

Superposition Without Interference? Towards Isolated Interventions via Almost Orthogonal Features in Language Models

Moritz Miller^{124*}, Florent Draye^{13*} & Bernhard Schölkopf¹²³⁴

¹ Max Planck Institute for Intelligent Systems, Tübingen, Germany

² ETH Zurich, Switzerland

³ ELLIS Institute Tübingen, Germany

⁴ Max Planck ETH Center for Learning Systems

moritz.miller@tuebingen.mpg.de

Abstract

A central premise in mechanistic interpretability is that meaningful concepts in language models are represented by linear features in activation space. For such features to support reliable interventions, manipulating one feature should not substantially alter the effects of others. In practice, however, feature entanglement leads to interference such that localized interventions can have unintended downstream effects. Motivated by the *Independent Causal Mechanisms* principle, we propose to constrain internal features to be almost orthogonal. We argue that this promotes modular representations amenable to causal intervention. We formalize this problem by characterizing the gap between an idealized isolated intervention and its realized effect on model outputs in terms of feature interference. We upper-bound the propagation of feature interference in terms of the self-coherence of the feature dictionary, and relate this discrepancy to an explicit orthogonality regularization on the dictionary itself. Empirically, we show that this regularization enables more isolated interventions on mathematical reasoning concepts while preserving model performance. Our code is available under <https://github.com/mrtzmlr/sae-icm>.

1 Introduction

The *Independent Causal Mechanisms* (ICM) principle posits that complex systems can be decomposed into autonomous modules that do not inform or influence each other (Janzing & Schölkopf, 2010). In the context of learned representations, this implies that intervening on one latent mechanism should not systematically alter the functional form of the remaining mechanisms. This perspective is especially appealing for mechanistic interpretability. If a feature corresponds to a meaningful concept, an intervention on that feature should induce a targeted change in model behavior without substantially altering the effects of other features. More broadly, this would move us closer to representations in which a conceptual intervention in the output corresponds to a localized intervention inside the model.

In practice, however, learned features in language models are often not isolated. Features that are strongly aligned in representation space can interfere with one another, so that modifying one feature in the residual stream also changes the influence of others (Elhage et al., 2022). This makes the effect of steering ambiguous, blurs the causal contribution of the targeted feature, and limits intervenability within the model. Prior work on superposition further suggests that such feature alignment can induce interference patterns that are difficult to control and may even support adversarial perturbations (Gorton & Lewis, 2025). In this paper, we formalize feature interference and study how reducing it improves the isolation of feature-level interventions.

*The authors contributed equally to this work. Author order was determined by a 60–40 coin flip.

Question :

Mike writes a 3-page letter to 2 different friends twice a week. How many pages does he write a year?

Standard Generation :

Mike writes a 3-page letter to 2 different friends, so he writes $3 \times 2 = 6$ pages per week.

He writes this letter twice a week, so he writes $6 \times 2 = 12$ pages per week.

There are 52 weeks in a year, so **Mike** writes $12 \times 52 = 624$ pages in a year.

624

The answer is: **624**

Post-Intervention Generation :

If **Aquaman** writes a 3-page letter to 2 different friends, he writes a total of $3 \times 2 = 6$ pages per week.

Since he writes this letter twice a week, he writes a total of $6 \times 2 = 12$ pages per week.

If there are 52 weeks in a year, then he writes a total of $12 \times 52 = 624$ pages in a year.

624

The answer is: **624**

Figure 1: **Intervention on SAE feature** After regularizing the SAE decoder under an orthogonality penalty, we perform an isolated intervention on the feature associated with **Mike** by exchanging it with the feature corresponding to the prefix **aqua**. The model accordingly substitutes **Mike** for **Aquaman** at inference while maintaining its reasoning capabilities.

Sparse autoencoders (SAEs) provide a natural setting in which to address this problem. They have become a central tool in mechanistic interpretability because they uncover large numbers of human-interpretable features in language model representations (Cunningham et al., 2023; Paulo et al., 2024). In our setting, SAEs serve not only to discover such features, but also to shape the geometry of the feature space itself. Concretely, we fine-tune an SAE with an orthogonality penalty, freeze the resulting weights, and then adapt the language model around it using low-rank adaptation, while preserving overall performance (Hu et al., 2021; Wang et al., 2025; Chen et al., 2025). This yields features that remain interpretable while exhibiting substantially less interference, allowing more localized interventions with limited spill-over onto other features. Figure 1 illustrates this setup: after replacing the feature corresponding to the concept **Mike** with one associated with usages of the prefix **aqua**, the model still reasons correctly toward **624** while adapting the intervention coherently to context, replacing the male first name **Mike** with **Aquaman** during generation. Despite this local concept substitution, the downstream effect on other concepts remains limited.

Given this methodology, our contributions are:

- We theoretically relate intervenability in the residual stream to the interference between features (Section 4). By a causal representation argument under a linear hypothesis, we upper-bound the expected discrepancy between true and model output in terms of the representation-space interference. Explicitly controlling the latter, we obtain bounds for the former.
- We fine-tune a LM around a fine-tuned SAE whose decoder matrix is regularized to be almost orthogonal. We demonstrate that this fine-tuning pipeline maintains similar performance and interpretability scores on mathematical reasoning as a non-penalized SAE (Section 5).
- To the best of our knowledge, we are the first to perform *isolated interventions* on specific concepts inside the SAE by explicitly limiting spill-over effects on other features (Section 5.2). We demonstrate that the increased orthogonality penalty improves intervenability in the model’s representation space.

2 Related work

Intervenability In applied causality, a set of interventions is commonly described as datasets from different environments (Peters et al., 2016). Under this notion, one main objective lies in identifying the causal mechanisms prevalent across environments (Guo et al., 2024b). Additionally, the unique disentanglement of hidden concepts appears in causal representation learning (Schölkopf et al., 2021). Here, identifiability concerns finding a unique model which can explain the observed data. While non-identifiability in causal representation learning represents a well-studied problem (Hyvärinen & Pajunen, 1999; Locatello et al., 2019), additional information such as access to interventional data can yield identifiability results (Buchholz et al., 2023). Thus, traditional causality studies interventions as distributional properties of the underlying data-generating process (Spirtes et al., 2000). Recent work utilizes the SAE to infer information about the data (Mencattini et al., 2026). In contrast, we study intervenability on the modeling level. This is similar to uncovering the causal mechanisms inside the model’s representation space (Miller et al., 2025). A general causal abstraction perspective on mechanistic interpretability is provided by Geiger et al. (2023).

Interpretability The linear superposition hypothesis (Elhage et al., 2022) posits that LMs represent features as directions in their activation space. In this way, the model can represent more features than it has dimensions. To disentangle the polysemantic residual stream into human-interpretable features, SAEs reconstruct the activations as sparse linear combinations of features (Cunningham et al., 2023). Further, the linear representation hypothesis (Park et al., 2024) hinges on the observation that semantically unrelated features are represented as almost orthogonal vectors in some representation space (Jiang et al., 2024). Moreover, recent empirical evidence supports the emergence and test-time use of such linear structure in LMs (Chen et al., 2026). In SAEs specifically, we empirically observe high self-coherence between features such that finding a canonical dictionary is impossible (Marks et al., 2024; Leask et al., 2025). In the context of sparsity, weight-sparse models (Gao et al., 2025b) and sparse attention fine-tuning (Draye et al., 2025) share our objective of constructing inherently interpretable architectures.

3 Background

Notation We employ an abbreviated set notation and write $[n] := \{1, \dots, n\}$ for integer n .

3.1 Causality

This section follows Peters et al. (2017). For additional details, we refer the reader there. The structural causal model (SCM) consists of a collection of m structural assignments

$$X_i := f_i(\text{Pa}_i, N_i), i \in [m]$$

for $\text{Pa}_i \subseteq \{X_1, \dots, X_m\} \setminus \{X_i\}$ the parents of X_i , and the distribution $\mathbb{P}_{\mathbf{N}} = \mathbb{P}_{N_1, \dots, N_m}$ over the jointly independent noise variables (Pearl, 2009). It entails a joint distribution $\mathbb{P}_{\mathbf{X}}$ over $\{X_1, \dots, X_m\}$ as well as the marginals $P(X_i | \text{Pa}_i)$ for $i \in [m]$. An intervention in the SCM is the replacement of at least one structural assignment, $l \in [m]$,

$$X_l := \tilde{f}_l(\tilde{\text{Pa}}_l, \tilde{N}_l)$$

inducing $\tilde{P}(X_l | \text{Pa}_l)$. We denote the intervention by $\text{do}(X_l := \tilde{f}_l(\tilde{\text{Pa}}_l, \tilde{N}_l))$. For any $i \in [m]$, $P(X_i | \text{Pa}_i)$ describes the causal mechanism generating X_i . With this, we can state the ICM principle (Janzing & Schölkopf, 2010; Guo et al., 2024a).

Principle 3.1 (Independent Causal Mechanisms). Causal mechanisms are independent of each other in the sense that a change in one mechanism $P(X_l | \text{Pa}_l)$ does not inform or influence any other mechanism $P(X_i | \text{Pa}_i)$, for $i \neq l$.

3.2 Sparse autoencoders

SAEs are trained on reconstructing the activations in the residual stream under a sparsity penalty. The output is then a sparse linear combination of vectors. Let $\mathbf{x} \in \mathbb{R}^m$ denote the Transformer state prior to the SAE. A TopK SAE (Makhzani & Frey, 2013; Gao et al., 2025a) then outputs the reconstruction $\hat{\mathbf{x}} \in \mathbb{R}^m$,

$$\begin{aligned}\mathbf{z} &= \text{TopK}(\mathbf{E}\mathbf{x} + \mathbf{b}_\mathbf{E}) \\ \hat{\mathbf{x}} &= \mathbf{D}\mathbf{z} + \mathbf{b}_\mathbf{D}\end{aligned}$$

for $\mathbf{z} \in \mathbb{R}^d$ the coefficients for the features and $\mathbf{E} \in \mathbb{R}^{d \times m}$, $\mathbf{D} \in \mathbb{R}^{m \times d}$ encoder and decoder matrices with biases $\mathbf{b}_\mathbf{E} \in \mathbb{R}^d$, $\mathbf{b}_\mathbf{D} \in \mathbb{R}^m$, respectively. We call $\mathbf{D}$ our feature dictionary with $d \gg m$ features.

We train our autoencoder on the normalized reconstruction loss with penalty on $\mathbf{D}^\top \mathbf{D}$,

$$\ell(\mathbf{x}) = \frac{\|\hat{\mathbf{x}} - \mathbf{x}\|_2}{\|\mathbf{x}\|_2} + \lambda \|\text{tril}(\mathbf{D}^\top \mathbf{D})\|_\mathbb{F}^2 \quad (1)$$

where $\|\text{tril}(\mathbf{D}^\top \mathbf{D})\|_\mathbb{F}^2$ enforces almost orthogonality between features with $\lambda > 0$. By $\text{tril}(\mathbf{D}^\top \mathbf{D})$ we denote the off-diagonal lower-triangular elements of $\mathbf{D}^\top \mathbf{D}$. During the forward pass, TopK only activates the features associated with the K highest coefficients with all other activations zeroed out.

3.3 Sparse dictionary learning

Similar to the features in the SAE decoder, the atoms compose the dictionary in sparse dictionary learning. If the number of atoms in the dictionary exceeds the number of dimensions per atom, we call the dictionary *overcomplete*. Thus, for $\mathbf{D} \in \mathbb{R}^{m \times d}$ with $d > m$, the atoms form an overcomplete spanning set. In the overcomplete setting, self-coherence determines how sparse the representations must be for identifiability to hold (Donoho et al., 2005). Relevant to our later work, we define the mean self-coherence here as

$$\mu(\mathbf{D}) := \frac{1}{d(d-1)} \sum_{j=1}^d \sum_{k \neq j} \langle \mathbf{f}_j, \mathbf{f}_k \rangle^2. \quad (2)$$

Welch bounds provide lower bounds on the similarity between any two atoms (Welch, 1974). Let $\{\mathbf{f}_j\}_{j \in [d]}$ unit-norm for $\mathbf{f}_j \in \mathbb{R}^m$ denote the d atoms composing the dictionary. For $\mathbb{N}_{>0}$ non-zero natural numbers and $s \in \mathbb{N}_{>0}$, we have

$$\frac{1}{d^2} \sum_{j=1}^d \sum_{k=1}^d |\langle \mathbf{f}_j, \mathbf{f}_k \rangle|^{2s} \geq \frac{1}{\binom{m+s-1}{s}}. \quad (3)$$

In Appendix A, we discuss self-coherence for identifiability and add further Welch bounds.

