

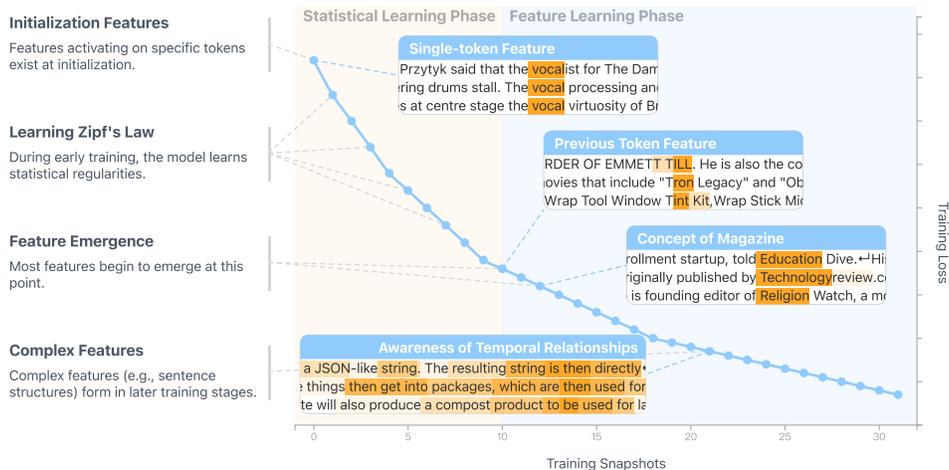
# EVOLUTION OF CONCEPTS IN LANGUAGE MODEL PRE-TRAINING

Xuyang Ge<sup>†</sup>, Wentao Shu<sup>†</sup>, Jiaxing Wu<sup>†</sup>, Yunhua Zhou<sup>‡</sup>, Zhengfu He<sup>†</sup>, Xipeng Qiu<sup>†</sup> \*

<sup>†</sup> OpenMOSS Team, Shanghai Innovation Institute; Fudan University

<sup>‡</sup> Shanghai AI Laboratory

xyge24@m.fudan.edu.cn, zfhel19@fudan.edu.cn



## ABSTRACT

Language models obtain extensive capabilities through pre-training. However, the pre-training dynamics remains a black box. In this work, we track linear interpretable feature evolution across pre-training snapshots using a sparse dictionary learning method called crosscoders. We find that most features begin to form around a specific point, while more complex patterns emerge in later training stages. Feature attribution analyses reveal causal connections between feature evolution and downstream performance. Our feature-level observations are highly consistent with previous findings on Transformer’s two-stage learning process, which we term a statistical learning phase and a feature learning phase. Our work opens up the possibility to track fine-grained representation progress during language model learning dynamics. Our code is available at <https://github.com/OpenMOSS/Language-Model-SAEs>.

## 1 INTRODUCTION

Pre-training (Radford et al., 2018; Devlin et al., 2019) has emerged as the dominant paradigm for developing frontier large language models (LLM) (OpenAI, 2023; Grattafiori et al., 2024; Yang et al., 2025). Despite its remarkable success, the pre-training process remains largely a black box. While scaling laws (Hestness et al., 2017; Kaplan et al., 2020; Bahri et al., 2021) reveal the predictable relationships between compute, data, and loss, they offer limited insight into the internal reorganization occurring within the model parameters. Other theoretical frameworks on learning dynamics, including neural tangent kernel (Jacot et al., 2018), information bottleneck theory (Tishby & Zaslavsky, 2015; Shwartz-Ziv & Tishby, 2017), and singular learning theory (Watanabe, 2009; Lau

\*Corresponding author.

et al., 2024; Wang et al., 2025a) — provide high-level explanations on *why* we can observe generalization and grokking (Power et al., 2022; Nanda et al., 2023). However, a fundamental question — *how* do LLMs actually develop their capabilities internally during pre-training — has remained largely opaque.

Recent advances in mechanistic interpretability, particularly through dictionary learning methods based on sparse autoencoders (SAEs), have begun to illuminate the internal representation of neural networks (Bricken et al., 2023; Templeton et al., 2024; Gao et al., 2025). By disentangling the phenomenon of superposition (Elhage et al., 2022), sparse autoencoders show their capabilities in extracting millions of features from LLM activations, demonstrating that LLMs encode human-interpretable concepts as linear directions in their activation spaces. However, these analyses have predominantly focused on studying fully-trained models. Consequently, the process of how features initially emerge and evolve throughout the training process remains largely unexplored.

In this paper, we propose to **track feature evolution** across pre-training snapshots using **cross-coders** (Lindsey et al., 2024), a variant of sparse autoencoders designed to simultaneously identify and align features from a family of correlated model activations. While originally introduced to resolve cross-layer superposition and track features distributed across layers, we adapt this approach to analyze activations from different training checkpoints. By applying cross-snapshot crosscoders, we can observe where features emerge, rotate, and degenerate, thereby providing deeper insight into the internal dynamics of model pre-training. Our main contributions are as follows:

1. To the best of our knowledge, this work is the first to adapt crosscoders to study training dynamics (Section 3). We evaluate completeness and faithfulness of interpretations provided by crosscoders in Appendix A.
2. We perform in-depth analyses on cross-snapshot features in Section 4. We empirically show that the decoder norms can serve as proxies of feature evolution status, and showcase both general and per-feature evolutionary properties.
3. Our method successfully connects the microscopic features to the macroscopic downstream task metrics using attribution-based circuit tracing techniques (Section 5).
4. We show evidence for the phase transition from a statistical learning phase to a feature learning phase in Section 6.

## 2 RELATED WORKS

**Learning dynamics and phase transitions.** Multiple theoretical frameworks explain neural network training dynamics. Neural Tangent Kernel (NTK) theory (Jacot et al., 2018) establishes that infinite-width networks evolve as kernel machines with fixed kernels during gradient descent, with extensions to finite-width corrections (Dyer & Gur-Ari, 2020) and modern architectures (Yang, 2020; Yang & Littwin, 2021). Recent work identifies phase transitions and lazy regimes in training dynamics (Kumar et al., 2024; Zhou et al., 2025). Information Bottleneck (IB) theory (Tishby & Zaslavsky, 2015) formulates deep learning as optimizing compression-prediction tradeoffs. Schwartz-Ziv & Tishby (2017) empirically demonstrates two distinct training phases—initial fitting followed by compression—which aligns with our findings in Section 6. Singular Learning Theory (SLT) treats neural networks as singular statistical models. Watanabe (1999; 2009) introduces the Real Log Canonical Threshold to provide geometric complexity measures predicting phase transitions, with recent advances enabling practical estimation at scale through Local Learning Coefficients (Lau et al., 2024; Furman & Lau, 2024; Wang et al., 2025a). Our work complements these theoretical frameworks by providing detailed mechanistic accounts of feature evolution during transformer pre-training, bridging high-level theory with empirical observations.

