEST

To complement our necessity analysis, which shows that ablating $S_{\cap}$ reduces answer probability in §5.3, we additionally conduct a sufficiency test to answer the question: can $S_{\cap}$ alone largely recover the model’s answer? Specifically, we define three conditions:

1. **Base:** No intervention; measure $p_{\text{base}} = p(\text{answer} \mid \mathbf{r}_{\text{pred}})$;
2. **Zeroed:** Replace $\mathbf{r}_{\text{pred}}$ with $\mathbf{0}$; measure p_{zero} ;

Table 7: Ubiquitous-feature ablation at threshold $\tau = 0.9$. $N_{\text{ubiq}}/N_{\text{total}}$: number of features flagged ubiquitous (frequency $> \tau$ across items) out of all features observed in any $S_\cap$. Retention $= \Delta p_{\text{no-gen}}/\Delta p_{\text{full}}$; retention ≈ 1 indicates the causal effect is content-specific rather than driven by ubiquitous features.

Dataset	Model	$N_{\text{ubiq}}/N_{\text{total}}$	Δp_{full}	$\Delta p_{\text{no-gen}}$	Δp_{gen}	Retention
ARC-Easy	Gemma2-2B	24/296	0.49 ± 0.39	0.44 ± 0.29	0.11 ± 0.44	0.88
	Gemma2-9B	23/476	0.88 ± 0.17	0.61 ± 0.33	0.83 ± 0.30	0.69
	Llama3.1-8B	9/53	0.46 ± 0.23	0.30 ± 0.19	0.28 ± 0.17	0.64
	Qwen3-1.7B	14/382	0.81 ± 0.37	0.81 ± 0.35	0.00 ± 0.04	1.00
	Qwen3-8B	14/402	0.97 ± 0.14	0.97 ± 0.14	0.19 ± 0.25	1.00
LogiQA	Gemma2-2B	18/206	0.18 ± 0.21	0.18 ± 0.08	-0.04 ± 0.31	1.00
	Gemma2-9B	22/360	0.41 ± 0.21	0.20 ± 0.16	0.39 ± 0.21	0.49
	Llama3.1-8B	10/48	0.20 ± 0.15	0.06 ± 0.14	0.13 ± 0.10	0.30
	Qwen3-1.7B	14/277	0.53 ± 0.45	0.53 ± 0.44	0.02 ± 0.09	1.00
	Qwen3-8B	16/317	0.70 ± 0.37	0.70 ± 0.37	0.19 ± 0.28	1.00
OpenbookQA	Gemma2-2B	23/262	0.31 ± 0.31	0.31 ± 0.23	-0.07 ± 0.35	1.00
	Gemma2-9B	20/495	0.66 ± 0.28	0.41 ± 0.31	0.48 ± 0.38	0.63
	Llama3.1-8B	9/45	0.32 ± 0.23	0.22 ± 0.21	0.20 ± 0.16	0.68
	Qwen3-1.7B	17/278	0.61 ± 0.46	0.61 ± 0.44	0.01 ± 0.06	1.00
	Qwen3-8B	15/346	0.82 ± 0.34	0.82 ± 0.34	0.11 ± 0.23	1.00

3. **Patched:** Replace $\mathbf{r}_{\text{pred}}$ with only the $S_\cap$ contribution: $\mathbf{r}_{\text{patch}} = \text{Dec}(\mathbf{a}_{\text{pred}} \odot \mathbf{m}_{S_\cap})$, where $\mathbf{m}_{S_\cap}$ is a binary mask selecting only shared features; measure p_{patch} .

Recall that $\mathbf{r}_{\text{pred}} \in \mathbb{R}^{\text{model}}$ denotes the residual stream activation and $\mathbf{a}_{\text{pred}} = \text{Enc}(\mathbf{r}_{\text{pred}})$ represents the corresponding SAE latent activation (§3.5).

We quantify sufficiency via the *recovery* metric:

$$\text{Recovery} = \frac{p_{\text{patch}} - p_{\text{zero}}}{p_{\text{base}} - p_{\text{zero}}} \quad (12)$$

A recovery of 1.0 indicates that $S_\cap$ fully restores the original answer probability. To verify that $S_\cap$ is *specifically* sufficient (not just any k features), we compare against a random baseline: for each example, we sample $k = |S_\cap|$ random features from S_{pred} , compute recovery, and repeat 20 times. If $S_\cap$ captures causally important features, its recovery should noticeably exceed the random baseline.

Results. Table 8 presents the sufficiency results across all model–dataset combinations. On average, patching only $S_\cap$ achieves a recovery of 1.05 ± 0.09 , indicating that these shared features not only restore but sometimes *exceed* the original answer probability, likely because the patched representation removes interfering non-shared features. In contrast, random features of equal size achieve only 0.45 ± 0.35 recovery, with a substantial gap of 0.60.

Together, these findings suggest that $S_\cap$ is **both necessary and sufficient**: removing it disrupts the answer probability, and retaining it alone largely recovers the answer compared to the random baseline, indicating that the shared concepts between prediction and CoT passes capture the causally relevant reasoning.

E ALTERNATIVES TO MEASURING FAITHFULNESS

Lack of Ground-truth. We want to highlight that there is no gold metric for measuring CoT faithfulness in the absence of a ground truth for the model’s underlying reasoning (Parcalabescu & Frank, 2024; Sun et al., 2026). Therefore, the existing metrics are used only as proxies for faithfulness.

Infeasibility of Human Evaluation. The effectiveness of human evaluation depends on the assumption that people can meaningfully assess whether a CoT faithfully represents a model’s internal reasoning. Yet this assumption becomes questionable in the context of modern LLMs. With billions of parameters composing intricate, nonlinear interactions that map inputs to outputs, the

Table 8: Sufficiency test results: patching only $S_\cap$ features into zeroed activations. Recovery measures how much of the original answer probability is restored. $S_\cap$ noticeably outperforms random features of the same size.

Dataset	Model	$S_\cap$ Recov.	Rand Recov.	Gap	$S_\cap > \text{Rand}$
LogiQA	Llama-3.1-8B	1.018	0.919	0.099	63.5%
	Qwen3-1.7B	0.978	0.068	0.910	83.9%
	Qwen3-8B	1.147	0.081	1.066	97.3%
	Gemma-2-2B	1.000	0.711	0.289	65.5%
	Gemma-2-9B	0.950	0.750	0.200	60.5%
OpenbookQA	Llama-3.1-8B	1.079	0.939	0.140	75.5%
	Qwen3-1.7B	1.213	0.140	1.073	92.8%
	Qwen3-8B	1.226	0.125	1.101	97.3%
	Gemma-2-2B	0.954	0.782	0.172	58.5%
	Gemma-2-9B	0.992	0.447	0.545	90.0%
ARC-Easy	Llama-3.1-8B	1.075	0.923	0.152	88.5%
	Qwen3-1.7B	1.085	0.084	1.001	97.7%
	Qwen3-8B	1.045	0.077	0.967	100.0%
	Gemma-2-2B	0.996	0.502	0.494	78.0%
	Gemma-2-9B	0.950	0.136	0.814	96.5%

actual computational pathways are essentially opaque to human observers. We currently lack both the tools and the theoretical understanding required to trace how these vast networks collectively produce any given prediction. As a result, human judgments of CoT faithfulness, though intuitively compelling, may be methodologically unsound given our present limited visibility into LLM internals.

Contextual Metrics. We consider contextual CoT faithfulness metrics from (Lanham et al., 2023), which are based on context perturbation: (i) **Early Answering**, which truncates the CoT before the original answer is produced and forces an early response; (ii) **Adding Mistake**, which injects an error into one CoT step and regenerates the remaining steps; (iii) **Paraphrasing**, which restates the initial portion of the CoT while preserving its meaning and regenerates the rest; and (iv) **Filler Token**, which substitutes the CoT with meaningless symbols such as ellipses. The results of these metrics are binary, i.e., faithful or unfaithful, which is coarse-grained compared to our proposed metrics, which range from 0 to 1. Therefore, we follow Zaman & Srivastava (2025) and instead measure the prediction score $\hat{z}_i$ changes, ranging from -1 to 1.

$$\hat{z}_i = \frac{\exp(p_\theta(L_i | x))}{\sum_{L_j \in L} \exp(p_\theta(L_j | x))} \quad (13)$$

where L denotes the set of possible labels and x represents the input text.

The hint injection method (Turpin et al., 2023) is intentionally excluded. It produces binary labels (faithful vs. unfaithful), whereas our metric yields continuous values. As a result, we cannot compute a meaningful rank (Spearman) correlation.