4 Aligned features reduce intervenability

We consider the causal graph

$$\begin{array}{ccc} Z & \longrightarrow & Y \\ \downarrow & & \\ X & \longrightarrow & \hat{Y} \end{array}$$

for $Z \in \mathbb{R}^d$ the latent state, $X \in \mathbb{R}^m$ the model representation, Y the true output, and $\hat{Y}$ the model output. We assume that the model representation is a sparse linear combination of features, i.e.,

$$X = \mathbf{D}Z.$$

In our theoretical setup, we assume that Z is exogenous with distribution $\mathbb{P}_Z$, and, following Park et al. (2024), we model the true and conditional output distributions as exponential-family distributions with natural parameters

$$\boldsymbol{\eta} := \mathbf{A}Z, \quad \hat{\boldsymbol{\eta}} := \mathbf{A}\hat{Z},$$

where $\mathbf{D} = [\mathbf{f}_1, \dots, \mathbf{f}_d] \in \mathbb{R}^{m \times d}$, $\mathbf{A} \in \mathbb{R}^{q \times d}$, and $\gamma(\mathbf{y}) \in \mathbb{R}^q$ is a fixed embedding of the output $\mathbf{y}$. The decoded latent representation is

$$\hat{Z} = \mathbf{D}^\top X = \mathbf{D}^\top \mathbf{D}Z = \mathbf{G}Z.$$

Accordingly, the true and model output distributions are given by

$$P(Y = \mathbf{y} \mid Z) \propto \exp(\boldsymbol{\eta}^\top \gamma(\mathbf{y})), \quad P(\hat{Y} = \mathbf{y} \mid Z) \propto \exp(\hat{\boldsymbol{\eta}}^\top \gamma(\mathbf{y})).$$

Thus, the true output depends on the latent state through $\boldsymbol{\eta} = \mathbf{A}Z$, whereas the model output depends on the decoded representation through $\hat{\boldsymbol{\eta}} = \mathbf{A}\hat{Z}$.

Now fix $j \in [d]$ and $\delta \in \mathbb{R}$, and consider the intervention

$$\text{do}(Z_j := Z_j + \delta),$$

which yields the intervened latent state

$$Z^{(j,\delta)} := Z + \delta \mathbf{e}_j.$$

Let $Y^{(j,\delta)}$ and $\hat{Y}^{(j,\delta)}$ denote the corresponding post-intervention output probability vectors. We measure the discrepancy between the true and model post-intervention outputs by

$$\mathbb{E}[\|Y^{(j,\delta)} - \hat{Y}^{(j,\delta)}\|_2^2]. \quad (4)$$

and upper-bound the expected discrepancy at the output level. We defer all proofs and details to Appendix B.

Theorem 4.1 (Upper bound on interference). *Let $\mathbf{D} = [\mathbf{f}_1, \dots, \mathbf{f}_d] \in \mathbb{R}^{m \times d}$ and assume that*

$$\|\mathbf{f}_j\|_2 = 1, \quad j \in [d].$$

Then, for $\mu(\mathbf{D})$ as defined in (2),

$$\frac{1}{d} \sum_{j=1}^d \mathbb{E}[\|Y^{(j,\delta)} - \hat{Y}^{(j,\delta)}\|_2^2] \leq \frac{1}{2} \|\mathbf{A}\|_{\text{op}}^2 \mathbb{E}[\|(\mathbf{G} - \mathbf{I})Z\|_2^2] + \frac{1}{2} \|\mathbf{A}\|_{\text{op}}^2 \delta^2 (d-1) \mu(\mathbf{D}).$$

The theorem shows that the post-intervention mismatch is controlled by a baseline term,

$$\mathbb{E}[\|(\mathbf{G} - \mathbf{I})Z\|_2^2],$$

and an intervention-specific term,

$$\delta^2 \sum_{k \neq j} \langle \mathbf{f}_k, \mathbf{f}_j \rangle^2.$$

Thus, when the columns of $\mathbf{D}$ are almost orthogonal, interventions on a latent feature induce only limited distortion in the model output. This becomes explicit under additional assumptions on the latent distribution. Suppose that Z has uniformly random support of size K , and that conditioned on its support the nonzero coordinates are independent, mean zero, and have variance σ^2 . Then, the averaged bound of the theorem yields

$$\frac{1}{d} \sum_{j=1}^d \mathbb{E}[\|Y^{(j,\delta)} - \hat{Y}^{(j,\delta)}\|_2^2] \leq \frac{1}{2} \|\mathbf{A}\|_{\text{op}}^2 (d-1) \mu(\mathbf{D}) (K\sigma^2 + \delta^2).$$

Empirical investigation In this regime, both the baseline term and the intervention-specific term are governed by the same coherence quantity $\mu(\mathbf{D})$, making the role of decoder geometry fully explicit. We design a short numerical example to demonstrate the effect of interference between features on the expected post-intervention mismatch. For $\mathbf{A}$ random matrix with standard Gaussian entries, we consider the case where $q > d > m$. If we understand q as the output dimension similar to the vocabulary size in language modeling, this follows the same hierarchy as in our SAE experiments. We compare an implementation with L2-normalized random Gaussian columns to settings with higher and lower self-coherence. Confirming our intuition, we observe that low self-coherence decreases the intervention mismatch. Details and results we defer to Appendix B.




Figure 2: **Orthogonality evaluation loss** We plot the orthogonality loss $\|\text{tril}(\mathbf{D}^\top \mathbf{D})\|_F^2$ for all values of λ . For each λ , we evaluate on a subset of 1’024 active features in the decoder. Error bars represent confidence intervals obtained from running 100 evaluations.

5 Experiments

Recent work demonstrates that fixing an SAE inside the LM during fine-tuning achieves low cross-entropy (Chen et al., 2025). We extend this two-step procedure (Wang et al., 2025) by restricting the SAE decoder to be almost orthogonal and fine-tuning the LM around the SAE. We insert the SAE after Transformer layer 13 of 26. In doing so, we add to the previous section by providing evidence that orthogonality regularization at an intermediate layer also improves intervenability.

1. We optimize a pre-trained TopK SAE (Gao et al., 2025a) for low self-reconstruction under an orthogonality penalty (1) on the decoder matrix.
2. Fixing the SAE decoder, we then insert the SAE into a Gemma 2 2B Transformer and low-rank adapt (Hu et al., 2021) the LM on cross-entropy.

In total, we fine-tune four SAEs with orthogonality penalty $\lambda \in \{0, 10^{-6}, 10^{-5}, 10^{-4}\}$. In particular, we demonstrate all our claims in comparison to the non-penalized case of $\lambda = 0$. Throughout, we choose $K = 20$. Both steps in our pipeline involve fine-tuning the model for one epoch on MetaMathQA (Yu et al., 2023). First, we fully fine-tune the SAE. Next, we low-rank adapt the attention matrices including the non-decoder SAE weights. Therefore, we restrict the decoder to contain almost orthogonal features, while allowing the rest of the architecture to adapt around this penalty. Adopting the SAEBench (Karvonen et al., 2025) module of size $2^{16} = 65’536$, we input the SAE into the 26-layer residual stream after layer 13. Thus, all information flows through the SAE. To limit the computational overhead induced by the penalty in (1), we evaluate the decoder regularization at every step on a subset of $2^{10} = 1’024$ randomly drawn features. This makes fine-tuning feasible while updating decoder weights progressively. All details including hyperparameter choice and training duration are deferred to Appendix C.

5.1 Almost orthogonality while keeping performance

Figure 2 displays the orthogonality loss $\|\text{tril}(\mathbf{D}^\top \mathbf{D})\|_F^2$ for increasing orthogonality penalty. Error bars represent the basic bootstrap (Efron, 1979) confidence intervals at the 95%-level for 100 sampled datasets. For computational efficiency, we compute the evaluation loss on 100 draws of $2^{10} = 1’024$ randomly selected active features. That is, we only evaluate on features

that activate at least once in the test set. First, we observe that no penalty yields the highest orthogonality loss. This confirms the intuition that features do not empirically tend toward orthogonality if not explicitly regularized to do so. Further, we observe monotonically decreasing loss as λ increases. Inserting the orthogonality penalty, therefore, has the desired effect.

Next, we evaluate the fine-tuned LM on the GSM8K (Cobbe et al., 2021) test set. *Inter alia*, GSM8K’s training set is augmented in MetaMathQA. Yu et al. (2023) observe performance in the range of $[0.665, 0.777]$ for 7B models fine-tuned on MetaMathQA. In our two-step fine-tuned 2B LM, we obtain competitive performance for increasing orthogonality penalty λ . Figure 3 represents the accuracy on the GSM8K test set. While just below the performance of the 7B models, our 2B models do not significantly differ in performance for increasingly strict penalty. Again, the evaluation accuracy on the model fine-tuned for all non-zero values of λ is comparable to the no-penalty LM. If anything, we observe a slight performance increase in the regularization parameter λ . Together with our findings from Figure 2, we show that stricter orthogonality does not come at the cost of performance.





Figure 3: **Evaluation on GSM8K** We evaluate on the GSM8K test set. Error bars represent the basic bootstrap confidence intervals (Efron, 1979) on 100 randomly drawn datasets.

5.2 Localized interventions work

In Section 4, we show that high self-coherence in the dictionary inhibits isolated intervenability on individual features. To confirm this idea, we intervene locally on known features in the SAE and record the downstream effect during generation. In particular, every time a certain feature is activated in the residual stream, we swap it for a different one. We then check if the model remains capable of reasoning about the overarching problem while adapting its generation output to the exchanged features.



(a) **Evaluation on mathematical reasoning** We plot mathematical reasoning performance on 3’960 hand-designed examples after intervening on the SAE. Error bars are basic bootstrap confidence intervals (Efron, 1979) with 100 samples.

(b) **Evaluation on inclusion of correct name** We plot the fraction of correctly included first names after intervening on the corresponding feature in the SAE. Error bars are basic bootstrap confidence intervals (Efron, 1979) of 100 random draws.

Figure 4: **Intervenability** We intervene on 12 features in the SAE corresponding to the concept of first names. Then, we evaluate the ability to maintain reasoning performance and the correct generation of the included first name.

Under the ICM principle, isolated intervention implies that altering one feature does not inform or influence the behavior of other features in the SAE. Translated to natural language modeling, we ideally expect the model to only swap the corresponding feature. To check this, we design the following experiment. Across the different SAEs, we find 12 features which we interpret as encoding conventionally male first names. We then choose 30 examples from GSM8K that involve male actors. In each example, we replace the original name with one of the 12 names and run the standard evaluation task.




Figure 5: **Interventions in the SAE** We perform interventions in the SAE by **turning off** the intervened on feature and **adding** to the residual stream the feature we replace with.

ture to be present in the output while keeping candidate values (Ameisen et al., 2025).

1. found an **aquarium** for \$10.00
2. in the two **aquariums** is twice
3. in the **fish tank**. One-third

Listing 1: **The aqua feature** The top-three input examples activating the **aqua** feature at $\lambda = 10^{-4}$ are related to aquaria.

Having access to the SAE features, however, we are able to swap the feature indices for another one of the 11 remaining first names. We then check if the model (a) *keeps performance* comparable to Figure 3, (b) *correctly drops* the name related to the feature we set to zero, and (c) *includes* the first name we inject into the residual stream through our intervention inside the SAE. Testing all combinations of names on each of the examples, we generate a dataset of $12 * 11 * 30 = 3'960$ examples.

A well-known issue in mechanistic interpretability, it is non-trivial which value is required to faithfully intervene in the representation space. As the coefficients associated with each feature are continuous, turning the feature on and off is a non-binary problem. In other words, pushing on a feature generation intact requires testing multiple

We thus sweep over hyperparameters between 10 and 500. For any combination of insertion value and orthogonality penalty, the accuracy only moves within an interval of 6 percentage points. Hence, we decide based on the models' ability to remove the intervened on concept and choose 200. Confirmation that these findings hold more broadly and the 30 selected examples are deferred to Appendix C.