**Sparse dictionary learning.** The superposition hypothesis posits that models use linear representations (Bengio et al., 2014; Alain & Bengio, 2017; Vargas & Cotterell, 2020) to embed more features than neurons (Olah et al., 2020; Elhage et al., 2022). Sparse Autoencoders (SAEs) (Bricken et al., 2023; Huben et al., 2024; He et al., 2024) extract monosemantic features from superposition using sparse dictionary learning, with subsequent improvements and scaling efforts (Gao et al., 2025; Rajamanoharan et al., 2024). Ge et al. (2024) and Dunefsky et al. (2024) propose transcoders, an SAE variant that predicts future activations to improve circuit tracability. Recently, Lindsey et al.

(2024) introduces crosscoders to simultaneously read and write to multiple layers. Mishra-Sharma et al. (2024); Minder et al. (2025) leverages crosscoders for capturing the difference between pre-trained models and their chat-tuned versions.

### 3 TRACKING THE EVOLUTION OF FEATURES

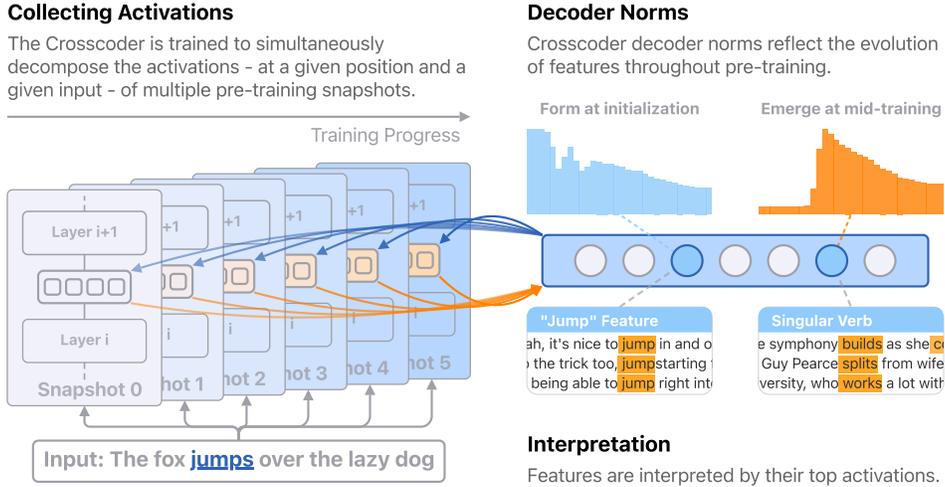


Figure 1: Overview of our method. The crosscoder is trained to decompose activations of multiple pre-training snapshots (left) into sparse features (right).

**Crosscoder architecture.** For a given text corpus  $\mathcal{C}$  and a family of training snapshots  $\Theta$  saved during LLM pre-training, let  $\mathcal{A} = \{a^\theta : \mathcal{C} \times \mathbb{N} \rightarrow \mathbb{R}^{d_{\text{model}}} \mid \theta \in \Theta\}$  denote a corresponding family of parameterized functions mapping input datapoints to model activations at a specific layer. Each datapoint  $x = (c, j) \in \mathcal{C} \times \mathbb{N}$  is a training token in  $\mathcal{C}$ , indexing the  $j$ -th token in sequence  $c$ .

A cross-snapshot crosscoder (Lindsey et al., 2024) operates on the parameterized function family  $\mathcal{A}$  over corpus  $\mathcal{C}$  (Figure 1). The crosscoder architecture is defined by:

$$f(x) = \sigma \left( \sum_{\theta \in \Theta} W_{\text{enc}}^\theta a^\theta(x) + b_{\text{enc}} \right) \quad (1)$$

$$\hat{a}^\theta(x) = W_{\text{dec}}^\theta f(x) + b_{\text{dec}}^\theta$$

with  $W_{\text{enc}}^\theta \in \mathbb{R}^{n_{\text{features}} \times d_{\text{model}}}$ ,  $b_{\text{enc}} \in \mathbb{R}^{n_{\text{features}}}$ ,  $W_{\text{dec}}^\theta \in \mathbb{R}^{d_{\text{model}} \times n_{\text{features}}}$ , and  $b_{\text{dec}}^\theta \in \mathbb{R}^{d_{\text{model}}}$ . The parameters  $W_{\text{enc}}^\theta$ ,  $W_{\text{dec}}^\theta$ , and  $b_{\text{dec}}^\theta$  correspond to snapshot-specific encoder and decoder weights for parameter  $\theta$ . The activation function  $\sigma(\cdot)$  produces sparse feature activations  $f(x)$  shared for all snapshots, and  $\hat{a}^\theta(x)$  denotes the reconstructed term of model activation  $a^\theta(x)$ .

**Training objectives.** The crosscoder is trained to minimize the loss:

$$\mathcal{L}(x) = \underbrace{\sum_{\theta \in \Theta} \|a^\theta(x) - \hat{a}^\theta(x)\|^2}_{\text{Reconstruction loss}} + \underbrace{\lambda_{\text{sparsity}} \sum_{\theta \in \Theta} \sum_{i=1}^{n_{\text{features}}} \Omega(f_i(x) \cdot \|W_{\text{dec},i}^\theta\|)}_{\text{Sparsity loss}} \quad (2)$$

where  $\lambda_{\text{sparsity}}$  is a hyperparameter to control the trade-off between reconstruction fidelity and feature sparsity. The regularization function  $\Omega(\cdot)$  serves as a differentiable substitute for L0 regularization, penalizing non-sparse feature activations. We include the decoder norm  $\|W_{\text{dec},i}^\theta\|$  in the regularization term to prevent the crosscoder from trivially reducing the feature activation  $f_i(x)$  while inflating the decoder norms under imperfect L0 approximations such as L1 regularization. Appendix A.1

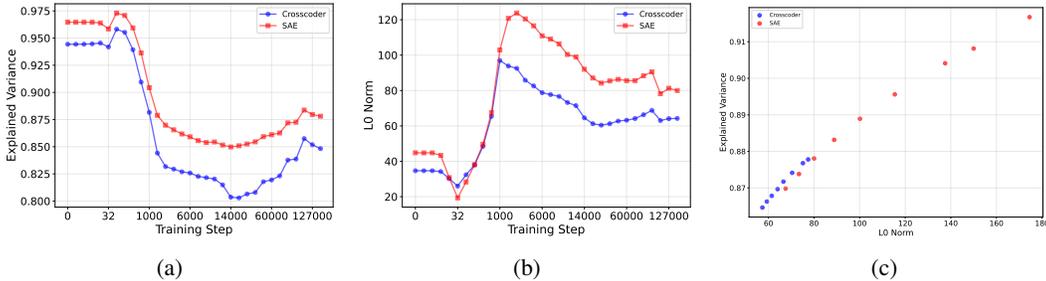


Figure 3: Comparison between crosscoders and per-snapshot SAEs. (a) The explained variance of crosscoders versus SAEs at each snapshot. (b) The L0 norm of crosscoders versus SAEs at each snapshot. (c) The Pareto frontier comparison of crosscoders and SAEs trained on the final snapshot.

shows the details for selecting the proper activation function  $\sigma(\cdot)$  and regularization function  $\Omega(\cdot)$  to optimize crosscoder feature sparsity.