Parametric Metric. Additionally, we follow Tutek et al. (2025); Sun et al. (2026) to assess faithfulness through interventions on the parametric knowledge of model $\mathcal{M}$. Specifically, given an input x , the target CoT c is first decomposed into T reasoning steps $c = (s_1, \dots, s_T)$. FUR provides a continuous answer as to whether a CoT is faithful by measuring the extent to which unlearning the knowledge used by models in CoT steps causes model to change their output prediction:

$$\text{FUR} = \frac{\sum_{i=1}^T \mathbb{I}[s_i \text{ such that } y \neq y^{(i)*}]}{T} \quad (14)$$

where s_i is the i -th reasoning step, $y^{(i)*}$ denotes the prediction made by $\mathcal{M}^{(i)*}$ after unlearning the knowledge involved in the i -th reasoning step, and $\mathbb{I}(\cdot)$ represents the indicator function.

Circuit-guided Internal-External Discrepancy (CIE). CIE compares the model’s internal computation against its displayed reasoning (Shen et al., 2026). For each trace, an internal representation is created by tracing sentence-level attribution circuits over informative tokens, and an external representation from the hidden states of the reasoning sentences is extracted and organized into sentence graphs. Their structural discrepancy is measured via the Fused Gromov-Wasserstein distance. A larger CIE score indicates greater internal-external divergence and thus lower faithfulness.

Table 9: Existing faithfulness metrics results with Llama3.1-8B on LogiQA.

Metric	Mean	Std	95% CI
Early Answering (Zaman & Srivastava, 2025)	0.6469	0.1119	[0.6371, 0.6567]
Filler Tokens (Zaman & Srivastava, 2025)	0.6476	0.1118	[0.6378, 0.6574]
Adding Mistakes (Zaman & Srivastava, 2025)	0.3953	0.1259	[0.3843, 0.4064]
Paraphrasing (Zaman & Srivastava, 2025)	-0.0447	0.0568	[-0.0497, -0.0398]
FUR (Tutek et al., 2025)	0.2217	0.1710	[0.1876, 0.2558]
CIE (Shen et al., 2026)	0.6160	0.0098	[0.6139, 0.6181]
CC-SAE	0.9142	0.0251	[0.9120, 0.9164]
Jaccard	0.6145	0.0588	[0.6093, 0.6196]
Recall	0.8580	0.0456	[0.8540, 0.8620]
Δp	0.1450	0.0990	[0.1400, 0.1585]

Relationship to Prior Metrics. This comparison serves as a sanity check to verify that our metric captures information not already redundant with existing approaches, rather than establishing superiority over prior methods, because no ground-truth can be provided as discussed in the prior paragraph. We only conduct a small-scale comparison with Llama3.1-8B on LogiQA, as most metrics are computationally too expensive for extensive evaluation (for example, per-instance faithfulness using FUR with Llama3.1-8B takes approximately 25 minutes). Therefore, we defer a more systematic and comprehensive comparison to future work. Table 9 reports correlations between our metrics and existing faithfulness measures. We observe slightly positive Spearman correlations between Δp and contextual metrics ($\rho \approx 0.10$), while noting weak-to-moderate correlations with the parametric faithfulness metric FUR ($\rho = 0.25$) and the MI-based metric CIE score ($\rho = 0.30$). The weak-to-moderate correlations are *expected*, not a deficiency of either approach, aligned with Parcalabescu & Frank (2024); Gur-Arieh et al. (2026). Prior metrics probe different dimensions: contextual metrics test sensitivity to textual perturbations; FUR tests dependence on unlearned knowledge; CIE measures circuit-trace discrepancy. Our Δp tests whether *shared internal features* causally drive the answer. These dimensions are conceptually orthogonal, i.e., a CoT can be contextually robust yet rely on different internal features, or vice versa. The low correlations confirm that our metric captures a complementary signal rather than redundantly measuring what existing methods already measure. Together, these metric families provide a more comprehensive assessment than any single approach.

F CAUSAL VALIDATION VIA CONDITIONS

F.1 PREDICTION PASS ABLATION

F.1.1 THREE BASE CONDITIONS

Table 10 shows the complete raw values of three base conditions (Conditions A–C) for all models and datasets. Conditions A–C yield $\Delta p \approx 0$ across all settings, while $\Delta p(S_{\cap})$ ranges from 0.1 to 0.9. This demonstrates that the causal effect is specific to intersection features rather than an artifact of set size or ablation magnitude. Even norm-matched random sampling (Condition C) fails to replicate the effect, indicating that $S_{\cap}$ captures semantically meaningful features beyond what magnitude alone would predict.

We further observe that Llama-3.1-8B conditions are magnitude-wise notably smaller than those of other models and exhibit markedly lower variance than those of the other models, indicating that random ablations have essentially no effect on Llama-3.1-8B.

Table 10: Prediction pass conditions and top- k upper bound comparison against $\Delta p(S_\cap)$ across models.

	Model	Condition A	Condition B	Condition C	top- k upper bound
LogiQA	Llama-3.1-8B	0.045 ± 0.150	0.068 ± 0.160	0.065 ± 0.141	0.156 ± 0.204
	Qwen3-1.7B	0.142 ± 0.311	0.178 ± 0.334	0.110 ± 0.308	0.530 ± 0.442
	Qwen3-8B	0.156 ± 0.302	0.172 ± 0.308	0.141 ± 0.319	0.678 ± 0.378
	Gemma-2-2B	0.034 ± 0.105	0.034 ± 0.130	0.021 ± 0.121	0.222 ± 0.100
	Gemma-2-9B	0.022 ± 0.109	0.035 ± 0.112	0.052 ± 0.124	0.378 ± 0.203
OpenbookQA	Llama-3.1-8B	0.075 ± 0.167	0.147 ± 0.192	0.140 ± 0.169	0.286 ± 0.289
	Qwen3-1.7B	0.135 ± 0.323	0.165 ± 0.349	0.072 ± 0.279	0.588 ± 0.464
	Qwen3-8B	0.164 ± 0.323	0.170 ± 0.327	0.166 ± 0.324	0.804 ± 0.347
	Gemma-2-2B	0.057 ± 0.122	0.057 ± 0.129	0.039 ± 0.111	0.351 ± 0.213
	Gemma-2-9B	0.035 ± 0.138	0.071 ± 0.150	0.067 ± 0.121	0.614 ± 0.273
ARC-Easy	Llama-3.1-8B	0.098 ± 0.173	0.195 ± 0.203	0.173 ± 0.184	0.366 ± 0.330
	Qwen3-1.7B	0.148 ± 0.311	0.189 ± 0.348	0.066 ± 0.224	0.766 ± 0.399
	Qwen3-8B	0.176 ± 0.313	0.178 ± 0.309	0.139 ± 0.282	0.935 ± 0.195
	Gemma-2-2B	0.082 ± 0.154	0.088 ± 0.167	0.076 ± 0.181	0.494 ± 0.270
	Gemma-2-9B	0.046 ± 0.141	0.066 ± 0.122	0.069 ± 0.127	0.819 ± 0.218
GSM8K	Llama-3.1-8B	0.019 ± 0.066	0.022 ± 0.061	0.023 ± 0.078	0.036 ± 0.086
	Qwen3-1.7B	0.063 ± 0.228	0.059 ± 0.220	0.034 ± 0.226	0.192 ± 0.269
	Qwen3-8B	0.091 ± 0.255	0.096 ± 0.258	0.066 ± 0.257	0.188 ± 0.299
	Gemma-2-2B	0.012 ± 0.073	0.017 ± 0.081	0.017 ± 0.080	0.153 ± 0.151
	Gemma-2-9B	0.029 ± 0.121	0.040 ± 0.140	0.031 ± 0.127	0.169 ± 0.209

F.1.2 CONDITION C: NORM-MATCHED CONTROL

To control for the magnitude of intervention rather than merely the number of ablated features, we introduce a norm-matched condition. Let $\mathbf{a}_{S_\cap}$ denote the SAE activation vector restricted to the intersection features $S_\cap$, and let $|\text{Dec}(\mathbf{a}_{S_\cap})|_2$ be the L2 norm of the corresponding contribution in residual stream space. For each random sample, we construct a feature set from $S_{\text{pred}} \setminus S_\cap$ by iteratively adding features in random order until the decoded ablation vector matches the target norm within a tolerance τ :

$$\text{Find } S_{\text{rand}} \subseteq S_{\text{pred}} \setminus S_\cap \text{ s.t. } ||\text{Dec}(\mathbf{a}_{S_{\text{rand}}})|_2 - |\text{Dec}(\mathbf{a}_{S_\cap})|_2| \leq \tau \cdot |\text{Dec}(\mathbf{a}_{S_\cap})|_2$$

This ensures that the **Condition C** ablation removes approximately the same amount of activation magnitude from the residual stream as ablating $S_\cap$, isolating the effect of which features are ablated from how much is removed.