Figure 4 depicts our results. As mentioned, we observe in Figure 4a that performance is comparable across different choices for λ . No setup performs significantly differently from the zero-penalty configuration. Relevant to our hypothesis, for stricter orthogonality penalty, the model significantly more often includes the correct first name in the generation. Figure 4b demonstrates that the orthogonality penalty monotonically increases in λ . Further, we recover the first name around 74.9% of the time for $\lambda = 10^{-4}$ while only recovering it around 62.2% of the time without penalty. Beyond first names, we extend this analysis to exchanging location and animal concepts in the LM output. We also run standard interpretability experiments and obtain constant performance across orthogonality penalties. Further, we run our intervenability experiments on Llama 3.2 1B to strengthen our results. For all this, we refer the reader to Appendix C.

To conclude this section, we discuss the insertion of a feature that corresponds to concepts which are not conventionally male first names. Listing 1 represents three input examples that activate the same feature. The corresponding explanation in the model with $\lambda = 10^{-4}$ reads
Aquarium capacity or fish population size is the latent concept shared among the spans.

Swapping this feature for the feature associated with the first name **Mike**, we observe that the model adapts this feature to be concordant with the preceding context. In fact, the

LM introduces **Aquaman**. We print the model outputs for this intervention in Figure 1. Moreover, Figure 5 schematically represents an intervention in the SAE. We note that the feature associated with the concept **Aquaman** has no incoming edge while the downstream effect of the feature corresponding to **Mike** is turned off.

5.3 Behavioral and internal measures of interference

To measure interference beyond accuracy, we report two additional metrics on our intervention experiments. First, we report the ROUGE-L RECALL (Lin, 2004) between the standard and the intervened generations as an external metric at the token level. This captures how closely the two generations are aligned as a function of the longest common subsequence between generations. In particular, for S and I the standard generation and intervened generation, respectively, we define the longest common subsequence LCS between S and I as the longest common sequence of tokens between the two. The sequence need not be consecutive. Under an ideal isolated intervention, therefore, only the first name corresponding to the intervened on feature changes while keeping general reasoning intact. For U_S the set of all tokens in our standard generation, ROUGE-L RECALL is then given by

$$\text{ROUGE-L RECALL}(S, I) = \frac{\text{LCS}(S, I)}{|U_S|}.$$

Second, we compute the symmetric Jensen-Shannon divergence (D_{JSD}) between the empirical distributions of the features used in both generations, P_S and P_I , as an internal, SAE-level statistic. For $D_{\text{KL}}(P||Q)$ the Kullback-Leibler divergence between P and Q and $M = \frac{1}{2}(P_S + P_I)$ a mixture between P_S and P_I , the Jensen-Shannon divergence is defined as

$$D_{\text{JSD}}(P_S||P_I) = \frac{1}{2}D_{\text{KL}}(P_S||M) + \frac{1}{2}D_{\text{KL}}(P_I||M).$$

Table 1 displays the two measures for all configurations. Significantly higher ROUGE-L RECALL and lower D_{JSD} at the strictest orthogonality level $\lambda = 10^{-4}$ relative to the no-penalty setting indicate lower interference. *Cum grano salis*, the results for $\lambda \in \{10^{-6}, 10^{-5}\}$ do not yield definitive conclusions. Studying this further, we present detailed statistics on the additional datasets in Appendix C. Across datasets, we observe significantly lower D_{JSD} for the strictest penalty compared to $\lambda = 0$ as well as significantly higher ROUGE-L RECALL on the location dataset.

λ	ROUGE-L RECALL		D_{JSD}	
	Mean	Confidence interval	Mean	Confidence interval
0	0.763	[0.756, 0.771]	0.222	[0.219, 0.225]
10^{-6}	0.746	[0.739, 0.750]	0.254	[0.252, 0.257]
10^{-5}	0.730	[0.723, 0.735]	0.235	[0.232, 0.239]
10^{-4}	0.773	[0.766, 0.779]	0.214	[0.211, 0.217]

Table 1: **Results for interference measures** We evaluate ROUGE-L RECALL and D_{JSD} on the tokens and features after intervening on the 3’960 examples defined above. We report mean and the basic bootstrap confidence interval over all observations.

6 Discussion

Sparse autoencoders (SAEs), which have become a standard tool in mechanistic interpretability, can uncover millions of human-interpretable features in a model’s representation space (Cunningham et al., 2023; Paulo et al., 2024). At the same time, recent work has raised concerns about the reliability of these features. In particular, SAEs trained on the same data with different random seeds often produce substantially different feature sets (Paulo

& Belrose, 2025), and features learned by an SAE can themselves be reconstructed from a smaller set of features using a meta-SAE (Leask et al., 2025). Also with respect to feature absorption (Chanin et al., 2025; Li & Ren, 2025), these results suggest that SAE features may not be truly atomic and may fail to be uniquely identifiable.

This concern is closely connected to classical results in dictionary learning, where identifiability depends critically on the geometry of the dictionary (Olshausen & Field, 1997; Lee et al., 2006). In particular, recovery guarantees are strongest when the dictionary has low *self-coherence*, meaning that its atoms are nearly orthogonal (Donoho et al., 2005). By contrast, when the dictionary is highly self-coherent, multiple distinct sparse combinations of features can generate the same observation (Schnass & Vandergheynst, 2008; Garfinkle & Hillar, 2016), so that even recovering the sparse coefficients may fail, let alone identifying the underlying features uniquely. Our work suggests that these representational desiderata need not come at the expense of performance: we show that it is possible to preserve fine-tuning performance on a specific task while enforcing activations to be represented as sparse combinations of nearly orthogonal features. Compared to previous work (Korznikov et al., 2025), we enforce global orthogonality across the full dictionary. This points to a possible route toward improving the identifiability and atomicity of SAE features without sacrificing downstream utility.

Finally, it has recently been hypothesized that adversarial examples are not bugs but a consequence of feature superposition (Gorton & Lewis, 2025). We view our method as demonstrating a fundamental constraint in this regard: feature interference limits intervention capabilities. When features are represented by non-orthogonal directions, interventions on one feature inevitably spill over to others. Our controlled model-in-the-loop setup makes this limitation explicit and shows that reducing interference improves the reliability of targeted interventions. A natural follow-up question is then whether enforcing greater feature orthogonality reduces superposition and, in turn, mitigates adversarial vulnerabilities. If so, orthogonality could act as a strong architectural bias toward more robust and safer models. Establishing this connection would motivate the development of more scalable versions of our approach.

Limitations Our theoretical study is focused on interventions in the residual stream. Choosing two off-the-shelf SAEs (Karvonen et al., 2025; Cho et al., 2025), we insert the module in Llama 3.2 1B and Gemma 2 2B only after layers 12 and 13, respectively. In consequence, we are constrained to uncovering features that tend to arise in the middle to later layers of the Transformer. Empirically, these concepts often represent high-level concepts relative to the low-level concepts encoded in earlier layers (Dorszewski et al., 2025). Additionally, we do not influence the arithmetic problem by changing just the subject of the sentence. In causal terms, we modify a feature that is not an ancestor of the final answer. Thus, the causal mechanism generating the mathematical solution remains unaffected. Extending this with tasks different from mathematical reasoning represents a potential avenue for future research.

7 Conclusion

We argue that feature entanglement inhibits the ability to perform isolated interventions. Formally relating intervenability to the ICM principle, we provide an upper bound for the expected discrepancy between true and model output in terms of the representation-space interference. We further show that it is possible to fine-tune a LM around a fixed middle layer SAE under orthogonality regularization. While keeping performance on the target dataset unchanged, we reduce the interference between features. In doing so, we perform precise local interventions by exchanging a single concept without affecting others. Moreover, we are able to swap two concepts with greatly improved performance. These results shed light on the practical consequences of feature interference in LMs and suggest that encouraging orthogonality may be a principled way to mitigate it.

Ethics statement

This work studies intervenability in language models, with the goal of enabling more precise and localized modifications of internal representations. Improving intervenability can support better control over model behavior and help identify how specific internal components contribute to outputs. These properties are relevant for developing more reliable and predictable AI systems.

At the same time, increasing the ability to intervene on internal features introduces potential dual-use concerns. More precise control mechanisms could be misused to steer model outputs in unintended or deceptive ways. However, our approach operates at the level of internal representations and requires direct access to model internals and technical expertise. This limits accessibility and reduces the likelihood of misuse in practice.

Our experiments are conducted on standard benchmarks and do not involve human subjects or sensitive data. The proposed method does not aim to enhance model capabilities, but rather to improve the structure and controllability of learned representations. We believe that advancing intervenability is a useful step toward safer and more transparent AI systems, while recognizing the importance of continued work on safeguards and responsible use.

Acknowledgements

F.D. acknowledges support through a fellowship from the *Hector Fellow Academy*.

References

- Emmanuel Ameisen, Jack Lindsey, Adam Pearce, Wes Gurnee, Nicholas L. Turner, Brian Chen, Craig Citro, David Abrahams, Shan Carter, Basil Hosmer, Jonathan Marcus, Michael Sklar, Adly Templeton, Trenton Bricken, Callum McDougall, Hoagy Cunningham, Thomas Henighan, Adam Jermyn, Andy Jones, Andrew Persic, Zhenyi Qi, T. Ben Thompson, Sam Zimmerman, Kelley Rivoire, Thomas Conerly, Chris Olah, and Joshua Batson. Circuit tracing: Revealing computational graphs in language models. *Transformer Circuits Thread*, 2025. URL <https://transformer-circuits.pub/2025/attribution-graphs/methods.html>.
- Simon Buchholz, Goutham Rajendran, Elan Rosenfeld, Bryon Aragam, Bernhard Schölkopf, and Pradeep Ravikumar. Learning linear causal representations from interventions under general nonlinear mixing. In *Proceedings of the 37th International Conference on Neural Information Processing Systems*, NIPS ’23, Red Hook, NY, USA, 2023. Curran Associates Inc.
- David Chanin, James Wilken-Smith, Tomáš Dulka, Hardik Bhatnagar, Satvik Golechha, and Joseph Bloom. A is for absorption: Studying feature splitting and absorption in sparse autoencoders, 2025. URL <https://arxiv.org/abs/2409.14507>.
- Matthew Chen, Joshua Engels, and Max Tegmark. Low-rank adapting models for sparse autoencoders. 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. 8077–8092. PMLR, 13–19 Jul 2025. URL <https://proceedings.mlr.press/v267/chen25r.html>.
- Michelle Chao Chen, Moritz Miller, Bernhard Schölkopf, and Siyuan Guo. On the emergence and test-time use of structural information in large language models. *arXiv preprint arXiv:2601.17869*, 2026.
- Seonglae Cho, Harryn Oh, Donghyun Lee, Luis Eduardo Rodrigues Vieira, Andrew Bermingham, and Ziad El Sayed. FaithfulSAE: Towards capturing faithful features with sparse autoencoders without external datasets dependency. In *ACL 2025 Student Research Workshop*, 2025. URL <https://openreview.net/forum?id=tBn9ChHGG9>.