**Experimental setup.** We use Pythia-160M and Pythia-6.9B (Biderman et al., 2023) for our experiments throughout this work. Pythia is a Transformer language model suite with well-controlled training settings and accessibility to training snapshots. We select the middle layers (Layer 6 of Pythia-160M and Layer 16 of Pythia-6.9B) for training crosscoders. To balance training cost with the granularity of feature evolution analysis, we select 32 snapshots out of 154 open-source snapshots ranging from step 0 to 143,000, with a stratified sampling approach: (1) all 20 snapshots before step 10,000 to capture early feature evolution with maximum temporal resolution, and (2) 12 evenly spaced snapshots from later training stages, containing 4 snapshots from steps 14,000–34,000 and 8 snapshots from steps 47,000–143,000. We use SlimPajama (Shen et al., 2023), a comprehensive text dataset covering a variety of data sources to sample activations.

**Results.** Figure 2 demonstrates crosscoder performance in decomposing model activations into sparse feature representations. We evaluate reconstruction quality (measured by activation variance explained) and sparsity (measured by L0 norm averaged across snapshots). Increasing dictionary size  $n_{\text{features}}$  yields significant Pareto improvements across both metrics.

To examine whether crosscoders can extract cross-snapshot features and align identical features in consistent directions within the feature space, we compare crosscoder performance against corresponding SAEs trained individually on each Pythia snapshot. We train SAEs using identical settings and hyperparameters as crosscoders on each snapshot. Figures 9a and 9b demonstrate that crosscoders achieve comparable performance in L0 norm and explained variance at each snapshot, even when SAEs are optimized for individual snapshots where they should have a theoretical advantage.

We further compare the Pareto frontiers (explained variance versus L0 norm) between crosscoders and SAEs trained on the final snapshot. Figure 9c shows that crosscoders exhibit a slightly superior Pareto frontier, demonstrating their effectiveness in sparse dictionary learning beyond their primary capability of tracking feature evolution across snapshots.

**Feature evolution revealed by crosscoders.** An important advantage of crosscoders is the unified feature space (or *sparse codes*) they reveal. The crosscoder encoder aggregates cross-snapshot information to produce shared feature activations. Then these activations are translated back to recover the original activations by a group of independent decoders (or *dictionaries*).

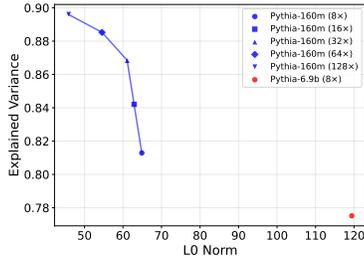


Figure 2: Explained variances versus L0 norms of our crosscoders.

If a feature activates but “exists” at only a subset of snapshots, the sparse penalty will suppress the decoder norms of this feature at irrelevant snapshots to near-zero so they won’t interfere with reconstruction on these snapshots and also reduce sparsity loss.

This design principle leads to a crucial observation: the decoder norm  $\|W_{\text{dec},i}^\theta\|$  directly reflects the strength and presence of feature  $i$  at snapshot  $\theta$ . Therefore, tracking decoder norm changes across snapshots provides a direct window into feature evolution dynamics. Appendix C provides a sanity check of whether  $\|W_{\text{dec},i}^\theta\|$  can indeed serve as a proxy of feature strength using linear probes.

**Potential failure modes of crosscoder feature alignment.** Before assessing crosscoder features, we further discuss two theoretical concerns about cross-snapshot feature alignment:

*Will crosscoders misalign unrelated features?* Such misalignment would strongly contradict the optimization objective of crosscoders. The distinct activation patterns of unrelated features would conflict when forced into the same dimension. Once a feature activates, it would cause others to activate simultaneously, introducing noise in the reconstruction. In terms of results, misalignment would lead to the polysemanticity of crosscoder features, damaging the consistency of their interpretation, which our interpretability evaluation demonstrates does not occur.

*Will crosscoders split shared features?* Suppose features from different snapshots represent the same underlying concept with identical activation patterns, splitting them across multiple dimensions would be suboptimal: it wastes representational capacity in feature space, as both previous works (Gao et al., 2025; Templeton et al., 2024) and Appendix A.4 prove that feature space dimensionality is crucial for better reconstruction fidelity. Nevertheless, in rare cases, we do observe feature splitting among highly active features. These features share semantic and directional similarity but exhibit subtly different activation patterns that justify separate dimensions. We detail this phenomenon in Appendix F.

These considerations suggest that crosscoders should naturally achieve effective feature alignment, mapping semantically equivalent features from different snapshots to consistent positions in the unified feature space.

## 4 ASSESSING CROSS-APSHOT FEATURES

### 4.1 OVERVIEW OF FEATURE EVOLUTION

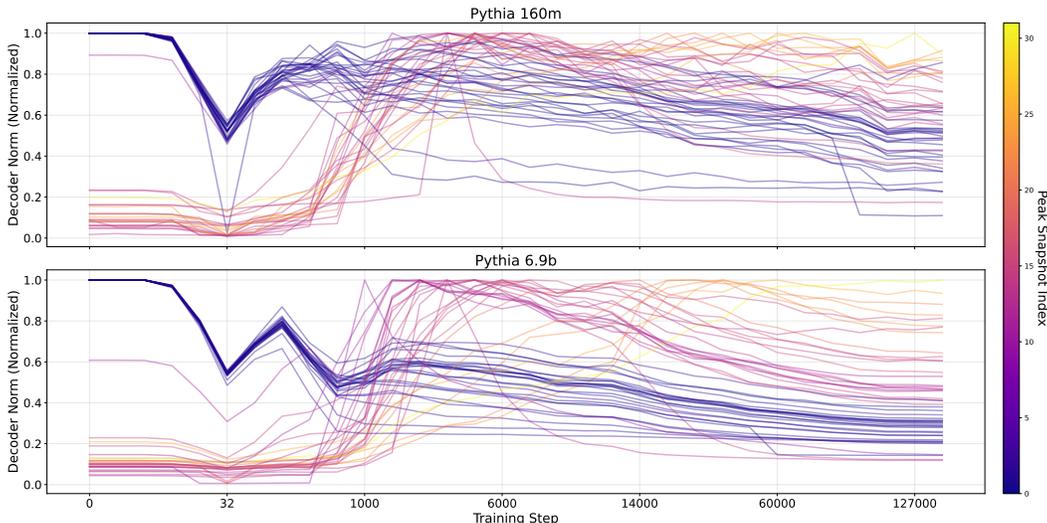


Figure 4: Overview of cross-snapshot feature decoder norm evolution. Features are extracted by a 98,304-feature crosscoder on Pythia-160M (top) and a 32,768-feature crosscoder on Pythia-6.9B (bottom).