F.1.3 TWO SPECIFIC CONDITIONS FOR SANITY CHECKS

Condition D. Table 11 shows that the incremental effect is negligible across all model-dataset combinations, suggesting that $S_\cap$ alone is sufficient for the causal effect. Features that are active only in the prediction pass but are not reused from CoT contribute minimally to the model’s final answer.

Condition E. Table 12 illustrates that ablating features exclusive to the CoT pass ($S_{\text{cot}} \setminus S_\cap$) yields $\Delta p \approx 0$ across all model-dataset combinations. In contrast, ablating intersection features $S_\cap$ produces substantial probability drops (Section 5.3). In other words, features activated only during CoT generation do not causally influence the final prediction. This indicates that the causal effect of $S_\cap$ is not an artifact of position-level activation overlap, but reflects genuine concept grounding between CoT reasoning and prediction. Condition E serves as an *implementation sanity check*, i.e., verifying that our hook correctly targets only the specified features and that the SAE encode-decode cycle introduces no spurious effects, rather than being evidence for any causal claim.

Table 11: Condition D analysis: incremental effect of ablating k pred-only features on top of $S_\cap$. Near-zero values confirm that $S_\cap$ is complete and prediction-only features contribute minimally. The averaged normalized ratio $\Delta p(S_\cap)/\Delta p(S_{\text{pred}})$ is 0.907 ± 0.108 , indicating negligible floor effect.

Dataset	Model	Incremental Effect
LogiQA	Llama-3.1-8B	$-6.8e^{-02} \pm 1.2e^{-01}$
	Qwen3-1.7B	$-3.5e^{-02} \pm 1.6e^{-01}$
	Qwen3-8B	$-2.8e^{-02} \pm 1.4e^{-01}$
	Gemma-2-2B	$-7.8e^{-02} \pm 1.2e^{-01}$
	Gemma-2-9B	$-3.1e^{-02} \pm 7.0e^{-02}$
OpenbookQA	Llama-3.1-8B	$-7.9e^{-02} \pm 1.4e^{-01}$
	Qwen3-1.7B	$-5.9e^{-02} \pm 2.3e^{-01}$
	Qwen3-8B	$-3.2e^{-02} \pm 1.6e^{-01}$
	Gemma-2-2B	$-6.3e^{-02} \pm 1.3e^{-01}$
	Gemma-2-9B	$-4.9e^{-02} \pm 9.7e^{-02}$
ARC-Easy	Llama-3.1-8B	$-1.1e^{-01} \pm 1.9e^{-01}$
	Qwen3-1.7B	$-6.4e^{-02} \pm 2.4e^{-01}$
	Qwen3-8B	$-4.3e^{-02} \pm 1.9e^{-01}$
	Gemma-2-2B	$-1.2e^{-01} \pm 1.9e^{-01}$
	Gemma-2-9B	$-7.3e^{-02} \pm 1.3e^{-01}$
GSM8K	Llama-3.1-8B	$-6.3e^{-03} \pm 2.5e^{-02}$
	Qwen3-1.7B	$-1.3e^{-02} \pm 8.5e^{-02}$
	Qwen3-8B	$-1.0e^{-02} \pm 8.5e^{-02}$
	Gemma-2-2B	$-3.9e^{-02} \pm 9.0e^{-02}$
	Gemma-2-9B	$-3.7e^{-02} \pm 8.5e^{-02}$

Table 12: Condition E analysis: ablating CoT-only features ($S_{\text{cot}} \setminus S_\cap$) yields $\Delta p \approx 0$, confirming these features do not affect the prediction. In contrast, ablating $S_\cap$ causes substantial probability drops (Table 10).

Dataset	Model	$\bar{\Delta p}$	σ
LogiQA	Llama-3.1-8B	0.030	0.122
	Qwen3-1.7B	0.013	0.306
	Qwen3-8B	0.016	0.275
	Gemma-2-2B	0.023	0.094
	Gemma-2-9B	0.030	0.105
OpenbookQA	Llama-3.1-8B	0.072	0.172
	Qwen3-1.7B	0.081	0.288
	Qwen3-8B	0.012	0.305
	Gemma-2-2B	0.036	0.096
	Gemma-2-9B	0.047	0.121
ARC-Easy	Llama-3.1-8B	0.019	0.192
	Qwen3-1.7B	0.096	0.279
	Qwen3-8B	0.014	0.267
	Gemma-2-2B	0.066	0.139
	Gemma-2-9B	0.052	0.114
GSM8K	Llama-3.1-8B	0.015	0.060
	Qwen3-1.7B	0.047	0.243
	Qwen3-8B	0.081	0.273
	Gemma-2-2B	0.008	0.074
	Gemma-2-9B	0.030	0.130

Table 13: CoT pass conditions and top- k upper bound comparison against $\Delta p(S_\cap)$ across models.

	Model	Condition A	Condition B	Condition C	top- k upper bound
LogiQA	Llama-3.1-8B	0.057 ± 0.143	0.072 ± 0.172	0.061 ± 0.170	0.187 ± 0.258
	Qwen3-1.7B	0.052 ± 0.212	0.038 ± 0.185	0.076 ± 0.265	0.521 ± 0.468
	Qwen3-8B	0.111 ± 0.287	0.074 ± 0.222	0.165 ± 0.296	0.739 ± 0.402
	Gemma-2-2B	0.019 ± 0.109	0.042 ± 0.159	0.023 ± 0.144	0.238 ± 0.328
	Gemma-2-9B	0.013 ± 0.088	0.018 ± 0.093	0.037 ± 0.108	0.374 ± 0.273
OpenbookQA	Llama-3.1-8B	0.077 ± 0.147	0.074 ± 0.140	0.076 ± 0.155	0.318 ± 0.333
	Qwen3-1.7B	0.095 ± 0.261	0.101 ± 0.275	0.125 ± 0.305	0.638 ± 0.453
	Qwen3-8B	0.193 ± 0.352	0.118 ± 0.267	0.185 ± 0.297	0.885 ± 0.281
	Gemma-2-2B	0.010 ± 0.109	0.012 ± 0.133	0.017 ± 0.109	0.292 ± 0.305
	Gemma-2-9B	0.027 ± 0.108	0.041 ± 0.127	0.069 ± 0.144	0.620 ± 0.345
ARC-Easy	Llama-3.1-8B	0.069 ± 0.143	0.074 ± 0.140	0.102 ± 0.158	0.392 ± 0.359
	Qwen3-1.7B	0.107 ± 0.267	0.103 ± 0.265	0.130 ± 0.311	0.706 ± 0.428
	Qwen3-8B	0.166 ± 0.329	0.089 ± 0.231	0.149 ± 0.274	0.919 ± 0.238
	Gemma-2-2B	0.018 ± 0.118	0.015 ± 0.148	0.011 ± 0.126	0.427 ± 0.358
	Gemma-2-9B	0.002 ± 0.106	0.020 ± 0.117	0.041 ± 0.097	0.755 ± 0.256
GSM8K	Llama-3.1-8B	0.036 ± 0.087	0.052 ± 0.107	0.030 ± 0.136	0.104 ± 0.241
	Qwen3-1.7B	0.145 ± 0.310	0.153 ± 0.316	0.118 ± 0.299	0.564 ± 0.411
	Qwen3-8B	0.226 ± 0.367	0.202 ± 0.333	0.081 ± 0.232	0.739 ± 0.369
	Gemma-2-2B	0.011 ± 0.082	0.020 ± 0.096	0.017 ± 0.093	0.190 ± 0.262
	Gemma-2-9B	0.010 ± 0.065	0.024 ± 0.092	0.016 ± 0.082	0.214 ± 0.299

F.2 CoT PASS ABLATION

F.2.1 SETUP AND CONDITIONS

Similar to the prediction pass ablation (§3.5), the CoT pass ablation is calculated as follows:

$$\mathbf{r}'_{\text{cot}} = \mathbf{r}_{\text{cot}} - \mathbf{W}_{\text{dec}}[\mathbf{a}_{\text{cot}} \odot \mathbf{m}_{S_\cap}] \quad (15)$$

where $\mathbf{m}_{S_\cap}$ is a binary mask that selects only the shared concepts. We then measure the change in answer probability:

$$\Delta p = P(\hat{y} | \mathbf{r}_{\text{cot}}) - P(\hat{y} | \mathbf{r}'_{\text{cot}}) \quad (16)$$

Analogously, we evaluate four base conditions for sanity checks on the CoT pass (§5.2.1):

- **Condition A:** Sample k random features from $(S_{\text{pred}} \cup S_{\text{cot}}) \setminus S_\cap$. It aims to test whether intersection features influence the answer probability more than other features that are active in either pass.
- **Condition B:** Sample k random features from $S_{\text{cot}} \setminus S_\cap$. It aims to test whether the causal effect is specific to features reused by CoT.
- **Condition C:** Norm-matched sampling from $S_{\text{cot}} \setminus S_\cap$, controlling for the L_2 norm of the ablation vector (Appendix F.1.2). It aims to test whether the effects are not driven by ablation magnitude.
- **Condition D:** Ablate $S_\cap$ together with k additional random features from $S_{\text{cot}} \setminus S_\cap$. If $S_\cap$ captures the causally important features, adding CoT-only features should not noticeably increase Δp beyond ablating $S_\cap$ alone, which serves as a completeness check.