- 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. *arXiv preprint arXiv:2110.14168*, 2021.
- Hoagy Cunningham, Aidan Ewart, Logan Riggs Smith, Robert Huben, and Lee Sharkey. Sparse autoencoders find highly interpretable features in language models. *ArXiv*, abs/2309.08600, 2023. URL <https://api.semanticscholar.org/CorpusID:261934663>.
- David L Donoho, Michael Elad, and Vladimir N Temlyakov. Stable recovery of sparse overcomplete representations in the presence of noise. *IEEE Transactions on information theory*, 52(1):6–18, 2005.
- Teresa Dorszewski, Lenka Tětková, Robert Jenssen, Lars Kai Hansen, and Kristoffer Knutsen Wickstrøm. From colors to classes: Emergence of concepts in vision transformers, 2025. URL <https://arxiv.org/abs/2503.24071>.
- Florent Draye, Anson Lei, Ingmar Posner, and Bernhard Schölkopf. Sparse attention post-training for mechanistic interpretability. *arXiv preprint arXiv:2512.05865*, 2025.
- B. Efron. Bootstrap methods: Another look at the jackknife. *The Annals of Statistics*, 7(1): 1–26, 1979. ISSN 00905364, 21688966. URL <http://www.jstor.org/stable/2958830>.
- Nelson Elhage, Tristan Hume, Catherine Olsson, Nicholas Schiefer, Tom Henighan, Shauna Kravec, Zac Hatfield-Dodds, Robert Lasenby, Dawn Drain, Carol Chen, Roger Grosse, Sam McCandlish, Jared Kaplan, Dario Amodei, Martin Wattenberg, and Christopher Olah. Toy models of superposition. *Transformer Circuits Thread*, 2022. URL https://transformer-circuits.pub/2022/toy_model/index.html.
- 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*, 2025a. URL <https://openreview.net/forum?id=tcsZt9ZNKD>.
- Leo Gao, Achyuta Rajaram, Jacob Coxon, Soham V Govande, Bowen Baker, and Dan Mossing. Weight-sparse transformers have interpretable circuits. *arXiv preprint arXiv:2511.13653*, 2025b.
- Charles J Garfinkle and Christopher J Hillar. Robust identifiability in sparse dictionary learning. *arXiv preprint arXiv:1606.06997*, 2016.
- Atticus Geiger, Duligur Ibeling, Amir Zur, Maheep Chaudhary, Sonakshi Chauhan, Jing Huang, Aryaman Arora, Zhengxuan Wu, Noah D. Goodman, Christopher Potts, and Thomas F. Icard. Causal abstraction: A theoretical foundation for mechanistic interpretability. 2023. URL <https://api.semanticscholar.org/CorpusID:255749463>.
- S. Gershgorin. Über die abgrenzung der eigenwerte einer matrix. *Izvestija Akademii Nauk SSSR, Serija Matematika*, 7(3):749–754, 1931.
- Liv Gorton and Owen Lewis. Adversarial examples are not bugs, they are superposition. *arXiv preprint arXiv:2508.17456*, 2025. URL <https://arxiv.org/abs/2508.17456>.
- Siyuan Guo, Viktor Tóth, Bernhard Schölkopf, and Ferenc Huszár. Causal de finetti: On the identification of invariant causal structure in exchangeable data, 2024a. URL <https://arxiv.org/abs/2203.15756>.
- Siyuan Guo, Chi Zhang, Karthika Mohan, Ferenc Huszár, and Bernhard Schölkopf. Do finetti: On causal effects for exchangeable data, 2024b. URL <https://arxiv.org/abs/2405.18836>.
- J. Edward Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yanzhi Li, Shean Wang, and Weizhu Chen. Lora: Low-rank adaptation of large language models. *ArXiv*, abs/2106.09685, 2021. URL <https://api.semanticscholar.org/CorpusID:235458009>.

- Aapo Hyvärinen and Petteri Pajunen. Nonlinear independent component analysis: Existence and uniqueness results. *Neural Networks*, 12(3):429–439, 1999. ISSN 0893-6080. doi: [https://doi.org/10.1016/S0893-6080\(98\)00140-3](https://doi.org/10.1016/S0893-6080(98)00140-3). URL <https://www.sciencedirect.com/science/article/pii/S0893608098001403>.
- D Janzing and B Schölkopf. Causal inference using the algorithmic Markov condition. *IEEE Transactions on Information Theory*, 56(10), 2010.
- Yibo Jiang, Goutham Rajendran, Pradeep Ravikumar, Bryon Aragam, and Victor Veitch. On the origins of linear representations in large language models. In *Proceedings of the 41st International Conference on Machine Learning, ICML’24*. JMLR.org, 2024.
- Adam Karvonen, Can Rager, Johnny Lin, Curt Tigges, Joseph Bloom, David Chanin, Yeu-Tong Lau, Eoin Farrell, Callum McDougall, Kola Ayonrinde, Demian Till, Matthew Wearden, Arthur Conmy, Samuel Marks, and Neel Nanda. Saebench: A comprehensive benchmark for sparse autoencoders in language model interpretability, 2025. URL <https://arxiv.org/abs/2503.09532>.
- Anton Korznikov, Andrey Galichin, Alexey Dontsov, Oleg Rogov, Elena Tutubalina, and Ivan Oseledets. Ortsae: Orthogonal sparse autoencoders uncover atomic features, 2025. URL <https://arxiv.org/abs/2509.22033>.
- Patrick Leask, Bart Bussmann, Michael Pearce, Joseph Bloom, Curt Tigges, Noura Al Moubayed, Lee Sharkey, and Neel Nanda. Sparse autoencoders do not find canonical units of analysis. *arXiv preprint arXiv:2502.04878*, 2025.
- Honglak Lee, Alexis Battle, Rajat Raina, and Andrew Ng. Efficient sparse coding algorithms. In B. Schölkopf, J. Platt, and T. Hoffman (eds.), *Advances in Neural Information Processing Systems*, volume 19. MIT Press, 2006. URL https://proceedings.neurips.cc/paper_files/paper/2006/file/2d71b2ae158c7c5912cc0bbde2bb9d95-Paper.pdf.
- Ed Li and Junyu Ren. Unlocking hierarchical concept discovery in language models through geometric regularization. In *ICLR 2025 Workshop on Building Trust in Language Models and Applications*, 2025. URL <https://openreview.net/forum?id=i31cKXiyim>.
- Chin-Yew Lin. ROUGE: A package for automatic evaluation of summaries. In *Text Summarization Branches Out*, pp. 74–81, Barcelona, Spain, July 2004. Association for Computational Linguistics. URL <https://aclanthology.org/W04-1013/>.
- Francesco Locatello, Stefan Bauer, Mario Lucic, Gunnar Raetsch, Sylvain Gelly, Bernhard Schölkopf, and Olivier Bachem. Challenging common assumptions in the unsupervised learning of disentangled representations. In Kamalika Chaudhuri and Ruslan Salakhutdinov (eds.), *Proceedings of the 36th International Conference on Machine Learning*, volume 97 of *Proceedings of Machine Learning Research*, pp. 4114–4124. PMLR, 09–15 Jun 2019. URL <https://proceedings.mlr.press/v97/locatello19a.html>.
- Alireza Makhzani and Brendan Frey. K-sparse autoencoders. *arXiv preprint arXiv:1312.5663*, 2013.
- Luke Marks, Alasdair Paren, David Krueger, and Fazl Barez. Enhancing neural network interpretability with feature-aligned sparse autoencoders. *arXiv preprint arXiv:2411.01220*, 2024.
- Tommaso Mencattini, Riccardo Cadei, and Francesco Locatello. Exploratory causal inference in saence, 2026. URL <https://arxiv.org/abs/2510.14073>.
- Moritz Miller, Bernhard Schölkopf, and Siyuan Guo. Counterfactual reasoning: an analysis of in-context emergence, 2025. URL <https://arxiv.org/abs/2506.05188>.
- Bruno A. Olshausen and David J. Field. Sparse coding with an overcomplete basis set: A strategy employed by v1? *Vision Research*, 37(23):3311–3325, 1997. ISSN 0042-6989. doi: [https://doi.org/10.1016/S0042-6989\(97\)00169-7](https://doi.org/10.1016/S0042-6989(97)00169-7). URL <https://www.sciencedirect.com/science/article/pii/S0042698997001697>.

- Kiho Park, Yo Joong Choe, and Victor Veitch. The linear representation hypothesis and the geometry of large language models. In *Proceedings of the 41st International Conference on Machine Learning, ICML'24*. JMLR.org, 2024.
- Gonalo Paulo and Nora Belrose. Sparse autoencoders trained on the same data learn different features. *arXiv preprint arXiv:2501.16615*, 2025.
- Gonalo Paulo, Alex Mallen, Caden Juang, and Nora Belrose. Automatically interpreting millions of features in large language models. *arXiv preprint arXiv:2410.13928*, 2024.
- Judea Pearl. *Causality: Models, Reasoning and Inference*. Cambridge University Press, 2nd edition, 2009. URL <https://bayes.cs.ucla.edu/B00K-2K/>.
- Jonas Peters, Peter B hlmann, and Nicolai Meinshausen. Causal inference by using invariant prediction: Identification and confidence intervals. *Journal of the Royal Statistical Society Series B: Statistical Methodology*, 78(5):947–1012, 10 2016. ISSN 1369-7412. doi: 10.1111/rssb.12167. URL <https://doi.org/10.1111/rssb.12167>.
- Jonas Peters, Dominik Janzing, and Bernhard Sch lkopf. *Elements of Causal Inference: Foundations and Learning Algorithms*. The MIT Press, 2017. ISBN 0262037319.
- Karin Schnass and Pierre Vanderghenst. Dictionary preconditioning for greedy algorithms. *IEEE Transactions on Signal Processing*, 56(5):1994–2002, 2008.
- B. Sch lkopf, F. Locatello, S. Bauer, N. R. Ke, N. Kalchbrenner, A. Goyal, and Y. Bengio. Toward causal representation learning. *Proceedings of the IEEE*, 109(5):612–634, 2021. doi: 10.1109/JPROC.2021.3058954. URL <https://ieeexplore.ieee.org/stamp/stamp.jsp?arnumber=9363924>. *equal contribution.
- P. Spirtes, C. Glymour, and R. Scheines. *Causation, Prediction, and Search*. MIT press, 2nd edition, 2000.
- S. Waldron. Generalized welch bound equality sequences are tight frames. *IEEE Transactions on Information Theory*, 49(9):2307–2309, 2003. doi: 10.1109/TIT.2003.815788.
- Shangshang Wang, Julian Asilis,  mer Faruk Akg l, Enes Burak Bilgin, Ollie Liu, Deqing Fu, and Willie Neiswanger. Resa: Transparent reasoning models via saes, 2025. URL <https://arxiv.org/abs/2506.09967>.
- L. Welch. Lower bounds on the maximum cross correlation of signals (corresp.). *IEEE Transactions on Information Theory*, 20(3):397–399, 1974. doi: 10.1109/TIT.1974.1055219.
- Longhui Yu, Weisen Jiang, Han Shi, Jincheng Yu, Zhengying Liu, Yu Zhang, James T Kwok, Zhenguo Li, Adrian Weller, and Weiyang Liu. Metamath: Bootstrap your own mathematical questions for large language models. *arXiv preprint arXiv:2309.12284*, 2023.

A Self-coherence and Welch bounds

A.1 Self-coherence

To illustrate the relevance of self-coherence for sparsity, we restate the below theorem (Donoho et al., 2005) and provide the proof. Theorem A.1 reflects a clear relationship between sparsity and self-coherence. Note that we define self-coherence here as the maximum interference between features, $\mu_{\max}$. Below, we discuss how they relate to (2) through the Welch bounds. First, low self-coherence corresponds to stricter almost orthogonality between dictionary atoms. Second, this yields a bound on the number of atoms required for input recovery K to ensure identifiability.

Theorem A.1 (Self-coherence bound for uniqueness). *Let $\mathbf{D} \in \mathbb{R}^{m \times d}$ have unit-norm columns and denote by $\mu_{\max}$ self-coherence as defined in (2). If*

$$K < \frac{1}{2} \left(1 + \frac{1}{\mu_{\max}} \right),$$

every K -sparse representation $\tilde{\mathbf{x}} = \mathbf{D}\mathbf{z}$ is unique. That is, for $\mathbf{D}\mathbf{z} = \mathbf{D}\tilde{\mathbf{z}}$ with $\|\mathbf{z}\|_0, \|\tilde{\mathbf{z}}\|_0 \leq K$, it holds that $\mathbf{z} = \tilde{\mathbf{z}}$.