Figure 4 shows 50 randomly sampled features and their decoder norm evolution across snapshots. Each feature’s decoder norms are linearly rescaled to a maximum of 1. Most cross-snapshot features exhibit two distinct developmental patterns:

1. *Initialization features* that exist from random initialization, exhibit a sudden drop and recovery around step 128, then gradually decay. The existence of these features has been established by Bricken et al. (2023) and Heap et al. (2025).
2. *Emergent features* that begin forming primarily around step 1000, reaching peak intensity at various subsequent training steps. There also exists emergent features only appearing in late training, which we discuss in later sections.

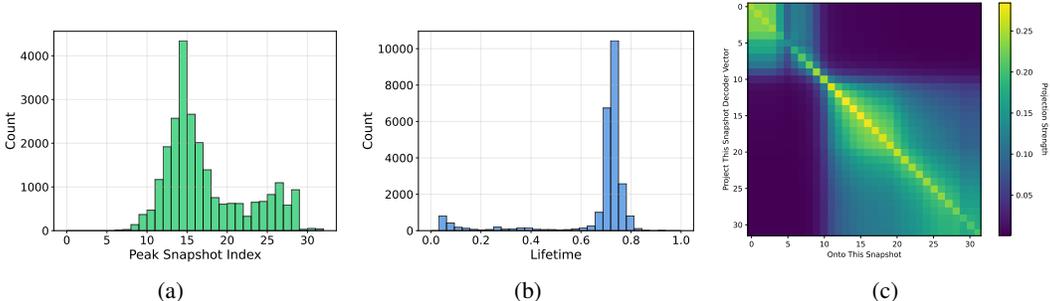


Figure 5: Statistics of emergent features in a 98,304-feature crosscoder on Pythia-160M. (a) Distribution of peak emergence times. (b) Distribution of feature lifetime. (c) Mean projection of each feature’s decoder vector of snapshot  $\theta_i$  onto its decoder vector of snapshot  $\theta_j$ .

**Feature emergence steepness.** Previous studies suggest that loss curves comprise discrete phase changes, each contributing to different circuits at distinct training stages, with evidence from toy model features (Elhage et al., 2022), 5-digit addition (Nanda et al., 2023), and in-context learning (Olsson et al., 2022). From a feature evolution perspective, we investigate whether features emerge abruptly enough to constitute the fundamental units of these phase transitions. By evaluating the time spent by each feature from initial emergence to peakness, we find that emergence steepness exhibits considerable variation, demonstrating the coexistence of both gradually-formed and abruptly-appearing features.

**Feature persistence after emergence.** We next examine whether emergent features persist after formation. We define the lifetime of a crosscoder feature as  $|\{j \mid \|W_{\text{dec},i}^{\theta_j}\| > 0.3\}|$ , where a threshold is introduced to block out near-zero decoder norms. Figure 5b shows the lifetimes of all emergent features. We find that most features persist for extended periods (above 60% of snapshots) after formation, indicating that: (1) LLMs retain learned features robustly, and (2) our crosscoders successfully align and track these features across snapshots.

**A universal directional turning point.** To further investigate the geometry of feature evolution, we compute the projections between feature directions across snapshots. The directional evolution patterns prove remarkably consistent: most features undergo drastic directional shifts around step 1,000, rendering pre- and post-step 1,000 directions nearly orthogonal (Figure 5c). Subsequently, features continue to rotate more gradually, with directions at the final snapshot maintaining notable cosine similarities to their initial post-step 1,000 orientations—a significant finding given the high-dimensional activation space.

#### 4.2 CORRELATION BETWEEN EMERGENCE STEP AND COMPLEXITY

To gain deeper insights into features emerging at different training stages, we leverage LLMs to automatically assess their complexity based on top activation patterns.

We follow Cunningham & Conerly (2024) to score feature complexity for 100 randomly sampled emergent features based on their top activations, ranging from 1 (simple) to 5 (complex). The

result in Figure 6a reveals a non-trivial correlation with moderate strength between peak timing and complexity scores, suggesting that more complex features tend to emerge later in training. The scoring rubrics and complete prompt for automated complexity scoring can be found in Appendix B.

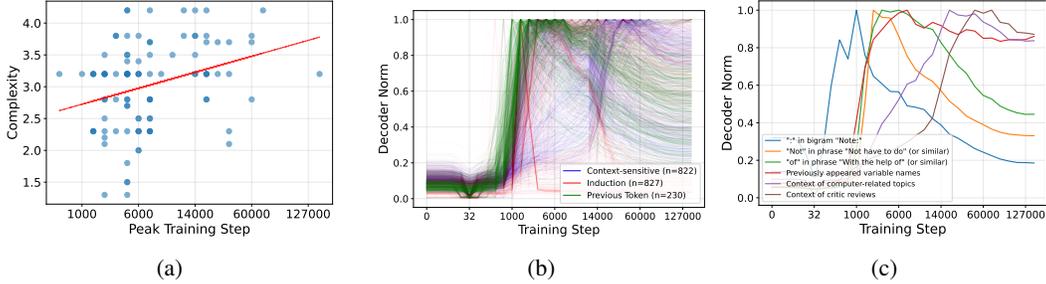


Figure 6: (a) Complexity scores (evaluated by Claude Sonnet 4) versus peak emergence times, showing a moderate positive correlation (Pearson  $r = 0.309$ ,  $p = 0.002$ ). (b) Decoder norm evolution trajectories for all previous token, induction, and context-sensitive features across training. (c) More cases of feature decoder norm evolution trajectories.

### 4.3 CASE STUDY ON TYPICAL CROSS-SNAPSHOT FEATURES

We employ simple rule-based approaches to find several well-studied feature types in a crosscoder with 32,768 features trained on Pythia-6.9B, including previous token features (which activate based on the preceding token), induction features (which activate on the second [A] in patterns [A][B]...[A][B] and help predict [B]), and context-sensitive features. These feature categories are extensively documented in prior research on neurons and SAE features (Gurnee et al., 2024; Huben et al., 2024; Ge et al., 2024).

Figure 6b demonstrates the distinct temporal pattern in feature emergence: previous token features arise early (around steps 1,000–5,000), while induction and context-sensitive features appear later and over a wider range (mostly steps 10,000–100,000). This suggests a general emergence hierarchy, from previous token to induction to context-sensitive features, which aligns with their increasing complexity. This finding is also consistent with the causal dependency between induction heads and previous token heads (Olsson et al., 2022).

For the majority of features that cannot fit into specific rule-based patterns, we further demonstrate additional cases of random emergent features with different evolutionary patterns (Figure 6c). Labels are annotated by summarizing the top activation samples.