Additionally, we define and compute a **top- k upper bound** (magnitude-ranked baseline), i.e., ablating the k features in S_{cot} with the largest activation magnitude.

F.2.2 RESULTS

Condition A–C. Table 13 reveals that across all models and datasets, Conditions A–C yield substantially lower Δp values compared to the true causal effect of $S_\cap$ (Table 14). For instance, on

Table 14: CoT ablation analysis. **Completeness:** $\Delta p(S_{\cap}) \approx \Delta p(S_{\text{pred}})$ on CoT pass indicates prediction-only features contribute minimally. **Cross-path:** similar Δp on both paths confirms $S_{\cap}$ is equally important for reasoning and prediction.

Dataset	Model	Completeness (CoT Pass)			Cross-path	
		$\Delta p(S_{\cap})$	$\Delta p(S_{\text{pred}})$	Ratio	Δp_{cot}	Δp_{direct}
LogiQA	Gemma2-2B	0.22 ± 0.32	0.24 ± 0.34	0.95 ± 0.84	0.22 ± 0.32	0.20 ± 0.12
	Gemma2-9B	0.37 ± 0.29	0.39 ± 0.30	0.96 ± 0.72	0.37 ± 0.29	0.38 ± 0.21
	Llama3.1-8B	0.16 ± 0.24	0.20 ± 0.26	0.83 ± 1.22	0.16 ± 0.24	0.15 ± 0.20
	Qwen3-1.7B	0.52 ± 0.47	0.58 ± 0.47	0.90 ± 3.28	0.52 ± 0.47	0.53 ± 0.44
	Qwen3-8B	0.72 ± 0.41	0.76 ± 0.39	0.96 ± 3.03	0.72 ± 0.41	0.67 ± 0.38
OpenbookQA	Gemma2-2B	0.29 ± 0.32	0.31 ± 0.32	0.96 ± 0.66	0.29 ± 0.32	0.32 ± 0.23
	Gemma2-9B	0.61 ± 0.35	0.63 ± 0.35	0.96 ± 0.75	0.61 ± 0.35	0.61 ± 0.28
	Llama3.1-8B	0.28 ± 0.32	0.35 ± 0.35	0.81 ± 0.75	0.28 ± 0.32	0.28 ± 0.29
	Qwen3-1.7B	0.70 ± 0.44	0.74 ± 0.42	0.94 ± 3.45	0.70 ± 0.44	0.61 ± 0.46
	Qwen3-8B	0.83 ± 0.34	0.88 ± 0.29	0.94 ± 3.23	0.83 ± 0.34	0.80 ± 0.35
ARC-Easy	Gemma2-2B	0.42 ± 0.36	0.44 ± 0.36	0.96 ± 0.42	0.42 ± 0.36	0.47 ± 0.30
	Gemma2-9B	0.77 ± 0.26	0.79 ± 0.25	0.97 ± 0.86	0.77 ± 0.26	0.82 ± 0.22
	Llama3.1-8B	0.37 ± 0.35	0.44 ± 0.38	0.83 ± 0.71	0.37 ± 0.35	0.36 ± 0.33
	Qwen3-1.7B	0.78 ± 0.39	0.85 ± 0.34	0.93 ± 3.97	0.78 ± 0.39	0.79 ± 0.39
	Qwen3-8B	0.86 ± 0.31	0.92 ± 0.23	0.94 ± 3.93	0.86 ± 0.31	0.94 ± 0.19

ARC-Easy, Qwen3-8B achieves $\Delta p(S_{\cap}) = 0.86$ while all three baseline conditions remain below 0.17, a gap exceeding $5\times$. Similarly, Gemma-2-9B shows $\Delta p(S_{\cap}) = 0.77$ versus baseline effects of merely 0.002–0.041. The consistently low baseline values across union-based (A), prediction-exclusive (B), and norm-matched (C) conditions provide evidence that the causal effect is specific to intersection features rather than arbitrary subsets. Crucially, similar to the prediction pass (§5.2.1), the top- k upper bound is marginally above $\Delta p(S_{\cap})$, and in several cases $S_{\cap}$ achieves comparable or even higher effects. This parallels our findings for the prediction pass (Table 10): intersection features, selected by conceptual reuse across reasoning paths rather than activation magnitude, are nearly as causally potent as the best magnitude-selected subset. The consistency across both CoT and prediction passes strengthens the conclusion that $S_{\cap}$ captures genuinely privileged causal structure.

Condition D: Completeness of $S_{\cap}$. We compare the causal effect of intersection features $\Delta p(S_{\cap})$ against all predictive features $\Delta p(S_{\text{pred}})$ on the CoT path. Across all model-dataset combinations, the ratio $\Delta p(S_{\cap})/\Delta p(S_{\text{pred}})$ averages 0.88 ± 0.11 , indicating that $S_{\cap}$ captures the vast majority of the causal effect. This suggests that *prediction-only features contribute minimally to CoT reasoning*.

Cross-path Comparison. We further examine whether $S_{\cap}$ is equally important for reasoning versus prediction by comparing Δp_{cot} and Δp_{pred} . Table 14 shows that both paths exhibit similar causal effects, suggesting that the model can leverage these features through either reasoning or direct retrieval. Additionally, the Pearson correlation between the ablation results of the two passes is 0.64 on average, suggesting that the two paths are moderately similar.

G COMPARISON BETWEEN CORRELATIONAL AND CAUSAL METRICS

G.1 CORRELATIONAL ANALYSIS

Table 15 illustrates that the averaged layer-wise Spearman correlation coefficient is generally low ($r=0.0$ – 0.2), indicating a negligible positive monotonic relationship between Δp and correlational metrics.

Table 15: Averaged layer-wise Spearman correlation between correlational metrics (CC-SAE, Jaccard, and Recall) and causal metric (Δp).

	Model	CC-SAE	Jaccard	Recall
LogiQA	Llama-3.1-8B	+0.04 $\pm$ 0.08	+0.06 $\pm$ 0.08	+0.03 $\pm$ 0.09
	Qwen3-1.7B	+0.14 $\pm$ 0.09	+0.05 $\pm$ 0.07	+0.13 $\pm$ 0.05
	Qwen3-8B	+0.07 $\pm$ 0.13	+0.00 $\pm$ 0.09	+0.11 $\pm$ 0.14
	Gemma-2-2B	+0.14 $\pm$ 0.15	+0.14 $\pm$ 0.15	+0.10 $\pm$ 0.15
	Gemma-2-9B	+0.09 $\pm$ 0.11	+0.08 $\pm$ 0.10	+0.08 $\pm$ 0.08
OpenbookQA	Llama-3.1-8B	+0.09 $\pm$ 0.08	+0.12 $\pm$ 0.09	+0.12 $\pm$ 0.09
	Qwen3-1.7B	+0.17 $\pm$ 0.11	+0.08 $\pm$ 0.07	+0.19 $\pm$ 0.07
	Qwen3-8B	+0.12 $\pm$ 0.20	+0.02 $\pm$ 0.13	+0.12 $\pm$ 0.16
	Gemma-2-2B	+0.10 $\pm$ 0.12	+0.08 $\pm$ 0.10	+0.00 $\pm$ 0.13
	Gemma-2-9B	+0.21 $\pm$ 0.19	+0.18 $\pm$ 0.16	+0.17 $\pm$ 0.15
ARC-Easy	Llama-3.1-8B	+0.10 $\pm$ 0.13	+0.11 $\pm$ 0.14	+0.10 $\pm$ 0.13
	Qwen3-1.7B	+0.11 $\pm$ 0.12	+0.07 $\pm$ 0.09	+0.08 $\pm$ 0.09
	Qwen3-8B	+0.08 $\pm$ 0.10	+0.02 $\pm$ 0.10	+0.07 $\pm$ 0.11
	Gemma-2-2B	+0.14 $\pm$ 0.19	+0.11 $\pm$ 0.13	+0.03 $\pm$ 0.14
	Gemma-2-9B	+0.22 $\pm$ 0.19	+0.19 $\pm$ 0.14	+0.12 $\pm$ 0.12
GSM8K	Llama-3.1-8B	+0.04 $\pm$ 0.11	+0.11 $\pm$ 0.13	+0.07 $\pm$ 0.11
	Qwen3-1.7B	+0.03 $\pm$ 0.06	+0.03 $\pm$ 0.06	+0.04 $\pm$ 0.06
	Qwen3-8B	+0.01 $\pm$ 0.07	+0.04 $\pm$ 0.07	-0.02 $\pm$ 0.08
	Gemma-2-2B	+0.13 $\pm$ 0.15	+0.12 $\pm$ 0.15	+0.11 $\pm$ 0.13
	Gemma-2-9B	+0.09 $\pm$ 0.13	+0.05 $\pm$ 0.09	+0.02 $\pm$ 0.08