Proof. Assume $\tilde{\mathbf{x}} = \mathbf{D}\mathbf{z} = \mathbf{D}\tilde{\mathbf{z}}$ with $\mathbf{z} \neq \tilde{\mathbf{z}}$, both K -sparse. Then $\mathbf{s} = \mathbf{z} - \tilde{\mathbf{z}} \neq \mathbf{0}$ satisfies $\mathbf{D}\mathbf{s} = \mathbf{0}$. Let $T = \text{supp}(\mathbf{s})$, then $|T| \leq 2K$. Let $\mathbf{G} = \mathbf{D}_T^\top \mathbf{D}_T$. All diagonal elements of $\mathbf{G}$ are equal to 1 with off-diagonal elements upper-bounded by $\mu_{\max}$. Formally, $\mathbf{G}_{j,j} = 1, \forall j \in [d]$ and $\mathbf{G}_{i,j} \leq \mu_{\max}$ for $i \neq j$ with $i, j \in [d]$. By Gershgorin's circle theorem (Gershgorin, 1931),

$$\lambda_{\min}(\mathbf{G}) \geq 1 - (2K - 1)\mu_{\max} > 0,$$

by assumption on K . Thus $\mathbf{G}$ is positive definite and the atoms in T are independent. We arrive at a contradiction. $\square$

A.2 Welch bounds

For $\{\mathbf{f}_j\}_{j \in [d]}$ unit-norm vectors with $\mathbf{f}_j \in \mathbb{R}^m$, define the maximum similarity by $\mu_{\max} = \max_{j \neq k} |\langle \mathbf{f}_j, \mathbf{f}_k \rangle|$. Then, for $s \in \mathbb{N}_{>0}$,

$$\mu_{\max}^{2s} \geq \frac{1}{d-1} \left[\frac{d}{\binom{m+s-1}{s}} - 1 \right],$$

with $\mathbb{N}_{>0}$ denoting the non-zero natural numbers. In addition, Welch obtains a bound on the mean similarity, which we state in (3),

$$\frac{1}{d^2} \sum_{j=1}^d \sum_{k=1}^d |\langle \mathbf{f}_j, \mathbf{f}_k \rangle|^{2s} \geq \frac{1}{\binom{m+s-1}{s}}$$

for which Waldron (2003) relaxes the unit-norm assumption,

$$\frac{\sum_{j=1}^d \sum_{k=1}^d |\langle \mathbf{f}_j, \mathbf{f}_k \rangle|^{2s}}{\left(\sum_{j=1}^d \|\mathbf{f}_j\| \right)^{2s}} \geq \frac{1}{\binom{m+s-1}{s}}.$$

with $\mathbf{f}_j$ and $\mathbf{f}_k$ arbitrary.

B Post-intervention interference between features

B.1 Causal intervention mismatch between the true output and the model output

Theorem B.1 (Upper bound on interference). *Let $\mathbf{D} = [\mathbf{f}_1, \dots, \mathbf{f}_d] \in \mathbb{R}^{m \times d}$ and assume that*

$$\|\mathbf{f}_j\|_2 = 1, \quad j \in [d].$$

Fix $j \in [d]$ and $\delta \in \mathbb{R}$. Then

$$\mathbb{E} \left[\|Y^{(j,\delta)} - \hat{Y}^{(j,\delta)}\|_2^2 \right] \leq L^2 \|\mathbf{A}\|_{\text{op}}^2 \mathbb{E} \left[\|(\mathbf{G} - \mathbf{I})(\mathbf{Z} + \delta \mathbf{e}_j)\|_2^2 \right].$$

where L is a bound on the operator norm of the Jacobian $\boldsymbol{\eta} \mapsto p_{\boldsymbol{\eta}}$. Consequently,

$$\mathbb{E} \left[\|Y^{(j,\delta)} - \hat{Y}^{(j,\delta)}\|_2^2 \right] \leq 2L^2 \|\mathbf{A}\|_{\text{op}}^2 \left(\mathbb{E} [\|(\mathbf{G} - \mathbf{I})\mathbf{Z}\|_2^2] + \delta^2 \sum_{k \neq j} \langle \mathbf{f}_k, \mathbf{f}_j \rangle^2 \right).$$

For $\mu(\mathbf{D})$ as defined in (2), averaging over j yields

$$\frac{1}{d} \sum_{j=1}^d \mathbb{E} \left[\|Y^{(j,\delta)} - \hat{Y}^{(j,\delta)}\|_2^2 \right] \leq 2L^2 \|\mathbf{A}\|_{\text{op}}^2 \mathbb{E}[\|(\mathbf{G} - \mathbf{I})Z\|_2^2] + 2L^2 \|\mathbf{A}\|_{\text{op}}^2 \delta^2 (d-1) \mu(\mathbf{D}),$$

In the standard softmax case, in particular, one may take $L = \frac{1}{2}$ such that

$$\mathbb{E} \left[\|Y^{(j,\delta)} - \hat{Y}^{(j,\delta)}\|_2^2 \right] \leq \frac{1}{4} \|\mathbf{A}\|_{\text{op}}^2 \mathbb{E} \left[\|(\mathbf{G} - \mathbf{I})(Z + \delta \mathbf{e}_j)\|_2^2 \right].$$

Proof. Under the intervention $\text{do}(Z_j := Z_j + \delta)$, the latent state becomes

$$Z^{(j,\delta)} = Z + \delta \mathbf{e}_j.$$

The true post-intervention output distribution is therefore parameterized by

$$\boldsymbol{\eta}_Y^{(j,\delta)} := \mathbf{A}(Z + \delta \mathbf{e}_j),$$

while the model post-intervention output distribution is parameterized by

$$\boldsymbol{\eta}_{\hat{Y}}^{(j,\delta)} := \mathbf{A}\mathbf{G}(Z + \delta \mathbf{e}_j).$$

Subtracting gives

$$\boldsymbol{\eta}_{\hat{Y}}^{(j,\delta)} - \boldsymbol{\eta}_Y^{(j,\delta)} = \mathbf{A}(\mathbf{G} - \mathbf{I})(Z + \delta \mathbf{e}_j).$$

Hence,

$$\|\boldsymbol{\eta}_{\hat{Y}}^{(j,\delta)} - \boldsymbol{\eta}_Y^{(j,\delta)}\|_2 \leq \|\mathbf{A}\|_{\text{op}} \|(\mathbf{G} - \mathbf{I})(Z + \delta \mathbf{e}_j)\|_2.$$

We now pass from parameters to output distributions. Since the Jacobian of $\boldsymbol{\eta} \mapsto p_{\boldsymbol{\eta}}$ is bounded by L , the mean value theorem implies

$$\|p_{\boldsymbol{\eta}} - p_{\boldsymbol{\eta}'}\|_2 \leq L \|\boldsymbol{\eta} - \boldsymbol{\eta}'\|_2 \quad \text{for all } \boldsymbol{\eta}, \boldsymbol{\eta}' \in \mathbb{R}^q.$$

Applying this with

$$\boldsymbol{\eta} = \boldsymbol{\eta}_Y^{(j,\delta)}, \quad \boldsymbol{\eta}' = \boldsymbol{\eta}_{\hat{Y}}^{(j,\delta)},$$

yields

$$\|Y^{(j,\delta)} - \hat{Y}^{(j,\delta)}\|_2 \leq L \|\boldsymbol{\eta}_{\hat{Y}}^{(j,\delta)} - \boldsymbol{\eta}_Y^{(j,\delta)}\|_2.$$

Combining the two inequalities, we obtain

$$\|Y^{(j,\delta)} - \hat{Y}^{(j,\delta)}\|_2 \leq L \|\mathbf{A}\|_{\text{op}} \|(\mathbf{G} - \mathbf{I})(Z + \delta \mathbf{e}_j)\|_2.$$

Squaring both sides and taking expectation proves the first bound,

$$\mathbb{E} \left[\|Y^{(j,\delta)} - \hat{Y}^{(j,\delta)}\|_2^2 \right] \leq L^2 \|\mathbf{A}\|_{\text{op}}^2 \mathbb{E} \left[\|(\mathbf{G} - \mathbf{I})(Z + \delta \mathbf{e}_j)\|_2^2 \right].$$

To obtain the more explicit estimate, expand

$$(\mathbf{G} - \mathbf{I})(Z + \delta \mathbf{e}_j) = (\mathbf{G} - \mathbf{I})Z + \delta(\mathbf{G} - \mathbf{I})\mathbf{e}_j.$$

Using

$$\|\mathbf{u} + \mathbf{v}\|_2^2 \leq 2\|\mathbf{u}\|_2^2 + 2\|\mathbf{v}\|_2^2,$$

we get

$$\|(\mathbf{G} - \mathbf{I})(Z + \delta \mathbf{e}_j)\|_2^2 \leq 2\|(\mathbf{G} - \mathbf{I})Z\|_2^2 + 2\delta^2\|(\mathbf{G} - \mathbf{I})\mathbf{e}_j\|_2^2.$$

Substituting into the previous inequality gives

$$\mathbb{E} \left[\|Y^{(j,\delta)} - \hat{Y}^{(j,\delta)}\|_2^2 \right] \leq 2L^2 \|\mathbf{A}\|_{\text{op}}^2 \left(\mathbb{E}[\|(\mathbf{G} - \mathbf{I})Z\|_2^2] + \delta^2\|(\mathbf{G} - \mathbf{I})\mathbf{e}_j\|_2^2 \right).$$

It remains to compute $\|(\mathbf{G} - \mathbf{I})\mathbf{e}_j\|_2^2$. Since $\mathbf{G} = \mathbf{D}^\top \mathbf{D}$, we have

$$\mathbf{G}_{kj} = \langle \mathbf{f}_k, \mathbf{f}_j \rangle.$$

Because the columns are unit norm, $\mathbf{G}_{jj} = 1$. Thus

$$((\mathbf{G} - \mathbf{I})\mathbf{e}_j)_k = \begin{cases} 0, & k = j, \\ \langle \mathbf{f}_k, \mathbf{f}_j \rangle, & k \neq j, \end{cases}$$

and therefore

$$\|(\mathbf{G} - \mathbf{I})\mathbf{e}_j\|_2^2 = \sum_{k \neq j} \langle \mathbf{f}_k, \mathbf{f}_j \rangle^2.$$

Substituting this identity proves the second bound.

Finally, averaging over j , we obtain

$$\frac{1}{d} \sum_{j=1}^d \mathbb{E} \left[\|Y^{(j,\delta)} - \hat{Y}^{(j,\delta)}\|_2^2 \right] \leq 2L^2 \|\mathbf{A}\|_{\text{op}}^2 \mathbb{E}[\|(\mathbf{G} - \mathbf{I})Z\|_2^2] + 2L^2 \|\mathbf{A}\|_{\text{op}}^2 \delta^2 \frac{1}{d} \sum_{j=1}^d \sum_{k \neq j} \langle \mathbf{f}_k, \mathbf{f}_j \rangle^2.$$

By the definition of $\mu(\mathbf{D})$ in (2),

$$\sum_{j=1}^d \sum_{k \neq j} \langle \mathbf{f}_k, \mathbf{f}_j \rangle^2 = d(d-1)\mu(\mathbf{D}),$$

we obtain the averaged inequality. $\square$

B.2 K-sparse latent states.

Suppose that Z has uniformly random support of size K , and that conditioned on its support the nonzero coordinates are independent, mean zero, and with variance σ^2 . Then,

$$\mathbb{E}[Z_i^2] = \frac{K}{d}\sigma^2, \quad \mathbb{E}[Z_i Z_j] = 0 \quad (i \neq j),$$

so

$$\mathbb{E}[ZZ^\top] = \frac{K}{d}\sigma^2 \mathbf{I}.$$

Therefore

$$\mathbb{E}[\|(\mathbf{G} - \mathbf{I})Z\|_2^2] = \mathbb{E}[Z^\top (\mathbf{G} - \mathbf{I})^2 Z] = \text{tr}((\mathbf{G} - \mathbf{I})^2 \mathbb{E}[ZZ^\top]) = \frac{K}{d}\sigma^2 \text{tr}((\mathbf{G} - \mathbf{I})^2).$$

Since $\mathbf{G} - \mathbf{I}$ is symmetric,

$$\text{tr}((\mathbf{G} - \mathbf{I})^2) = \|\mathbf{G} - \mathbf{I}\|_F^2,$$

and because $\mathbf{G} = \mathbf{D}^\top \mathbf{D}$ with unit-norm columns,

$$\|\mathbf{G} - \mathbf{I}\|_F^2 = \sum_{j=1}^d \sum_{k \neq j} \langle \mathbf{f}_k, \mathbf{f}_j \rangle^2 = d(d-1)\mu(\mathbf{D}).$$

Hence

$$\mathbb{E}[\|(\mathbf{G} - \mathbf{I})Z\|_2^2] = K\sigma^2(d-1)\mu(\mathbf{D}).$$

Substituting this into the averaged bound of the theorem yields

$$\frac{1}{d} \sum_{j=1}^d \mathbb{E} \left[\|Y^{(j,\delta)} - \hat{Y}^{(j,\delta)}\|_2^2 \right] \leq 2L^2 \|\mathbf{A}\|_{\text{op}}^2 (d-1)\mu(\mathbf{D}) (K\sigma^2 + \delta^2).$$

In particular, in the standard softmax case,

$$\frac{1}{d} \sum_{j=1}^d \mathbb{E} \left[\|Y^{(j,\delta)} - \hat{Y}^{(j,\delta)}\|_2^2 \right] \leq \frac{1}{2} \|\mathbf{A}\|_{\text{op}}^2 (d-1)\mu(\mathbf{D}) (K\sigma^2 + \delta^2).$$




Figure 6: **Intervention mismatch** We plot the discrepancy between predicted and ground truth output for different levels of self-coherence in an overcomplete dictionary. The dotted lines represent the bound derived in Theorem 4.1 for the three self-coherence levels.