## 5 CONNECTING MICROSCOPIC EVOLUTION TO MACROSCOPIC BEHAVIORS

One of the ambitious missions of mechanistic interpretability research is to connect feature-level findings with the model’s downstream performance. To this end, we employ attribution-based circuit tracing techniques (Syed et al., 2023; Ge et al., 2024; Marks et al., 2025) to investigate the causal effects of crosscoder features formation on downstream task metrics improvements.

**Method.** Let metric  $m : \mathbb{R}^{d_{\text{model}}} \rightarrow \mathbb{R}$  be an arbitrary scalar-valued function of model activations  $a^\theta(x)$ .<sup>1</sup> To quantify the causal effect of each feature activation  $f_i(x)$  on metric  $m$ , we first decompose model activations into per-feature representations:

$$a^\theta(x) = \hat{a}^\theta(x) + \epsilon^\theta(x) = \sum_{i=1}^{n_{\text{features}}} f_i(x) \cdot W_{\text{dec},i}^\theta + b_{\text{dec}}^\theta + \epsilon^\theta(x) \quad (3)$$

<sup>1</sup>The original notation in Marks et al. (2025) uses  $m$  as a function of input  $x$ , where the computation graph flows through nodes of interest. Since we focus on feature effects within a single layer (e.g., layer 6 of Pythia-160M), we simplify  $m$  as a function of model activations to avoid complex intervention notation.

where  $\epsilon^\theta(x) \in \mathbb{R}^{d_{\text{model}}}$  represents the crosscoder reconstruction error. This decomposition incorporates crosscoder features into the computation graph, enabling gradient computation with respect to features. We then estimate each feature’s causal effect using its attribution score:

$$\text{attr}_i^\theta(x) = f_i(x) \cdot \frac{\partial m(a^\theta(x))}{\partial f_i(x)} \quad (4)$$

where the gradient  $\frac{\partial m}{\partial f_i}$  flows through the decomposition in Eq. 3. This attribution score employs first-order Taylor expansion as a linear approximation of model computation.

For structured downstream tasks with clean/corrupted input pairs, such as subject-verb agreement (Finlayson et al., 2021), we can employ the full framework of attribution patching (Syed et al., 2023; Marks et al., 2025) by emphasizing differences between clean and corrupted inputs:

$$\text{attr}_i^\theta(x, \tilde{x}) = [f_i(x) - f_i(\tilde{x})] \cdot \frac{\partial m(a^\theta(x))}{\partial f_i(x)} \quad (5)$$

where  $\tilde{x}$  is the corrupted version of input  $x$ . Attribution patching isolates components explaining the transition from corrupted to clean performance, excluding the majority of components that contribute to baseline model function. This focused approach improves the validity of the linear approximation.

In practice, we employ the integrated gradient (IG) version of the attribution scores defined in Eq. 4 and Eq. 5, which compute gradients at evenly-spaced points from  $x$  to  $\tilde{x}$  ( $x$  to zero vector in the pure attribution case). Details of IG computation can be found in Appendix E.

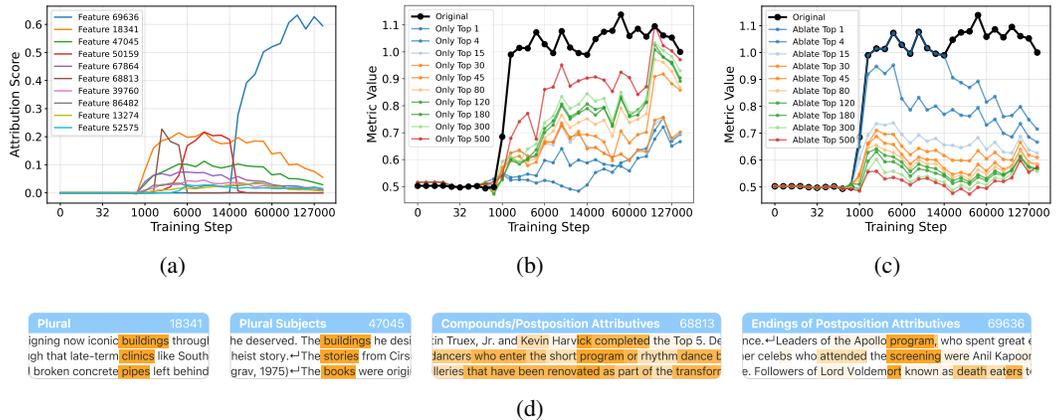


Figure 7: Crosscoder feature attribution on the Across-PP variant of the subject-verb agreement task, e.g. “The teachers near the desk are”. We use a crosscoder train at layer 6 of Pythia-160M with 98,304 features. (a) The attribution scores of top contributing features over time. (b) The metric recovery when ablating all features except the top  $k$  contributing features. (c) The metric recovery when ablating the top  $k$  contributing features. (d) Top activation samples of key features contributing to this task. Features recognizing plural nouns appear before features recognizing postposition attributives.

**Experimental setup.** We evaluate our method on subject-verb agreement (SVA) (Finlayson et al., 2021), induction, and indirect object identification (IOI) (Wang et al., 2023) tasks, using 1000 samples each to identify critical features.

For SVA and IOI, we create corrupted controls by swapping singular/plural forms (SVA) or altering the second subject (IOI), using logit differences between clean and corrupted answers as metrics. For induction, which lacks natural corruptions, we use target answer log probabilities. We compute IG attribution scores for all features on a 98,304-feature crosscoder trained on Pythia-160M, and rank crosscoder feature contribution by mean attribution scores across all snapshots. We then perform

complementary ablations: (1) removing top-ranked features, and (2) removing all except top-ranked features.

**Results.** Figure 7 shows results for the Across-PP variant of SVA, where postpositional attributives separate subjects and verbs, e.g. “The teachers near the desk are”. We identify key contributing features ordered by emergence time: (1) Features 18341 and 47045 capture plural nouns, with 47045 specialized for plural subjects; (2) Feature 68813 marks compound subjects and postpositional attributives; (3) Features 50159 and 69636 identify endings of postpositional attributives, with 69636 showing higher accuracy. Additional features include subject-specific and context-specialized plural markers with lower task contributions. Notably, Features 68813, 50159, and 69636 alternately dominate the metric, revealing circuit-level model evolution through component iteration.

Ablation experiments across all tasks demonstrate that within tens of features, we can consistently disrupt or recover model performance on downstream tasks across training snapshots, confirming our method identifies necessary and sufficient task components.