G.2 RELATIONSHIP BETWEEN CONCEPT ALIGNMENT AND CAUSAL EFFECT

We further examine the relationship between CC-SAE alignment and causal effect (Δp) at the instance level using per-dataset median thresholds to account for cross-dataset difficulty differences. Table 16) illustrates that:

- $\bar{P}(A|B) = 0.53$ (range: 0.47–0.59): When causal effect is above median, alignment is equally likely to be above or below median.
- $\bar{P}(B|A) = 0.54$ (range: 0.47–0.59): High alignment does not reliably predict high causal effect.
- $\overline{\text{Lift}} = 1.08$ (range: 0.94–1.19): High alignment increases the probability of high causal effect by only 8% relative to the base rate.
- $\bar{\phi} = 0.08$ (range: -0.06–0.18): The association between binarized alignment and causal effect is weak.

These findings suggest that CC-SAE measures *representational overlap* between SAE features and model activations, which is conceptually distinct from *causal importance*. A feature set may align well with model activations (high CC-SAE) yet contribute minimally to the output when ablated, e.g., due to redundant computational pathways. Conversely, causally critical features may be partially missed by the SAE (low CC-SAE, high Δp).

H CoT POST-HOC RATIONALIZATION EXAMPLES

Table 17 presents a CoT instance yielding a Δp of 0.22. The illustrated CoT glosses over a critical logical gap: in the original prompt, meat prices lack any logical connection to grain, vegetables, or cooking oil. The post-hoc explanation completely ignores this flaw and instead praises option B for “elegantly connecting the hidden variables”.

A similar pattern appears in Table 18 ($\Delta p = 0.18$), where the CoT opens with a measurement that is not among the choices (“the student should measure the type of soil”), declares every option

Table 16: Association between Concept Alignment and Causal Effect under Per-Dataset Median Thresholds. θ^* , τ^* : per-dataset median thresholds for CC-SAE and Δp . ϕ : Pearson correlation coefficient for 2x2 table. Normalization ensures comparable metrics across datasets with different difficulty levels.

Model	Dataset	Thresholds		Conditional Prob.			Association	
		θ^*	τ^*	$P(A B)$	$P(B \neg A)$	$P(B A)$	Lift	ϕ
Gemma-2-2B	ARC-Easy	0.88	0.52	0.52	0.48	0.52	1.04	0.04
	GSM8K	0.78	0.10	0.55	0.44	0.56	1.12	0.12
	LogiQA	0.89	0.21	0.53	0.47	0.53	1.07	0.07
	OpenbookQA	0.88	0.28	0.51	0.48	0.52	1.04	0.04
Gemma-2-9B	ARC-Easy	0.90	0.91	0.59	0.41	0.59	1.18	0.18
	GSM8K	0.66	0.07	0.47	0.53	0.47	0.94	-0.06
	LogiQA	0.89	0.32	0.52	0.47	0.53	1.06	0.06
	OpenbookQA	0.88	0.67	0.56	0.43	0.57	1.14	0.14
Llama-3.1-8B	ARC-Easy	0.89	0.29	0.55	0.45	0.55	1.10	0.10
	GSM8K	0.82	0.01	0.51	0.49	0.51	1.03	0.03
	LogiQA	0.93	0.12	0.52	0.48	0.52	1.03	0.03
	OpenbookQA	0.90	0.21	0.53	0.47	0.53	1.06	0.06
Qwen3-1.7B	ARC-Easy	0.82	1.00	0.59	0.41	0.59	1.18	0.18
	GSM8K	0.65	0.07	0.51	0.49	0.51	1.02	0.02
	LogiQA	0.84	0.56	0.53	0.47	0.53	1.06	0.06
	OpenbookQA	0.79	0.92	0.56	0.44	0.56	1.12	0.12
Qwen3-8B	ARC-Easy	0.85	1.00	0.55	0.44	0.57	1.14	0.13
	GSM8K	0.62	0.04	0.50	0.50	0.50	1.01	0.01
	LogiQA	0.87	0.90	0.53	0.45	0.55	1.10	0.10
	OpenbookQA	0.83	0.99	0.55	0.44	0.57	1.13	0.13

unrelated to the question, and then selects one of them regardless. Nothing in the displayed reasoning provides a basis for that selection.

In Table 19 ($\Delta p = 0.16$), the CoT correctly states that skills are “learned through experience” and that height is “determined by genetics”. Both point to (B), yet the CoT answers (D): the answer does not follow from what the CoT itself says.

I SAE CONFIGURATION ABLATION STUDY

I.1 SPARSITY AND WIDTH ABLATION

Apart from Gemma-2-2B (§5.4), we further conduct an ablation study on the remaining models. For Gemma-2-9B, similarly, we use an increased width ($16k \rightarrow 131k$) and an increased sparsity ($\ell_0 = 129 \rightarrow 50$). For Qwen3 models, as QwenScope only provides SAE variants with different ℓ_0 values, we only consider increased sparsity ($\ell_0 = 100 \rightarrow 50$).¹¹

Figure 15–18 display layer-wise Δ across three SAE configurations. Table 21–23 further show the Δ raw values and Pearson’s r correlation values across three SAE configurations.

I.2 ℓ_0 VARIATION ACROSS LAYERS

Qwen3 and Llama3.1 models use TopK-ReLU activations, which therefore have the same ℓ_0 value across layers. However, Gemma-2 models employ JumpReLU and the ℓ_0 values vary across layers. As a sanity check, we investigate how ℓ_0 correlates with Δp . In our experiments, we consider the ℓ_0 value at the first layer as the anchor and consistently choose the closest ℓ_0 from the available ℓ_0 set at the other layers to largely maintain the ℓ_0 consistency.¹²

¹¹<https://huggingface.co/collections/Qwen/qwen-scope>

¹²<https://huggingface.co/google/gemma-scope>

Table 17: Qualitative analysis of CoT faithfulness: Gemma-2-9B on LogiQA.

Gemma-2-2B	LogiQA	$\Delta p = 0.22$
Question. Professionals predict that if grain prices remain stable, vegetable prices will also remain stable; if edible oil prices are unstable, vegetable prices will also fluctuate. Lao Li concluded that grain prices will remain stable, but meat prices will rise. According to the above professional forecasts, which of the following, if true, can best question Lao Li’s point of view? A: If the price of cooking oil is stable, then the price of meat will rise B: If the price of cooking oil is stable, then the price of meat will not rise ★ C: If the price of meat does not rise, then the price of cooking oil will rise D: If the price of cooking oil fluctuates, then the price of meat will not rise CoT Pass: B ✓ Gold: B ★		
CoT: “The experts establish a delicate market balance: grain stability ensures vegetable stability, while cooking oil instability disrupts vegetables. Lao Li leaps to a bold conclusion that grains stay stable but meat will rise. To best challenge Lao Li, we need an option that disrupts his certainty. We can immediately eliminate (A) because it agrees with Lao Li’s prediction that meat will rise, offering no skepticism. Options (C) and (D) get tangled up in meat not rising triggering oil rises, or oil fluctuating, which completely misses the core premise of market stability. This leaves us with (B). Why is (B) correct? Because it perfectly connects the hidden variable: it states that if cooking oil is stable, meat will not rise. By introducing oil stability as a stabilizing anchor for meat, it directly invalidates Lao Li’s claim that meat is guaranteed to rise under stable grain conditions. It uses the implicit harmony of the market to dismantle his view. Therefore, (B) is logically the only answer that fits.”		

Table 18: Qualitative analysis of CoT faithfulness: Gemma-2-2B on ARC-Easy.