B.3 Empirical evidence

In the interference context, we use the terms “discrepancy” and “mismatch” interchangeably. Figure 6 depicts the discrepancy for different levels of self-coherence. The **medium** self-coherence of 0.138 has been achieved by a decoder $\mathbf{D}$ with random Gaussian, L2-normalized columns. The **high** self-coherence of 0.855 we obtain by creating columns which are all small perturbations of a shared direction. Constructing the overcomplete dictionary in a structured way, we reach **low** self-coherence of 0.071. In addition, we display the corresponding self-coherence bounds derived in Theorem 4.1. In this example, the bounds are not tight. We observe, however, that the bound for the **high** self-coherence setting is up to an order of magnitude above the upper bound for **low** self-coherence.

In spirit, this is similar to our approach of fine-tuning the SAE under an orthogonality penalty. There, too, we aim to add low self-coherence as an inductive bias to enable isolated interventions. Here, we take a full orthonormal basis in $\mathbb{R}^{m \times m}$ for $m = 8$ and add the remaining $d - m = 12 - 8 = 4$ columns through Givens-rotated combinations. In particular, for $j \in [d - m]$, we define the angle $\theta_j = \frac{j\pi}{d-m+1}$ and set

$$\mathbf{e}_j = \cos(\theta_j) \mathbf{f}_{j \bmod m} + \sin(\theta_j) \mathbf{f}_{(j+1) \bmod m} \quad (5)$$

for $\{\mathbf{f}_j\}_{j \in [m]}$ the orthonormal basis vectors. We then fill the $d - m$ columns of $\mathbf{D}$ through $\{\mathbf{e}_j\}_{j \in [d-m]}$ and obtain the final dictionary by renormalizing each column to unit norm.

Finally, Figure 7 represents the gram matrix $\mathbf{G}$ for each of the three levels of self-coherence as a heatmap. While the interference between features increases with self-coherence, we can attribute the **low** self-coherence behavior on the off-diagonals to the generation of the dictionary $\mathbf{D}$ in (5).




Figure 7: Coherence of decoders for different levels of self-coherence

C Experimental Details

C.1 Setup

All experiments were conducted using multi-GPU training. Fine-tuning the SAE on the MetaMathQA dataset required approximately 8 hours using 4 GPUs with 80GB of memory each. Subsequent fine-tuning of the language model with the SAE integrated into the architecture required an additional 16 hours under the same hardware configuration.

We employed parameter-efficient fine-tuning via LoRA with rank 128 and dropout rate 0.05. The underlying model has a hidden dimension of 2304. Training was performed with a per-device batch size of 4. We used an initial learning rate of 5×10^{-5} , combined with 200 warm-up steps followed by a cosine decay schedule. Fine-tuning the LLM was carried out with gradient checkpointing enabled to reduce memory usage, and gradient norms were clipped to a maximum value of 1 to ensure training stability. During generation, we greedily decode all tokens.

C.2 Interpretability

For 2’000 features per orthogonality penalty, we generate a natural-language description and evaluate whether it correctly explains the feature’s activation behavior. Descriptions are constrained to a maximum length of 20 words. For each feature, we provide the evaluator with up to 20 text spans sampled from the top 100 most strongly activating snippets in the GSM8K test set. The dataset follows a 95-1-4 train-validation-test split over 395’000 examples. Each span consists of a window of ± 10 tokens around the activation point.

You are evaluating the interpretability of an explanation with respect to a set of text snippets.

You will be given:

- A list of `{total_snippets}` labeled text snippets (`{snippet_labels}`)
- A single explanation describing a specific concept, feature, or pattern

Your task is to identify which snippet the explanation applies to MOST STRONGLY.

Exactly `{correct_snippets}` of the provided snippets is correct.

Guidelines:

- Select ONLY the single snippet for which the explanation is clearly and directly applicable.
- You must output exactly one label.
- Do not explain your reasoning.

Output format requirements (MANDATORY):

- Output must be a Python-style list of snippet labels.
- Example valid outputs: `[1]`, `[{total_snippets}]`

- Do NOT include any additional text, punctuation, or explanation.
- Do NOT include quotes around labels.
- Do NOT include reasoning or commentary.

If the explanation best fits snippet 2, your entire output must be:
[2]

Listing 2: System Prompt for Interpretability Score

We evaluate interpretability over 300 feature explanations using a multiple-choice setup, where the evaluator selects the span best matching the provided description from five candidates. All evaluations are performed using Llama 3.1 8B-Instruct. We include the full system prompt in Listing 2. A representative example of feature explanation is given in Listing 1.




Figure 8: **Interpretability score** We plot the interpretability score of correctly identifying one out of five examples, which relates closest to the provided explanation. Error bars are basic bootstrap confidence intervals (Efron, 1979) with 100 resamples.

C.3 Interpretability results

SAEs are primarily used for mechanistic interpretability. By uncovering human-understandable concepts, we aim to track the Transformer’s internal computations. We therefore check the interpretability of our SAEs under different orthogonality penalties. For this, we create LM-generated explanations for 2’000 activated features in each of our fine-tuned SAEs. We do so providing 20 text snippets which activate the respective feature and asking Llama 3.1 8B-Instruct to describe the common theme across snippets. For details and prompt templates, see below.

We then ask Llama 3.1 8B-Instruct to match one of five text snippets to the corresponding feature explanation. Figure 8 demonstrates our results. Without allowing the model to reason, we observe interpretability performance on this task

between 37% and 47%. Given the basic bootstrap error bars (Efron, 1979), we observe non-significant difference between orthogonality parameters. We also notice significantly better interpretability than random guessing at 20%.

Finally, we note that any feature we report has the same feature index across all tested values of λ . Hence, LoRA adapts the LM around the SAE without repurposing feature indices. This indicates that the orthogonality regularization reshapes the geometric relationships between features while leaving the encoded concepts unchanged. Together with the stable performance in Figure 3 and interpretability scores in Figure 8, this provides evidence that our pipeline does not induce unintended semantic drift.

You are labeling latent features of a language model trained on grade-school math word problems.

You will see {num_spans} numbered spans from these problems. In each span, exactly one token is wrapped in double square brackets, like [[this]]. The feature you are labeling fires on precisely the bracketed token; the surrounding text is context only.

Procedure:

1. Resolve each bracketed token to the full word or entity it belongs to. A bracketed fragment of a longer word stands for the whole word: [[umbr]] inside "umbrella" means umbrella.
2. Find the single most specific concept shared by these resolved

- words across (nearly) all {num_spans} spans.
3. Name that concept as specifically as the evidence allows: a concrete entity (a particular first name, animal, food or drink, place, object); else a specific mathematical element (a particular unit, operation word, number format, quantity type); else a structural role.
- Constraints:
- Your answer must name something you can point to in the bracketed tokens themselves. Never answer with a word that does not appear in, or directly behind, the bracketed tokens.
 - The illustration word in these instructions ("umbrella") is not evidence; no word from these instructions may appear in your answer unless it genuinely occurs in the spans.
 - State the most specific level the spans support: the particular name, animal, or unit itself, not its category ("first names", "animals", "units").
 - Do NOT describe features as "mathematical", "math-related", "word problems", "numbers", or "reasoning".
 - Do not mention the brackets, tokens, models, or activations in your answer.
 - Output exactly one sentence (max {max_words} words) in the exact form: "The shared latent concept is ...".
 - Your sentence must contain the specific shared word or name itself, copied verbatim from the spans and wrapped in quotation marks – category plus instance, e.g.: The shared latent concept is the first name "X". / the unit "X". / the animal "X".
 - If the bracketed tokens resolve to different words that share a category, name the category and quote two or three of the actual words, e.g.: first names such as "X" and "Y".
 - A category alone ("a first name", "an animal", "a unit") is never an acceptable answer. If you cannot quote the specific word(s), reply exactly: "No coherent concept."

Listing 3: System Prompt for Feature Explanations


C.4 Intervenability

To evaluate the extent to which learned features support localized and controllable interventions, we conduct targeted manipulation experiments. We identify 12 SAE features corresponding to male first names. We observe that the model assigns the same feature indices to different variations of a name. We thus list in [Listing 4](#) all variations of names which we accept as correct. As our intervenability study hinges on string extraction, the spelling of the concept is relevant. Across all models considered, these concepts are consistently represented by the same feature indices. Before running the experiment, we check that all names are recoverable in at least some contexts.


For each name, we intervene by swapping the corresponding feature activation at every token position where the name appears. We sweep over the following values for the associated coefficient $z_j, j \in [d]$,

$$\{10, 20, 50, 100, 150, 200, 300, 500\}. \quad (6)$$


By doing so, we find that the accuracy ranges from 0.601 to 0.658 for any combination of insertion value and orthogonality penalty. We then decide the hyperparameter based on the models' ability to remove the intervened on concept. For sufficiently high insertion value, the model near-perfectly drops the intervened on concept from the generation. Hence, we choose among the set of hyperparameters that includes at most 1% of concepts we intervene on. This leaves us with hyperparameter values greater than or equal to 100. For our experiments, we choose hyperparameter 200. We show similar behavior for $z_j \in \{100, 150, 300\}$ in [Figure 9](#).




(a) **Evaluation on mathematical reasoning** We plot mathematical reasoning performance for insertion hyperparameter 100.




(b) **Evaluation on inclusion of correct name** We plot the fraction of correctly included names for insertion hyperparameter 100.




(c) **Evaluation on mathematical reasoning** We plot mathematical reasoning performance for insertion hyperparameter 150.



(d) **Evaluation on inclusion of correct name** We plot the fraction of correctly included names for insertion hyperparameter 150.




(e) **Evaluation on mathematical reasoning** We plot mathematical reasoning performance for insertion hyperparameter 300.



(f) **Evaluation on inclusion of correct name** We plot the fraction of correctly included names for insertion hyperparameter 300.

Figure 9: **Intervenability across hyperparameter sweep** We intervene on 12 features in the SAE corresponding to the concept of first names. Then, we evaluate the ability to maintain reasoning performance and the correct generation of the included first name.



(a) **Evaluation on mathematical reasoning** We plot mathematical reasoning performance on animal-related queries after intervening on the SAE. Error bars are basic bootstrap confidence intervals (Efron, 1979) with 100 samples.

(b) **Evaluation on inclusion of correct animal** We plot the fraction of correctly included animal names after intervening on the corresponding feature in the SAE. Error bars are basic bootstrap confidence intervals of 100 random draws.

Figure 10: **Intervenability** We intervene on 5 features in the SAE corresponding to the concept of animal names. Then, we evaluate the ability to maintain reasoning performance and the correct generation of the included animal name. We choose an insertion value of 200.

Choosing $z_j = 200$, we evaluate both *drop* and *include* interventions. While all models successfully suppress the original concept under drop interventions, including a new concept proves more challenging for models with weaker orthogonality. We hypothesize that this asymmetry arises from differences in the alignment of feature directions within the embedding space. Beyond names, we also observe successful semantic interventions on other concepts; for example, activating an [aqua](#) feature causes the model to transform the character [Mike](#) into [Aquaman](#) while leaving the surrounding context unchanged.

```
{
  "Jason": ["Jason", "Jase"],
  "Mike": ["Michael", "Mike", "Mikey"],
  "Jacob": ["Jacob", "Jake", "Jakob"],
  "Jerry": ["Jerry", "Jeremy", "Jermey", "Jerome"],
  "James": ["James", "Jim", "Jimmy", "Jamie"],
  "Robert": ["Robert", "Rob", "Robbie"],
  "Jordan": ["Jordan", "Jordy"],
  "Jackson": ["Jackson", "Jack", "Jax"],
  "Paul": ["Paul", "Pauly"],
  "David": ["David", "Dave", "Davey"],
  "Andrew": ["Andrew", "Andy"],
  "Gary": ["Gary", "Garre", "Garret", "Garrett", "Garry"]
}
```

Listing 4: **First names used for interventions**

C.5 Intervenability on animals and locations

```
{
  "animals": ["dog", "rabbit", "bear", "goat", "horse"],
  "locations": ["park", "room", "school", "library", "fast-food
    restaurant"],
}
```

Listing 5: **Additional concepts used for interventions**

Similar to our analysis on first names, we run experiments on two additional groups. We notice now that for sentences where correct retrieval is simple for all models, the model with

the strictest orthogonality penalty of 10^{-4} obtains higher accuracy than the non-penalized model. This can be seen for the animal dataset at insertion value 200 in Figure 10, for instance, at an accuracy of 72.5% compared to 62.8%. The five animal concepts are shown in Listing 5. While slightly higher at 93.4% relative to 90.6%, correct feature inclusion between strictest penalty and no-penalty setup are comparable. Under all setups, all features corresponding to dropped concepts are adequately excluded.