## 6 OBSERVATIONS OF A STATISTICAL-TO-SUPERPOSITION TRANSITION

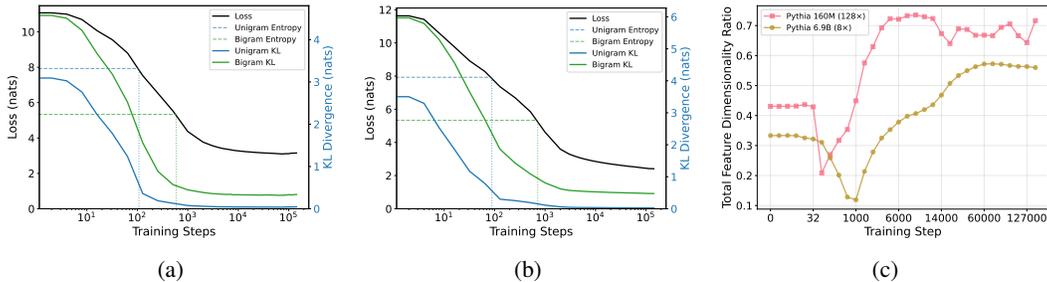


Figure 8: Observations supporting a statistical-to-superposition transition in Pythia-160M and Pythia-6.9B. (a) Unigram and bigram KL divergence evolution in Pythia-160M. The convergence timing of KL divergences coincides with when training loss approaches the theoretical minimum (unigram/bigram entropy). (b) Unigram and bigram KL divergence evolution in Pythia-6.9B. (c) Total feature dimensionality ratio (relative to activation space dimension) over training time.

What is the model learning at the beginning of training while the training loss rapidly decreases, if no features are formed at this period? We hypothesize that the rapid loss decrease in early training is driven by learning coarse statistical patterns, rather than by forming distinct features. After this initial optimization nears completion, sparse features emerge in superposition for further loss reduction.

This two-stage structure is deeply consistent with the fitting-to-compressing phase transition predicted by the information bottleneck theory (Shwartz-Ziv & Tishby, 2017).

**Early training almost exclusively learns uni- and bi-gram distributions.** Following previous work on language model learning statistical patterns (Takahashi & Tanaka-Ishii, 2017; Xu et al., 2019; Choshen et al., 2022; Belrose et al., 2024; Svete & Cotterell, 2024; Nguyen, 2024; Chen et al., 2024), we compute the KL divergence between the true token distribution  $Q$  and the model’s predicted distribution  $P$ . We randomly sample 10M tokens from SlimPajama to approximate both distributions. We then evaluate unigram KL divergence  $D_{\text{KL}}(P(x) \parallel Q(x))$  and bigram KL divergence  $D_{\text{KL}}(P(x_i | x_{i-1}) \parallel Q(x_i | x_{i-1}))$ .

The results show that both unigram and bigram KL divergences rapidly converge to low values during early training (Figure 8a and 8b). Furthermore, the training losses during this period approach the theoretical minimum achievable if the model perfectly matched the true token distributions—i.e., the entropy of these distributions. This suggests that the model primarily learns to fit statistical regularities (Zipf’s law (Zipf, 1935; Piantadosi, 2014)) during the early training stage, which explains the dense nature of internal representations at this phase.

**Total feature dimensionality undergoes compression and expansion.** To directly measure superposition status at each snapshot, we adapt the approach from Elhage et al. (2022) and compute the dimensionality of each feature at snapshot  $\theta$  as:

$$D_i = \frac{\|W_{\text{dec},i}^\theta\|^2}{\sum_{j=1}^{n_{\text{features}}} \left(\hat{W}_{\text{dec},i}^\theta \cdot W_{\text{dec},j}^\theta\right)^2} \quad (6)$$

where  $\hat{W}_{\text{dec},i}^\theta$  is the normalized version of decoder vector  $W_{\text{dec},i}^\theta$ . In contrast to its original application in toy models with ground-truth features, we apply this metric to crosscoder features due to their consistent alignment across training snapshots, providing insight into superposition dynamics.

We compute total feature dimensionalities for crosscoders trained on Pythia-160M and Pythia-6.9B (Figure 8c). Under ideal superposition, features should form symmetric arrangements with total dimensionalities summing to the activation space dimension. Feature dimensionalities should go down if large interference exists among features. We observe that total feature dimensionality first decreases then increases around the turning point, eventually reaching approximately 70% of available dimensions in the crosscoder with 98,304 features for Pythia-160M. Features from Pythia-6.9B account for a smaller proportion of activation space dimensions, likely due to the limited representational capacity of the 32,768-feature crosscoder. Nevertheless, both models exhibit the same trend, suggesting that initialization features initially form weak superposition, then become compressed to accommodate emergent features. This indicates that the model develops into a feature learning phase.

## 7 LIMITATIONS

**Scope and generalizability.** Superposition has been proven to be a general phenomenon in deep neural networks (Elhage et al., 2022). However, our analysis focuses on feature evolution in Pythia suite models using their open-source training snapshots. While previous research on feature universality (Wang et al., 2025b) suggest our method might be able to generalize to different architectures, datasets, post training dynamics, and tasks beyond language modeling, the extent to which feature evolution patterns are consistent across diverse settings remains to be established. We leave broader generalization studies to future work.

**Limited downstream task complexity.** Section 5 establishes the connection between feature evolution and model behavior. However, the downstream tasks we examine are relatively simple, constrained by multiple factors including Pythia model capabilities, sparse dictionary quality, and the current state of circuit tracing methodologies. Scaling to more complex downstream tasks represents a natural direction for future work.

**Discrete snapshot requirement.** Crosscoder training requires activations from discrete training snapshots, with memory and computational costs scaling linearly with snapshot count, which limits observational granularity. Potential solutions include architectural modifications for online multi-snapshot processing or incorporating gradient information to capture continuous training dynamics.

## 8 CONCLUSION

We introduce crosscoders to study feature evolution in LLM pre-training. Our analysis reveals two patterns: initialization-dependent features and emergent features, with complex patterns emerging later. We establish causal connections between feature evolution and downstream performance. Supported by uni- and bi-gram distribution analysis and feature dimensionality dynamics, we propose that model pre-training can be roughly divided into a statistical learning phase and a feature learning phase. This work bridges mechanistic interpretability with training dynamics.

## REPRODUCIBILITY STATEMENT

All our crosscoders are based on open-source models (Pythia suite) and datasets (SlimPajama) with public accessibility. We provide source code for (1) generating model activations on each snapshot, (2) training crosscoders, and (3) analyzing trained crosscoders at <https://github.com/OpenMOSS/Language-Model-SAEs>. Detailed instructions for replicating our results are provided in the `examples/reproduce_evolution_of_concepts/README.md` file. We note that training large crosscoders requires substantial computational resources and disk space, as also listed in the `README.md` file.

## ACKNOWLEDGEMENT

This work was supported by the National Natural Science Foundation of China (No. 62525602).