Gemma-2-2B	ARC-Easy	$\Delta p = 0.18$
Question. A student is investigating in which type of soil a plant will grow best. Which should the student measure to answer the question? A: height of the plants ★ B: weight of the soil C: amount of water used D: temperature of the room CoT Pass: A ✓ Gold: A ★		
CoT: “The question is asking about the type of soil that will grow best. The student should measure the type of soil. The other options are not related to the question. The height of the plants is not related to the type of soil. The weight of the soil is not related to the type of soil. The amount of water used is not related to the type of soil. The temperature of the room is not related to the type of soil. Therefore, the correct answer is (A) height of the plants.”		

Figures 19 and 20 demonstrate that the correlation between ℓ_0 and Δp is generally negligible. This indicates that the variation of ℓ_0 should have minimal impact on the ability to detect shared concepts.

I.3 SAE RECONSTRUCTION

In order to separate model effects from SAE effects, we examine the SAE reconstruction with Qwen3-1.7B on LogiQA as a sanity check. Specifically, we evaluate SAE reconstruction quality using the Fraction of Variance Unexplained (FVU), defined as:

$$\text{FVU} = \frac{|x - \hat{x}|^2}{|x - \bar{x}|^2} \quad (17)$$

Table 19: Qualitative analysis of CoT faithfulness: Gemma-2-2B on OpenbookQA.

Gemma-2-2B OpenbookQA $\Delta p = 0.16$	
Question.	What type of characteristics are people not born with?
A:	genetics
B:	skills ★
C:	physical attributes
D:	height
CoT Pass:	D ✗ Gold: B ★
CoT: “(A) Genetics: This is a characteristic that is passed down from parents to children. (B) Skills: This is a characteristic that is learned through experience. (C) Physical attributes: This is a characteristic that is determined by the body’s physical makeup. (D) Height: This is a characteristic that is determined by genetics. Therefore, the correct answer is (D).”	

Table 20: SAE configuration ablation study results with layer-wise faithfulness results and averaged layer-wise Pearson’s r on Gemma-2-2B.

	Configuration	Δp	CC-SAE	Jaccard	Recall
LogQA	width = 16k, $\ell_0 = 105$ (baseline)	0.206 ± 0.124	0.863 ± 0.113	0.477 ± 0.166	0.780 ± 0.151
	width = 65k, $\ell_0 = 105$	0.198 ± 0.128	0.863 ± 0.107	0.450 ± 0.163	0.770 ± 0.156
	Pearson’s r	0.906 ± 0.031	0.804 ± 0.081	0.700 ± 0.093	0.717 ± 0.112
	width = 16k, $\ell_0 = 46$	0.191 ± 0.129	0.866 ± 0.113	0.511 ± 0.175	0.796 ± 0.156
OpenbookQA	Pearson’s r	0.860 ± 0.117	0.779 ± 0.080	0.633 ± 0.094	0.703 ± 0.052
	width = 16k, $\ell_0 = 105$ (baseline)	0.307 ± 0.234	0.845 ± 0.115	0.464 ± 0.161	0.716 ± 0.165
	width = 65k, $\ell_0 = 105$	0.283 ± 0.238	0.841 ± 0.119	0.405 ± 0.168	0.706 ± 0.178
	Pearson’s r	0.890 ± 0.024	0.812 ± 0.083	0.708 ± 0.076	0.747 ± 0.205
ARC-Easy	width = 16k, $\ell_0 = 46$	0.281 ± 0.242	0.856 ± 0.116	0.468 ± 0.174	0.744 ± 0.170
	Pearson’s r	0.901 ± 0.035	0.706 ± 0.079	0.577 ± 0.068	0.676 ± 0.071
	width = 16k, $\ell_0 = 105$ (baseline)	0.509 ± 0.291	0.854 ± 0.113	0.447 ± 0.159	0.729 ± 0.163
	width = 65k, $\ell_0 = 105$	0.490 ± 0.297	0.846 ± 0.116	0.411 ± 0.164	0.712 ± 0.179
GSM8K	Pearson’s r	0.800 ± 0.120	0.764 ± 0.077	0.698 ± 0.041	0.719 ± 0.170
	width = 16k, $\ell_0 = 46$	0.479 ± 0.301	0.857 ± 0.116	0.477 ± 0.174	0.748 ± 0.171
	Pearson’s r	0.824 ± 0.115	0.745 ± 0.087	0.635 ± 0.040	0.707 ± 0.060
	width = 16k, $\ell_0 = 105$ (baseline)	0.134 ± 0.149	0.753 ± 0.176	0.369 ± 0.159	0.646 ± 0.196
	width = 65k, $\ell_0 = 105$	0.127 ± 0.147	0.744 ± 0.176	0.324 ± 0.149	0.615 ± 0.199
	Pearson’s r	0.886 ± 0.087	0.888 ± 0.084	0.817 ± 0.115	0.824 ± 0.135
	width = 16k, $\ell_0 = 46$	0.129 ± 0.151	0.747 ± 0.186	0.385 ± 0.176	0.639 ± 0.211
	Pearson’s r	0.868 ± 0.114	0.796 ± 0.038	0.695 ± 0.080	0.732 ± 0.081

where $\hat{x} = \text{Dec}(\text{Enc}(x))$ is the SAE reconstruction of residual stream activation x . The metric is bounded meaningfully between 0 and 1. An FVU of 0 corresponds to perfect reconstruction, where the SAE recovers the input activations exactly and leaves no residual variance unexplained. An FVU of 1 corresponds to a reconstruction no better than trivially predicting the dataset mean for every input, i.e., the SAE explains none of the variance beyond a constant baseline.

Across layers, Figure 21 shows that FVU ranges from 0.01 to 0.18, with the strongest reconstruction in late layers (L20–L24, $\text{FVU} \approx 0.01$ –0.03) and higher residual error in early-to-middle layers (L2–L7, $\text{FVU} \approx 0.07$ –0.18). For most layers FVU falls below 0.1, corresponding to over 90% of the activation variance being recovered – a fidelity comparable to or better than that reported for SAEs in prior work (Lieberum et al., 2024; Rajamanoharan et al., 2024; Gao et al., 2025). This provides a favorable basis for the feature-extraction and causal-intervention experiments that follow.

Table 21: SAE configuration ablation study results on Gemma-2-9B.

	Configuration	Δp	CC-SAE	Jaccard	Recall
LogQA	width = 16k, $\ell_0 = 129$ (baseline)	0.366 ± 0.034	0.875 ± 0.069	0.466 ± 0.112	0.787 ± 0.102
	width=131k, $\ell_0 = 129$	0.340 ± 0.050	0.847 ± 0.076	0.449 ± 0.115	0.778 ± 0.105
	Pearson's r	0.880 ± 0.105	0.762 ± 0.119	0.693 ± 0.105	0.761 ± 0.080
	width = 16k, $\ell_0 = 50$	0.350 ± 0.051	0.872 ± 0.072	0.525 ± 0.120	0.819 ± 0.096
	Pearson's r	0.915 ± 0.095	0.756 ± 0.088	0.645 ± 0.121	0.733 ± 0.071
OpenbookQA	width = 16k, $\ell_0 = 129$ (baseline)	0.613 ± 0.059	0.865 ± 0.073	0.456 ± 0.120	0.784 ± 0.113
	width=131k, $\ell_0 = 129$	0.567 ± 0.094	0.829 ± 0.082	0.418 ± 0.116	0.761 ± 0.119
	Pearson's r	0.866 ± 0.121	0.828 ± 0.065	0.703 ± 0.079	0.751 ± 0.092
	width = 16k, $\ell_0 = 50$	0.595 ± 0.079	0.856 ± 0.077	0.498 ± 0.120	0.796 ± 0.105
	Pearson's r	0.914 ± 0.085	0.790 ± 0.092	0.618 ± 0.116	0.699 ± 0.091
Arc Easy	width=16k, $\ell_0 = 129$ (baseline)	0.822 ± 0.209	0.885 ± 0.078	0.482 ± 0.134	0.809 ± 0.119
	width=131k, $\ell_0 = 129$	0.765 ± 0.255	0.855 ± 0.088	0.449 ± 0.131	0.792 ± 0.128
	Pearson's r	0.829 ± 0.120	0.848 ± 0.066	0.667 ± 0.088	0.743 ± 0.094
	width=16k, $\ell_0 = 50$	0.798 ± 0.227	0.878 ± 0.083	0.528 ± 0.140	0.820 ± 0.117
	Pearson's r	0.797 ± 0.187	0.829 ± 0.093	0.598 ± 0.138	0.713 ± 0.109
GSM8K	width = 16k, $\ell_0 = 129$ (baseline)	0.132 ± 0.188	0.657 ± 0.219	0.293 ± 0.178	0.556 ± 0.222
	width = 131k, $\ell_0 = 129$	0.115 ± 0.177	0.561 ± 0.256	0.242 ± 0.181	0.469 ± 0.245
	Pearson's r	0.837 ± 0.075	0.916 ± 0.063	0.873 ± 0.089	0.889 ± 0.068
	width = 16k, $\ell_0 = 50$	0.123 ± 0.182	0.595 ± 0.256	0.307 ± 0.207	0.533 ± 0.242
	Pearson's r	0.864 ± 0.069	0.934 ± 0.048	0.865 ± 0.098	0.890 ± 0.067

Table 22: SAE configuration ablation study results on Qwen3-1.7B.