The results for the location concept are found in Table 5. On the one hand, we observe above-average evaluation scores of up to 95.7% for $\lambda = 10^{-4}$ at modest insertion value of 10. We suspect this to be an artifact of our dataset generation. As several sentences fall into the category of part-whole or other simple arithmetic problems, the evaluation scores tend to be higher than for the full dataset in Figure 4. We expect these problems to be easier to solve for the LM.

On the other hand, the inclusion values for the location concept are relatively small. Here, we visually inspect the ground truth reasoning traces in GSM8K of the original sentences. We notice that the answer often does not restate concept in the corpus. In other words, the model is not necessarily fine-tuned on restating the concept explicitly to arrive at the final answer. As the location regularly serves as the place description in the sentence, it is only tangentially relevant to the computation of the right solution. Therefore, it is natural to omit this during generation. However, we find these results relevant because they enhance our understanding beyond the subject-level first names. We refer the reader to Listing 8.

C.6 Behavioral and internal measures of interference

We extend our above analysis to the animal and location dataset. The results are reported in Tables 2 and 3.

λ	ROUGE-L RECALL		D_{JSD}	
	Mean	Confidence interval	Mean	Confidence interval
0	0.761	[0.741, 0.774]	0.259	[0.248, 0.272]
10^{-6}	0.793	[0.780, 0.804]	0.243	[0.233, 0.252]
10^{-5}	0.797	[0.785, 0.809]	0.229	[0.220, 0.241]
10^{-4}	0.819	[0.803, 0.835]	0.246	[0.234, 0.259]

Table 2: **Results for interference measures** We evaluate ROUGE-L RECALL and D_{JSD} on the tokens and features of the animal dataset. We report mean and the basic bootstrap confidence interval over all observations.

λ	ROUGE-L RECALL		D_{JSD}	
	Mean	Confidence interval	Mean	Confidence interval
0	0.898	[0.879, 0.914]	0.236	[0.202, 0.269]
10^{-6}	0.894	[0.867, 0.923]	0.246	[0.214, 0.276]
10^{-5}	0.871	[0.843, 0.893]	0.251	[0.220, 0.278]
10^{-4}	0.894	[0.879, 0.905]	0.182	[0.167, 0.198]

Table 3: **Results for interference measures** We evaluate ROUGE-L RECALL and D_{JSD} on the tokens and features of the location dataset. We report mean and the basic bootstrap confidence interval over all observations.

C.7 Dead features

Figure 11 demonstrates that the configuration with orthogonality penalty $\lambda = 10^{-4}$ only uses approximately two thirds of the features of the non-regularized configuration. In

Section 6, we hypothesize that this constitutes an inherent characteristic of the strict orthogonality penalty.





Figure 11: **Dead features** We plot the fraction of dead features with increasing orthogonality penalty. Error bars are basic bootstrap confidence intervals (Efron, 1979).

C.8 Experiments on Llama 3.2 1B

Building on the above experiments, we run the pipeline on Llama 3.2 1B to confirm our results. Our starting point is the FaithfulSAE (Cho et al., 2025) optimized for $K = 48$. For easier comparability, we still run our two-stage fine-tuning pipeline on $K = 20$. Similar to Cho et al. (2025), we insert the SAE after layer 12. Architecturally, Llama 3.2 1B consists of 16 layers instead of the 26 Gemma 2 2B layers. Further, the SAE dimension is lower at 14/336 compared to 65/536. On top of intervening on one of the layers in the early half of the Gemma-2-2B transformer, the Llama 3.2 1B experiments strengthen our claim by inserting the SAE three-quarters of the way through the network.



(a) **Evaluation on mathematical reasoning** We plot mathematical reasoning performance after intervening on the SAE.

(b) **Evaluation on including correct name** We plot the fraction of correctly included first names after intervention in the SAE.

(c) **Evaluation on dropping the first name** We plot the fraction of mistakenly included first names after dropping feature.

Figure 12: **Intervenability** We intervene on 180 examples containing male first names. Then, we evaluate the ability to maintain reasoning performance and the correct generation of the included first name. We choose an insertion value of 20. Error bars are basic bootstrap confidence intervals (Efron, 1979) with 100 samples.

In contrast to the Gemma 2 2B experiments, we observe non-zero inclusion of the intervened on features for all orthogonality penalties and insertion values. We therefore do not compare the models conditional on properly excluding the dropped feature but choose the insertion value which yields highest accuracy. We thus also report that statistic in Figure 12c. Sweeping over the values for the coefficient associated with the intervention feature indices in (6), we note highest accuracy for an insertion value of 20 for all configurations except $\lambda = 10^{-5}$, for which this is 10. We thus choose 20 but observe similar results for 10. We again intervene on gendered first names and evaluate on 180 snippets. In line with the relatively smaller model, we observe lower accuracy across all configurations. Figure 12a indicates

better performance for the no-penalty and $\lambda = 10^{-6}$ configurations compared to stricter setups. Due to the width of the basic bootstrap confidence intervals, however, we cannot claim significantly worse performance of the $\lambda = 10^{-4}$ configuration. Primarily interested in intervenability, we note in [Figure 12b](#) that first name inclusion monotonically increases with stricter orthogonality penalty. According to the basic bootstrap intervals reported, this is also significant. Moreover, [Figure 12c](#) plots the fraction of included first names after removing the corresponding SAE feature from the residual stream. Note that lower inclusion of incorrect indices indicates crisper interventions. While non-zero, the two strictest configurations drop significantly more first names than the settings with $\lambda \in \{0, 10^{-6}\}$.

We understand this as additional support for our claim that orthogonality regularization improves intervenability inside the SAE. Tested on two different models at different SAE sizes and positions, we obtain qualitatively similar results. On the one hand, mathematical reasoning accuracy neither improves nor falls significantly. Interventions, on the other hand, are significantly crisper for stricter orthogonality. We show this both in terms of dropping the intervened on SAE feature as well as including the feature we manually add to the residual stream.

C.9 Intervenability datasets

```
{
  "0": {"sentence": [f"{first_name} buys twice as many red ties as blue
    ties. The red ties cost 50% more than blue ties. He spent $200 on
    blue ties that cost $40 each. How much did he spend on ties?"],
    "ground_truth": 800},
  "1": {"sentence": [f"{first_name} sold clips to 48 of his friends in
    April, and then he sold half as many clips in May. How many clips
    did {first_name} sell altogether in April and May?"],
    "ground_truth": 72},
  "2": {"sentence": [f"{first_name} writes a 3-page letter to 2 different
    friends twice a week. How many pages does he write a year?"],
    "ground_truth": 624},
  "3": {"sentence": [f"{first_name} is wondering how much pizza he can eat
    in one day. He buys 2 large pizzas and 2 small pizzas. A large pizza
    has 16 slices and a small pizza has 8 slices. If he eats it all, how
    many pieces does he eat that day?"], "ground_truth": 48},
  "4": {"sentence": [f"{first_name} created a care package to send to his
    brother, who was away at boarding school. {first_name} placed a box
    on a scale, and then he poured into the box enough jelly beans to
    bring the weight to 2 pounds. Then, he added enough brownies to
    cause the weight to triple. Next, he added another 2 pounds of jelly
    beans. And finally, he added enough gummy worms to double the weight
    once again. What was the final weight of the box of goodies, in
    pounds?"], "ground_truth": 16},
  "5": {"sentence": [f"{first_name} can read 8 pages of a book in 20
    minutes. How many hours will it take him to read 120 pages?"],
    "ground_truth": 5},
  "6": {"sentence": [f"{first_name} is fundraising for his charity by
    selling brownies for $3 a slice and cheesecakes for $4 a slice. If
    {first_name} sells 43 brownies and 23 slices of cheesecake, how much
    money does {first_name} raise?"], "ground_truth": 221},
  "7": {"sentence": [f"{first_name} is putting together 4 tables. Each
    table has 4 legs and each leg needs 2 screws. He has 40 screws. How
    many screws will he have left over?"], "ground_truth": 8},
  "8": {"sentence": [f"{first_name} buys 10 packs of magic cards. Each
    pack has 20 cards and 1/4 of those cards are uncommon. How many
    uncommon cards did he get?"], "ground_truth": 50},
  "9": {"sentence": [f"{first_name} creates a media empire. He creates a
    movie for $2000. Each DVD cost $6 to make. He sells it for 2.5 times
    that much. He sells 500 movies a day for 5 days a week. How much
    profit does he make in 20 weeks?"], "ground_truth": 448000},
  "10": {"sentence": [f"{first_name} went to a shop to buy some groceries.
    He bought some bread for $2, butter for $3, and juice for two times
    the price of the bread. He had $15 for his shopping. How much money
    did {first_name} have left?"], "ground_truth": 6},
  "11": {"sentence": [f"{first_name} has 2 dogs, 3 cats and twice as many
    fish as cats and dogs combined. How many pets does {first_name} have
    in total?"], "ground_truth": 15},
  "12": {"sentence": [f"{first_name} is buying a new pair of shoes that
    costs $95. He has been saving up his money each month for the past
    three months. He gets a $5 allowance a month. He also mows lawns and
    shovels driveways. He charges $15 to mow a lawn and $7 to shovel.
    After buying the shoes, he has $15 in change. If he mows 4 lawns,
    how many driveways did he shovel?"], "ground_truth": 5},
  "13": {"sentence": [f"{first_name} is three times older than Ben. Ben is
    two times older than Chris. If Chris is 4, how old is Caroline?"],
    "ground_truth": 24},
  "14": {"sentence": [f"{first_name} eats 1 apple a day for two weeks.
    Over the next three weeks, he eats the same number of apples as the
    total of the first two weeks. Over the next two weeks, he eats 3
    apples a day. Over these 7 weeks, how many apples does he average a
    week?"], "ground_truth": 10},
}
```

"15": {"sentence": [f"{first_name} bought 2 soft drinks for\$ 4 each and 5 candy bars. He spent a total of 28 dollars. How much did each candy bar cost?"], "ground_truth": 4,},

"16": {"sentence": [f"{first_name} has a stack of books that is 12 inches thick. He knows from experience that 80 pages is one inch thick. If he has 6 books, how many pages is each one on average?"], "ground_truth": 160,},

"17": {"sentence": [f"{first_name} is throwing a huge Christmas party. He invites 30 people. Everyone attends the party, and half of the guests bring a plus one (one other person). He plans to serve a 3-course meal for the guests. If he uses a new plate for every course, how many plates does he need in total for his guests?"], "ground_truth": 135,},

"18": {"sentence": [f"{first_name} volunteers at a shelter twice a month for 3 hours at a time. How many hours does he volunteer per year?"], "ground_truth": 72,},

"19": {"sentence": [f"{first_name} puts \$25 in his piggy bank every month for 2 years to save up for a vacation. He had to spend \$400 from his piggy bank savings last week to repair his car. How many dollars are left in his piggy bank?"], "ground_truth": 200,},

"20": {"sentence": [f"{first_name} has 16 toy cars, and the number of cars he has increases by 50% every year. How many toy cars will {first_name} have in three years?"], "ground_truth": 54,},

"21": {"sentence": [f"{first_name} just turned 12 and started playing the piano. His friend Sheila told him about the 10,000-hour rule which says, after 10,000 hours of practice, you become an expert or master in your field. If {first_name} wants to become a piano expert before he is 20, how many hours a day will he need to practice if he practices every day, Monday - Friday, and takes two weeks off for vacation each year?"], "ground_truth": 5,},

"22": {"sentence": [f"In 6 months Bella and {first_name} will be celebrating their 4th anniversary. How many months ago did they celebrate their 2nd anniversary?"], "ground_truth": 18,},