## REFERENCES

- Guillaume Alain and Yoshua Bengio. Understanding intermediate layers using linear classifier probes. In *5th International Conference on Learning Representations, ICLR 2017, Toulon, France, April 24-26, 2017, Workshop Track Proceedings*. OpenReview.net, 2017. URL <https://openreview.net/forum?id=HJ4-rAVtL>.
- Yasaman Bahri, Ethan Dyer, Jared Kaplan, Jaehoon Lee, and Utkarsh Sharma. Explaining neural scaling laws. *CoRR*, abs/2102.06701, 2021. URL <https://arxiv.org/abs/2102.06701>.
- Nora Belrose, Quintin Pope, Lucia Quirke, Alex Mallen, and Xiaoli Z. Fern. Neural networks learn statistics of increasing complexity. In *Forty-first International Conference on Machine Learning, ICML 2024, Vienna, Austria, July 21-27, 2024*. OpenReview.net, 2024. URL <https://openreview.net/forum?id=IGdpKP0N6w>.
- Yoshua Bengio, Nicholas Léonard, and Aaron C. Courville. Estimating or propagating gradients through stochastic neurons for conditional computation. *CoRR*, abs/1308.3432, 2013. URL <http://arxiv.org/abs/1308.3432>.
- Yoshua Bengio, Aaron Courville, and Pascal Vincent. Representation learning: A review and new perspectives, 2014. URL <https://arxiv.org/abs/1206.5538>.
- Stella Biderman, Hailey Schoelkopf, Quentin Gregory Anthony, Herbie Bradley, Kyle O’Brien, Eric Hallahan, Mohammad Aflah Khan, Shivanshu Purohit, USVSN Sai Prashanth, Edward Raff, Aviya Skowron, Lintang Sutawika, and Oskar van der Wal. Pythia: A suite for analyzing large language models across training and scaling. In Andreas Krause, Emma Brunskill, Kyunghyun Cho, Barbara Engelhardt, Sivan Sabato, and Jonathan Scarlett (eds.), *International Conference on Machine Learning, ICML 2023, 23-29 July 2023, Honolulu, Hawaii, USA*, volume 202 of *Proceedings of Machine Learning Research*, pp. 2397–2430. PMLR, 2023. URL <https://proceedings.mlr.press/v202/biderman23a.html>.
- Trenton Bricken, Adly Templeton, Joshua Batson, Brian Chen, Adam Jermy, Tom Conerly, Nick Turner, Cem Anil, Carson Denison, Amanda Askell, Robert Lasenby, Yifan Wu, Shauna Kravec, Nicholas Schiefer, Tim Maxwell, Nicholas Joseph, Zac Hatfield-Dodds, Alex Tamkin, Karina Nguyen, Brayden McLean, Josiah E Burke, Tristan Hume, Shan Carter, Tom Henighan, and Christopher Olah. Towards monosemanticity: Decomposing language models with dictionary learning. *Transformer Circuits Thread*, 2023. <https://transformer-circuits.pub/2023/monosemantic-features/index.html>.
- Tom B. Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel M. Ziegler, Jeffrey Wu, Clemens Winter, Christopher Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. Language models are few-shot learners. *CoRR*, abs/2005.14165, 2020. URL <https://arxiv.org/abs/2005.14165>.

- Yihong Chen, Xiangxiang Xu, Yao Lu, Pontus Stenetorp, and Luca Franceschi. Jet expansions of residual computation. *CoRR*, abs/2410.06024, 2024. doi: 10.48550/ARXIV.2410.06024. URL <https://doi.org/10.48550/arXiv.2410.06024>.
- Leshem Choshen, Guy Hacohen, Daphna Weinshall, and Omri Abend. The grammar-learning trajectories of neural language models. In Smaranda Muresan, Preslav Nakov, and Aline Villavicencio (eds.), *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers), ACL 2022, Dublin, Ireland, May 22-27, 2022*, pp. 8281–8297. Association for Computational Linguistics, 2022. doi: 10.18653/V1/2022.ACL-LONG.568. URL <https://doi.org/10.18653/v1/2022.acl-long.568>.
- Hoagy Cunningham and Tom Conerly. Circuits updates - june 2024. *Transformer Circuits Thread*, 2024. URL <https://transformer-circuits.pub/2024/june-update/index.html#hurdles>.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. BERT: pre-training of deep bidirectional transformers for language understanding. In Jill Burstein, Christy Doran, and Thamar Solorio (eds.), *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, NAACL-HLT 2019, Minneapolis, MN, USA, June 2-7, 2019, Volume 1 (Long and Short Papers)*, pp. 4171–4186. Association for Computational Linguistics, 2019. doi: 10.18653/V1/N19-1423. URL <https://doi.org/10.18653/v1/n19-1423>.
- Jacob Dunefsky, Philippe Chlenski, and Neel Nanda. Transcoders find interpretable LLM feature circuits. In Amir Globersons, Lester Mackey, Danielle Belgrave, Angela Fan, Ulrich Paquet, Jakub M. Tomczak, and Cheng Zhang (eds.), *Advances in Neural Information Processing Systems 38: Annual Conference on Neural Information Processing Systems 2024, NeurIPS 2024, Vancouver, BC, Canada, December 10 - 15, 2024*, 2024. URL [http://papers.nips.cc/paper\\_files/paper/2024/hash/2b8f4db0464cc5b6e9d5e6bea4b9f308-Abstract-Conference.html](http://papers.nips.cc/paper_files/paper/2024/hash/2b8f4db0464cc5b6e9d5e6bea4b9f308-Abstract-Conference.html).
- Ethan Dyer and Guy Gur-Ari. Asymptotics of wide networks from feynman diagrams. In *8th International Conference on Learning Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26-30, 2020*. OpenReview.net, 2020. URL <https://openreview.net/forum?id=S1gFvANKDS>.
- 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. [https://transformer-circuits.pub/2022/toy\\_model/index.html](https://transformer-circuits.pub/2022/toy_model/index.html).
- Matthew Finlayson, Aaron Mueller, Sebastian Gehrmann, Stuart Shieber, Tal Linzen, and Yonatan Belinkov. Causal analysis of syntactic agreement mechanisms in neural language models. In Chengqing Zong, Fei Xia, Wenjie Li, and Roberto Navigli (eds.), *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pp. 1828–1843, Online, August 2021. Association for Computational Linguistics. doi: 10.18653/v1/2021.acl-long.144. URL <https://aclanthology.org/2021.acl-long.144/>.
- Zach Furman and Edmund Lau. Estimating the local learning coefficient at scale, 2024. URL <https://arxiv.org/abs/2402.03698>.
- Leo Gao, Tom Dupré 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, ICLR 2025, Singapore, April 24-28, 2025*. OpenReview.net, 2025. URL <https://openreview.net/forum?id=tcsZt9ZNKD>.
- Tianyu Gao, Adam Fisch, and Danqi Chen. Making pre-trained language models better few-shot learners. In Chengqing Zong, Fei Xia, Wenjie Li, and Roberto Navigli (eds.), *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing, ACL/IJCNLP 2021, (Volume 1: Long*

*Papers*), *Virtual Event, August 1-6, 2021*, pp. 3816–3830. Association for Computational Linguistics, 2021. doi: 10.18653/V1/2021.ACL-LONG.295. URL <https://doi.org/10.18653/v1/2021.acl-long.295>.

Xuyang Ge, Fukang Zhu, Wentao Shu, Junxuan Wang, Zhengfu He, and Xipeng Qiu. Automatically identifying local and global circuits with linear computation graphs, 2024.