	Configuration	Δp	CC-SAE	Jaccard	Recall
LogQA	width = 16k, $\ell_0 = 129$ (baseline)	0.511 ± 0.442	0.832 ± 0.112	0.404 ± 0.132	0.673 ± 0.125
	width = 16k, $\ell_0 = 50$	0.512 ± 0.441	0.831 ± 0.110	0.381 ± 0.126	0.656 ± 0.130
	Pearson's r	0.958 ± 0.075	0.607 ± 0.182	0.493 ± 0.063	0.624 ± 0.070
OpenbookQA	width = 16k, $\ell_0 = 129$ (baseline)	0.603 ± 0.452	0.743 ± 0.196	0.358 ± 0.125	0.627 ± 0.148
	width = 16k, $\ell_0 = 50$	0.621 ± 0.445	0.731 ± 0.216	0.326 ± 0.114	0.609 ± 0.155
	Pearson's r	0.914 ± 0.107	0.604 ± 0.160	0.411 ± 0.063	0.576 ± 0.065
ARC-Easy	width = 16k, $\ell_0 = 129$ (baseline)	0.811 ± 0.366	0.768 ± 0.186	0.371 ± 0.132	0.643 ± 0.145
	width = 16k, $\ell_0 = 50$	0.824 ± 0.356	0.757 ± 0.206	0.341 ± 0.122	0.626 ± 0.152
	Pearson's r	0.852 ± 0.151	0.596 ± 0.204	0.415 ± 0.047	0.578 ± 0.088
GSM8K	width = 16k, $\ell_0 = 129$ (baseline)	0.165 ± 0.240	0.670 ± 0.180	0.240 ± 0.152	0.497 ± 0.199
	width = 16k, $\ell_0 = 50$	0.169 ± 0.239	0.657 ± 0.177	0.220 ± 0.140	0.479 ± 0.208
	Pearson's r	0.941 ± 0.058	0.800 ± 0.145	0.673 ± 0.155	0.786 ± 0.123




Figure 15: Ablation study on Gemma-2-2B with increased width and increased sparsity.

J SHARED CONCEPT TRANSITION ACROSS LAYERS

J.1 SAE FEATURE DESCRIPTION GENERATION

We detail the automated interpretability pipeline used to map the shared concepts $S_\cap$ to natural-language descriptions, following Paulo et al. (2025) and using `gemini-3-flash-preview`¹³ as the explainer. Specifically, for each shared SAE feature we mine its top-activating exemplars from the dataset’s question and option texts together with the generated CoT traces. Each exemplar is a 32-token window centered on the peak-activating token, within which we mark the contiguous high-activation span and quantize its per-token activations to an integer scale from 0 to 10. We then prompt the explainer to summarize these exemplars into a single concise description. All case studies use Gemma-2-9B with a compact width-16k SAE applied at every layer.

¹³<https://ai.google.dev/gemini-api/docs/models/gemini-3-flash-preview>






Figure 16: SAE configuration ablation study on Gemma-2-9B with increased width and reduced sparsity.

J.2 CASE STUDIES


Tables 24 – 27 report one qualitative instance each from LogiQA, OpenbookQA, ARC-Easy, and GSM8K. For every instance, we show the question and its options, the predicted answers under the prediction and CoT passes together with the gold answer, the generated CoT, and, for each layer, the shared SAE features $S_\cap$ ablated in the causal intervention (§3.5), reporting their descriptions, pre-ablation activations, and the per-layer Δp . $|S_\cap|$ is typically large, so we cannot list every layer or every shared concept in each table. We instead hand-select a few features from $S_\cap$ whose descriptions are topically relevant to the question. Across all four datasets, the shared SAE features trace a consistent trajectory: early layers encode SAE features tied to the question topic (e.g., cultural entities in LogiQA, spatial orientation in OpenbookQA), while middle and late layers activate




(a) LogiQA



(b) OpenbookQA



(c) ARC-Easy



(d) GSM8K

Figure 17: Ablation study on Qwen3-1.7B with reduced sparsity.

more abstract, answer-evaluation concepts. Additionally, Δp is largest at layers whose SAE features are semantically relevant to the answer. Additionally, we observe that not all SAE features are verbalized in the CoT, while not all verbalizations correspond to highly activated concepts across layers.




Figure 18: Ablation study on Qwen3-8B with reduced sparsity.

Table 23: SAE configuration ablation study results on Qwen3-8B.

	Configuration	Δp	CC-SAE	Jaccard	Recall
LogQA	width = 16k, $\ell_0 = 129$ (baseline)	0.659 ± 0.396	0.872 ± 0.060	0.434 ± 0.092	0.759 ± 0.081
	width = 16k, $\ell_0 = 50$	0.665 ± 0.392	0.877 ± 0.065	0.410 ± 0.091	0.754 ± 0.089
	Pearson's r	0.935 ± 0.091	0.464 ± 0.157	0.381 ± 0.091	0.551 ± 0.101
ARC-Easy/OpenbookQA	width = 16k, $\ell_0 = 129$ (baseline)	0.807 ± 0.334	0.841 ± 0.063	0.379 ± 0.074	0.706 ± 0.081
	width = 16k, $\ell_0 = 50$	0.818 ± 0.327	0.853 ± 0.064	0.356 ± 0.073	0.699 ± 0.082
	Pearson's r	0.929 ± 0.085	0.722 ± 0.181	0.541 ± 0.123	0.720 ± 0.176
ARC-Easy/GSM8K	width=16k, $\ell_0 = 129$ (baseline)	0.939 ± 0.192	0.853 ± 0.063	0.388 ± 0.081	0.732 ± 0.086
	width=16k, $\ell_0 = 50$	0.952 ± 0.166	0.861 ± 0.064	0.365 ± 0.079	0.727 ± 0.089
	Pearson's r	0.806 ± 0.166	0.778 ± 0.218	0.600 ± 0.152	0.768 ± 0.187
GSM8K	width = 16k, $\ell_0 = 129$ (baseline)	0.175 ± 0.284	0.689 ± 0.182	0.266 ± 0.134	0.520 ± 0.183
	width = 16k, $\ell_0 = 50$	0.175 ± 0.284	0.636 ± 0.195	0.209 ± 0.116	0.450 ± 0.186
	Pearson's r	0.946 ± 0.040	0.632 ± 0.121	0.483 ± 0.117	0.594 ± 0.114





Figure 19: ℓ_0 correlation analysis for Gemma-2-2B across four different datasets.





Figure 20: ℓ_0 correlation analysis for Gemma-2-9B across four different datasets.




Figure 21: Layer-wise SAE reconstruction quality (FVU) for Qwen3-1.7B on LogiQA. FVU is the ratio of the SAE’s reconstruction error to the total variance of the original activations, so lower values indicate better reconstruction.

Table 24: Qualitative instance from Gemma-2-9B on LogiQA: question, choices, CoT rationale, and prediction under both conditions, together with the shared concepts $S_\cap$ ablated in the causal intervention (§3.5) and their feature descriptions and pre-ablation activation magnitudes.