"23": {"sentence": [f"{first_name} was preparing for a party to be held in four days. So, he made 24 gallons of root beer on the first day and put them in the refrigerator cooler. But later that evening, his children discovered the delicious nectar and robbed the cooler, drinking 4 of those gallons of root beer. On the second day, his wife Barbie also discovered the root beer and accidentally spilled 7 gallons. On the third day, {first_name}'s friend Ronnie visited {first_name}'s house and helped himself to the root beer, further reducing the amount remaining by 5 gallons. On the fourth day, 3 people showed up for the party. If {first_name} and the others shared the remaining root beer equally, how much was available for each to drink during the party?"], "ground_truth": 350,},

"24": {"sentence": [f"{first_name} owns an ice cream shop and every sixth customer gets a free ice cream cone. Cones cost \$2 each. If he sold \$100 worth of cones, how many free ones did he give away?"], "ground_truth": 10,},

"25": {"sentence": [f"{first_name} likes to collect model trains. He asks for one for every birthday of his, and asks for two each Christmas. {first_name} always gets the gifts he asks for, and asks for these same gifts every year for 5 years. At the end of the 5 years, his parents give him double the number of trains he already has. How many trains does {first_name} have now?"], "ground_truth": 45,},

"26": {"sentence": [f"{first_name} has \$5000. He spends \$2800 on a new motorcycle, and then spends half of what's left on a concert ticket. {first_name} then loses a fourth of what he has left. How much money does he have left?"], "ground_truth": 825,},

"27": {"sentence": [f"{first_name} has some coins. He has 2 more quarters than nickels and 4 more dimes than quarters. If he has 6 nickels, how much money does he have?"], "ground_truth": 350,},

"28": {"sentence": [f"{first_name} is at a train station and is waiting for his train. He isn't sure how long he needs to wait, but he knows

```

    that the fourth train scheduled to arrive at the station is the one
    he needs to get on. The first train is scheduled to arrive in 10
    minutes, and this train will stay in the station for 20 minutes. The
    second train is to arrive half an hour after the first train leaves
    the station, and this second train will stay in the station for a
    quarter of the amount of time that the first train stayed in the
    station. The third train is to arrive an hour after the second train
    leaves the station, and this third train is to leave the station
    immediately after it arrives. The fourth train will arrive 20
    minutes after the third train leaves, and this is the train
    {first_name} will board. In total, how long, in minutes, will
    {first_name} wait for his train?"] , "ground_truth": 27,},
"29": {"sentence": [f"{first_name} decides to take up juggling to
perform at the school talent show a month in the future. He starts
off practicing juggling 3 balls, and slowly gets better adding 1
ball to his juggling act each week. After the end of the fourth week
the talent show begins, but when {first_name} walks on stage he
slips and drops three of his balls. 2 of them are caught by people
in the crowd as they roll off the stage, but one gets lost
completely since the auditorium is dark. With a sigh, {first_name}
starts to juggle on stage with how many balls?"] , "ground_truth":
4,},
}

```

Listing 6: Intervenability dataset for first names

```

"0": {"sentence": [f"Out of the 60 {animal}s in a kennel, 9 {animal}s
enjoy watermelon, 48 {animal}s enjoy salmon, and 5 {animal}s enjoy
both watermelon and salmon. How many {animal}s in the kennel do not
eat either watermelon or salmon?"] , "ground_truth": 8},
"1": {"sentence": [f"James needs to get more toys for his {animal}
shelter. Each {animal} needs one toy. James currently has 4 toys on
hand for 4 {animal}s, but there are 8 more {animal}s in the shelter
now. After buying the toys, he went back to see that there are twice
as many more {animal}s than when he left so he had to buy some more
toys. When James came back yet again, 3 {animal}s were gone so he no
longer needed those toys. How many toys in total does James need?"] ,
"ground_truth": 33},
"2": {"sentence": [f"If Elise purchased a 15kg bag of {animal} food and
then another 10kg bag, resulting in a total of 40kg of {animal}
food, how many kilograms of {animal} food did Elise already have
before her recent purchases?"] , "ground_truth": 15},
"3": {"sentence": [f"John takes care of 10 {animal}s. Each {animal}
takes .5 hours a day to walk and take care of their business. How
many hours a week does he spend taking care of {animal}s?"] ,
"ground_truth": 35},
"4": {"sentence": [f"Cecilia just bought a new {animal}. According to
her veterinarian, she has to feed the {animal} 1 cup of {animal}
food every day for the first 180 days. Then she has to feed the
{animal} 2 cups of {animal} food every day for the rest of its life.
If one bag of {animal} food contains 110 cups, how many bags of
{animal} food will Cecilia use in the first year?"] , "ground_truth":
5},
"5": {"sentence": [f"What is the total weight, in ounces, of the pet
food that Mrs. Anderson bought if she purchased 2 bags of 3-pound
cat food and 2 bags of {animal} food that weigh 2 pounds more than
each bag of cat food?"] , "ground_truth": 256},
"6": {"sentence": [f"If Belle consumes 4 {animal} biscuits and 2 rawhide
bones every evening, and each rawhide bone costs $1 and each
{animal} biscuit costs $0.25, what is the total cost, in dollars, to
feed Belle these treats for a week?"] , "ground_truth": 21},
"7": {"sentence": [f"If Ben's potato gun can launch a potato 6 football
fields, and each football field is 200 yards long, and Ben's
{animal} can run at a speed of 400 feet per minute, how many minutes

```

```

    will it take for his {animal} to fetch a potato that he launches?"]],
    "ground_truth": 9},
"8": {"sentence": [f"John adopts a {animal}. He takes the {animal} to
    the groomer, which costs $100. The groomer offers him a 30% discount
    for being a new customer. How much does the grooming cost?"]],
    "ground_truth": 70},
"9": {"sentence": [f"It takes 2.5 hours to groom a {animal} and 0.5
    hours to groom a cat. What is the number of minutes it takes to
    groom 5 {animal}s and 3 cats?"]], "ground_truth": 840},
"10": {"sentence": [f"A cat has nine lives. A {animal} has x less lives
    than a cat. A mouse has 7 more lives than a {animal}. A mouse has
    13 lives. What is the value of unknown variable x?"]],
    "ground_truth": 3},
"11": {"sentence": [f"If seven more {animal}s are added to the thirteen
    in the cage, the number of {animal}s in the cage will be 1/3 the
    number of {animal}s Jasper saw in the park today. How many {animal}s
    did Jasper see in the park today?"]], "ground_truth": 60},
"12": {"sentence": [f'Chris has a two-speed lawn mower. He can mow his
    entire lawn in "turtle" mode in 1 hour, or 40 minutes in "{animal}"
    mode. Today, he experimented by mowing half in turtle mode and half
    in {animal} mode. How many minutes did it take him to mow the
    lawn?'], "ground_truth": 50},
"13": {"sentence": [f"At a James Bond movie party, each guest is either
    male (M) or female (F. 40% of the guests are women, 80% of the women
    are wearing {animal} ears, and 60% of the males are wearing {animal}
    ears. If the total number of guests at the party is 200, given all
    this information, what is the total number of people wearing
    {animal} ears?"]], "ground_truth": 136},
"14": {"sentence": [f"Every month, Madeline has to buy food, treats, and
    medicine for her {animal}. Food costs $25 per week. Treats cost $20
    per month. Medicine costs $100 per month. How much money does
    Madeline spend on her {animal} per year if there are 4 weeks in a
    month?"]], "ground_truth": 2640},
"15": {"sentence": [f"Russell works at a pet store and is distributing
    straw among the rodents. The rats are kept in 3 cages in equal
    groups and each rat is given 6 pieces of straw. There are 10 cages
    of hamsters that are kept alone and each hamster is given 5 pieces
    of straw. There is also a pen of {animal} where 20 pieces of straw
    are distributed among the {animal}. No straw is used anywhere else
    in the store. If 160 pieces of straw have been distributed among the
    small rodents, how many rats are in each cage?"]], "ground_truth": 5},

```

Listing 7: Intervenability dataset for animals

```

{
"0": {"sentence": [f"If the total cost of all the {location} supplies
    that Pauline wants to purchase is $150 before sales tax, and the
    sales tax is 8% of the total amount purchased, what will be the
    total amount that Pauline will spend on all the items, including
    sales tax?"]], "ground_truth": 162},
"1": {"sentence": [f"If adding seven more rabbits to the thirteen
    already in the cage would result in a total number of rabbits that
    is 1/3 of the number of rabbits Jasper saw in the {location} today,
    what is the total number of rabbits that Jasper saw in the
    {location} today?"]], "ground_truth": 60},
"2": {"sentence": [f"If there were initially 250 books in the {location}
    and 120 books were taken out on Tuesday to be read by children,
    followed by 35 books being returned on Wednesday and another 15
    books being withdrawn on Thursday, what is the current total number
    of books in the {location}?"]], "ground_truth": 150},
"3": {"sentence": [f"Jason goes to the {location} 4 times more often
    than William goes. If William goes 2 times per week to the
    {location}, how many times does Jason go to the {location} in 4
    weeks?"]], "ground_truth": 32},

```

```
"4": {"sentence": [f"If there were initially 235 books in the {location}
and 227 books were borrowed on Tuesday, then 56 books were returned
on Thursday, and 35 books were borrowed again on Friday, how many
books are currently in the {location}?"], "ground_truth": 29},
"5": {"sentence": [f"A {location} has a number of books. 35% of them are
intended for children and 104 of them are for adults. How many books
are in the {location}?"], "ground_truth": 160},
"6": {"sentence": [f"If Benjamin spent a total of $15 at {location},
buying a salad, a burger, and two packs of fries, and one pack of
fries cost $2, with the salad being three times that price, what is
the cost of the burger?"], "ground_truth": 5},
}
```

Listing 8: **Intervenability dataset for locations**

C.10 Sweep results

Table 4: Sweep results for different insertion values for intervening on the animals concept

metric	insertion value	acc (10^{-4})	acc (10^{-5})	acc (10^{-6})	acc (0)
eval	10	0.784	0.628	0.625	0.600
eval	20	0.769	0.662	0.644	0.619
eval	50	0.716	0.650	0.637	0.619
eval	100	0.709	0.669	0.659	0.631
eval	150	0.716	0.662	0.656	0.647
eval	200	0.725	0.656	0.666	0.628
eval	300	0.731	0.656	0.650	0.637
eval	500	0.662	0.578	0.647	0.597
include	10	0.487	0.572	0.656	0.653
include	20	0.734	0.744	0.772	0.781
include	50	0.912	0.925	0.934	0.922
include	100	0.938	0.928	0.931	0.916
include	150	0.934	0.925	0.925	0.916
include	200	0.934	0.916	0.912	0.906
include	300	0.912	0.831	0.816	0.797
include	500	0.678	0.547	0.531	0.544
drop	10	0.344	0.241	0.153	0.147
drop	20	0.138	0.113	0.084	0.078
drop	50	0.028	0.009	0.000	0.003
drop	100	0.000	0.000	0.000	0.000
drop	150	0.000	0.000	0.000	0.000
drop	200	0.000	0.000	0.000	0.000
drop	300	0.000	0.000	0.000	0.000
drop	500	0.003	0.000	0.009	0.006

Table 5: Sweep results for different insertion values for intervening on the locations concept

metric	insertion value	acc (10^{-4})	acc (10^{-5})	acc (10^{-6})	acc (0)
eval	10	0.957	0.879	0.893	0.929
eval	20	0.957	0.864	0.914	0.929
eval	50	0.957	0.871	0.921	0.914
eval	100	0.957	0.871	0.907	0.907
eval	150	0.929	0.864	0.907	0.921
eval	200	0.914	0.886	0.900	0.914
eval	300	0.893	0.843	0.907	0.871
eval	500	0.814	0.779	0.893	0.821
include	10	0.229	0.214	0.321	0.421
include	20	0.400	0.429	0.464	0.429
include	50	0.521	0.586	0.607	0.579
include	100	0.557	0.586	0.600	0.586
include	150	0.557	0.586	0.600	0.593
include	200	0.564	0.593	0.593	0.600
include	300	0.571	0.600	0.586	0.593
include	500	0.564	0.593	0.586	0.657
drop	10	0.379	0.400	0.393	0.243
drop	20	0.243	0.236	0.221	0.186
drop	50	0.107	0.057	0.021	0.000
drop	100	0.029	0.000	0.000	0.000
drop	150	0.007	0.000	0.000	0.000
drop	200	0.000	0.000	0.000	0.000
drop	300	0.000	0.000	0.000	0.000
drop	500	0.014	0.007	0.000	0.000