Gemma-2-9B LogiQA				
Question. Many elderly Beijings have never forgotten the unique skills they saw at temple fairs when they were young. Today, these unique skills have more formal names—folk art. However, as the modernization of society accelerates, Chinese folk culture is facing unprecedented crisis. The survival crisis of the city, the continuous change of the urban environment, the rapid diversion and transfer of people’s interests and hobbies, and the gradual loss of folk art talents, all of which make the development of folk art face difficulties. Which of the following can be derived from this text? A: Marketization is the way out for folk art B: Folk culture needs rescue protection ★ C: Urban construction should highlight cultural characteristics D: The social status of folk art talents should be improved Prediction Pass: B ✓ CoT Pass: B ✓ Gold: B ★				
CoT: “The text says that the survival of folk art is facing unprecedented crisis. The text also says that the development of folk art is facing difficulties. So, the answer is B. The text does not say that marketization is the way out for folk art. The text does not say that urban construction should highlight cultural characteristics. The text does not say that the social status of folk art talents should be improved. So, the answer is B. . . .”				
Layer	Δp	Feature ID ($S_\cap$)	Act.	Feature Description
L5	0.67	#15137	1.9	Proper nouns and specific cultural or geographic identifiers associated with China, Hong Kong, Taiwan, and Chinese political or social contexts.
		#4649	6.2	Specialized compound nouns and terminology related to social sciences, economics, and formal academic assessments.
L8	0.67	#4149	2.9	Specific entities, proper nouns, and key subjects within logical reasoning problems or translated Chinese cultural contexts.
		#16084	3.0	Formal and technical terminology related to business, industry trends, and economic development.
L11	0.67	#10504	2.9	Pinyin transliterations of Chinese proper nouns, including personal names, locations, and cultural terms.
		#5924	2.4	Proper nouns of Chinese or Taiwanese origin and structural formatting elements used in standardized test questions and logical premises.
L16	0.21	#6229	6.1	Abstract nouns and formal terminology associated with systemic processes, institutional mechanisms, and socio-economic concepts in academic or standardized testing contexts.
L21	0.36	#7862	18.1	Formal academic or examination-style text, often characteristic of Chinese-to-English translations in reading comprehension and logical reasoning contexts.
L32	0.66	#8910	15.8	Assertions or evaluations regarding the truth, validity, or correctness of a specific statement, theory, or position.
L41	0.67	#13986	14.8	Academic and formal terminology related to cultural heritage, archaeological sites, and scientific or logical observations.
marks features whose description semantically overlaps with concepts explicitly verbalized in the CoT above.				

Table 25: Qualitative instance from Gemma-2-9B on OpenbookQA: question, choices, CoT rationale, and prediction under both conditions, together with the shared concepts $S_\cap$ ablated in the causal intervention (§3.5) and their feature descriptions and pre-ablation activation magnitudes.

Gemma-2-9B OpenbookQA				
Question. The middle of the day usually involves the bright star nearest to the earth to be straight overhead why? A: moons gravity B: human planet rotation ★ C: global warming D: moon rotation Prediction Pass: B ✓ CoT Pass: B ✓ Gold: B ★				
CoT: “The sun is the closest star to the earth. The earth rotates on its axis. The earth rotates from west to east. The sun is in the east in the morning and in the west in the evening. The sun is straight overhead at noon. The sun is straight overhead at noon because the earth rotates from west to east...”				
Layer	Δp	Feature ID ($S_\cap$)	Act.	Feature Description
L8	0.76	#12308	2.3	Terms describing spatial orientation, geographical directions, navigation, and physical positioning.
L23	0.75	#11033	12.7	Technical descriptions of physical motion, biological processes, and scientific concepts within educational or multiple-choice contexts.
L24	0.75	#15666	11.3	Scientific descriptions of planetary rotation and biological traits, frequently appearing as specific options or answers within a multiple-choice format.
L25	0.73	#4751	12.2	Technical descriptions of physical rotation, directional movement, and specific nouns used as options in multiple-choice questions.
L26	0.46	#1718	18.1	Technical terminology and specific answer choices within scientific explanations or multiple-choice questions, particularly regarding planetary rotation and physical properties.
L29	0.76	#10712	23.6	Scientific facts and descriptions concerning animal biology, behaviors, and astronomical or physical processes, frequently appearing in educational or multiple-choice contexts.
L31	0.76	#15964	17.2	Factual descriptions and multiple-choice selections concerning biological characteristics, animal types, and physical or astronomical motions.
L33	0.75	#5895	19.9	Scientific and biological terminology describing animal species, ecological behaviors, and physical or astronomical processes.
L38	0.53	#1529	20.1	Key terms and phrases that provide the specific factual or logical completion to a statement or question.
L41	0.76	#1957	17.4	Scientific descriptions and multiple-choice questions regarding celestial bodies, lunar cycles, and astronomical events.
 marks features whose description semantically overlaps with concepts explicitly verbalized in the CoT above.				

Table 26: Qualitative instance from Gemma-2-9B on ARC-Easy: question, choices, CoT rationale, and prediction under both conditions, together with the shared concepts $S_{\cap}$ ablated in the causal intervention (§3.5) and their feature descriptions and pre-ablation activation magnitudes.

Gemma-2-9B ARC-Easy				
Question. A student has just completed a laboratory activity. What is the last action that the student should perform before leaving the lab area? A: wash hands ★ B: turn off all equipment C: put away all glassware D: wash instruments and table tops Prediction Pass.: A ✓ CoT Pass.: A ✓ Gold: A ★				
CoT. “The first thing that the student should do is to wash hands. The second thing that the student should do is to turn off all equipment. The third thing that the student should do is to put away all glassware. The fourth thing that the student should do is to wash instruments and table tops. Therefore, the correct answer is (A) wash hands.”				
Layer	Δp	Feature ID ($S_{\cap}$)	Act.	Feature Description
L1	0.21	#14397	2.2	Nouns identifying the participants, safety protocols, and formal designations of scientific activities within an educational laboratory or classroom setting.
L11	0.29	#3433	3.3	Scientific terminology and formal phrasing used in educational assessments, laboratory observations, and multiple-choice questions.
L15	0.15	#8576	3.9	Structural elements of multiple-choice science questions, including question framing, student-led scenarios, and the identification of correct answers.
L19	0.29	#12737	5.3	Technical terminology and vocabulary used in scientific educational materials, standardized test questions, and laboratory contexts.
L22	0.28	#7909	14.7	Nouns and descriptive phrases identifying physical features, biological structures, or scientific equipment in academic contexts.
L32	0.28	#8551	13.3	Procedural steps, safety equipment, and experimental methodologies found in scientific laboratory instructions and multiple-choice science assessment questions.
L38	0.18	#6835	12.8	Key scientific terms and concepts serving as correct answers or critical components within educational multiple-choice questions and explanations.
L40	0.29	#15047	39.5	Symbols and notation characters used in scientific formulas, temperature units, and structured multiple-choice options.
marks features whose description semantically overlaps with concepts explicitly verbalized in the CoT above.				

Table 27: Qualitative instance from Gemma-2-9B on GSM8K: question, CoT rationale, and prediction under both conditions, together with the shared concepts $S_{\cap}$ ablated in the causal intervention (§3.5) and their feature descriptions and pre-ablation activation magnitudes.

Gemma-2-9B GSM8K				
Question. In a dance class of 20 students, 20% enrolled in contemporary dance, 25% of the remaining enrolled in jazz dance, and the rest enrolled in hip-hop dance. What percentage of the entire students enrolled in hip-hop dance? Prediction Pass.: 60 ✓ CoT Pass.: 60 ✓ Gold: 60 ★				
CoT: “First, we need to find out how many students enrolled in contemporary dance. 20% of 20 students = $0.2 * 20 = 4$ students. Now, we need to find out how many students are left after 4 students enrolled in contemporary dance. $20 - 4 = 16$ students. Next, we need to find out how many students enrolled in jazz dance. 25% of 16 students = $0.25 * 16 = 4$ students. Now, we need to find out how many students are left after 4 students enrolled in jazz dance. $16 - 4 = 12$ students. Finally, we need to find out what percentage of the entire students enrolled in hip-hop dance. 12 students out of 20 students = $12/20 = 0.6 = 60\%$. Therefore, 60% of the entire students enrolled in hip-hop dance.”				
Layer	Δp	Feature ID ($S_{\cap}$)	Act.	Feature Description
L8	0.00	#10482	3.0	Mathematical terminology and symbols used to define proportions, percentages, and quantitative relationships within word problems.
L14	0.03	#6364	9.8	Mathematical terminology and symbols used to express probability, percentages, and statistical likelihood within word problems.
L19	0.01	#760	23.2	Numerical values and percentages representing proportions, ratios, or fractional parts within mathematical contexts.
L25	0.07	#2610	19.5	Mathematical contexts involving the calculation, phrasing, or conversion of percentages and proportions.
L28	0.06	#10761	24.1	Phrases following a fractional or numerical value that specify a portion of a remaining quantity in mathematical word problems.
L31	0.07	#16056	40.5	Transitions and qualifiers used in mathematical word problems to introduce the final, remaining, or resulting quantity in a sequence of calculations.
L33	0.06	#7582	19.9	Clauses in mathematical word problems that introduce specific constraints or define the final question to be solved.
L34	0.07	#3298	55.5	The transition point and leading digit of a numerical value representing a secondary quantity or proportion in a mathematical word problem.
marks features whose description semantically overlaps with concepts explicitly verbalized in the CoT above.				