Aaron Grattafiori, Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Alex Vaughan, Amy Yang, Angela Fan, Anirudh Goyal, Anthony Hartshorn, Aobo Yang, Archi Mitra, Archie Sravankumar, Artem Korenev, Arthur Hinsvark, Arun Rao, Aston Zhang, Aurelien Rodriguez, Austen Gregerson, Ava Spataru, Baptiste Roziere, Bethany Biron, Binh Tang, Bobbie Chern, Charlotte Caucheteux, Chaya Nayak, Chloe Bi, Chris Marra, Chris McConnell, Christian Keller, Christophe Touret, Chunyang Wu, Corinne Wong, Cristian Canton Ferrer, Cyrus Nikolaidis, Damien Allonsius, Daniel Song, Danielle Pintz, Danny Livshits, Danny Wyatt, David Esiobu, Dhruv Choudhary, Dhruv Mahajan, Diego Garcia-Olano, Diego Perino, Dieuwke Hupkes, Egor Lakomkin, Ehab AlBadawy, Elina Lobanova, Emily Dinan, Eric Michael Smith, Filip Radenovic, Francisco Guzmán, Frank Zhang, Gabriel Synnaeve, Gabrielle Lee, Georgia Lewis Anderson, Govind Thattai, Graeme Nail, Gregoire Mialon, Guan Pang, Guillem Cucurell, Hailey Nguyen, Hannah Korevaar, Hu Xu, Hugo Touvron, Iliyan Zarov, Imanol Arrieta Ibarra, Isabel Kloumann, Ishan Misra, Ivan Evtimov, Jack Zhang, Jade Copet, Jaewon Lee, Jan Geffert, Jana Vranes, Jason Park, Jay Mahadeokar, Jeet Shah, Jelmer van der Linde, Jennifer Billock, Jenny Hong, Jenya Lee, Jeremy Fu, Jianfeng Chi, Jianyu Huang, Jiawen Liu, Jie Wang, Jiecao Yu, Joanna Bitton, Joe Spisak, Jongsoo Park, Joseph Rocca, Joshua Johnstun, Joshua Saxe, Junteng Jia, Kalyan Vasuden Alwala, Karthik Prasad, Kartikeya Upasani, Kate Plawiak, Ke Li, Kenneth Heafield, Kevin Stone, Khalid El-Arini, Krithika Iyer, Kshitiz Malik, Kuenley Chiu, Kunal Bhalla, Kushal Lakhotia, Lauren Rantala-Yearly, Laurens van der Maaten, Lawrence Chen, Liang Tan, Liz Jenkins, Louis Martin, Lovish Madaan, Lubo Malo, Lukas Blecher, Lukas Landzaat, Luke de Oliveira, Madeline Muzzi, Mahesh Pasupuleti, Mannat Singh, Manohar Paluri, Marcin Kardas, Maria Tsimpoukelli, Mathew Oldham, Mathieu Rita, Maya Pavlova, Melanie Kambadur, Mike Lewis, Min Si, Mitesh Kumar Singh, Mona Hassan, Naman Goyal, Narjes Torabi, Nikolay Bashlykov, Nikolay Bogoychev, Niladri Chatterji, Ning Zhang, Olivier Duchenne, Onur Çelebi, Patrick Alrassy, Pengchuan Zhang, Pengwei Li, Petar Vasic, Peter Weng, Prajjwal Bhargava, Pratik Dubal, Praveen Krishnan, Punit Singh Koura, Puxin Xu, Qing He, Qingxiao Dong, Ragavan Srinivasan, Raj Ganapathy, Ramon Calderer, Ricardo Silveira Cabral, Robert Stojnic, Roberta Raileanu, Rohan Maheswari, Rohit Girdhar, Rohit Patel, Romain Sauvestre, Ronnie Polidoro, Roshan Sumbaly, Ross Taylor, Ruan Silva, Rui Hou, Rui Wang, Saghar Hosseini, Sahana Chennabasappa, Sanjay Singh, Sean Bell, Seohyun Sonia Kim, Sergey Edunov, Shaoliang Nie, Sharan Narang, Sharath Rapparthi, Sheng Shen, Shengye Wan, Shruti Bhosale, Shun Zhang, Simon Vandenhende, Soumya Batra, Spencer Whitman, Sten Sootla, Stephane Collet, Suchin Gururangan, Sydney Borodinsky, Tamar Herman, Tara Fowler, Tarek Sheasha, Thomas Georgiou, Thomas Scialom, Tobias Speckbacher, Todor Mihaylov, Tong Xiao, Ujjwal Karn, Vedanuj Goswami, Vibhor Gupta, Vignesh Ramanathan, Viktor Kerkez, Vincent Gonguet, Virginie Do, Vish Vogeti, Vitor Albiero, Vladan Petrovic, Weiwei Chu, Wenhan Xiong, Wenyin Fu, Whitney Meers, Xavier Martinet, Xiaodong Wang, Xiaofang Wang, Xiaoqing Ellen Tan, Xide Xia, Xinfeng Xie, Xuchao Jia, Xuwei Wang, Yaelle Goldschlag, Yashesh Gaur, Yasmine Babaei, Yi Wen, Yiwen Song, Yuchen Zhang, Yue Li, Yuning Mao, Zacharie Delpierre Coudert, Zheng Yan, Zhengxing Chen, Zoe Papanikos, Aaditya Singh, Aayushi Srivastava, Abha Jain, Adam Kelsey, Adam Shajnfeld, Adithya Gangidi, Adolfo Victoria, Ahuva Goldstand, Ajay Menon, Ajay Sharma, Alex Boesenberg, Alexei Baevski, Allie Feinstein, Amanda Kallet, Amit Sangani, Amos Teo, Anam Yunus, Andrei Lupu, Andres Alvarado, Andrew Caples, Andrew Gu, Andrew Ho, Andrew Poulton, Andrew Ryan, Ankit Ramchandani, Annie Dong, Annie Franco, Anuj Goyal, Aparajita Saraf, Arkabandhu Chowdhury, Ashley Gabriel, Ashwin Bharambe, Assaf Eisenman, Azadeh Yazdan, Beau James, Ben Maurer, Benjamin Leonhardi, Bernie Huang, Beth Loyd, Beto De Paola, Bhargavi Paranjape, Bing Liu, Bo Wu, Boyu Ni, Braden Hancock, Bram Wasti, Brandon Spence, Brani Stojkovic, Brian Gamido, Britt Montalvo, Carl Parker, Carly Burton, Catalina Mejia, Ce Liu, Changhan Wang, Changkyu Kim, Chao Zhou, Chester Hu, Ching-Hsiang Chu, Chris Cai, Chris Tindal, Christoph Feichtenhofer, Cynthia Gao, Damon Civin, Dana Beaty, Daniel Kreymer, Daniel Li, David Adkins, David Xu, Davide Testuggine, Delia David, Devi Parikh, Diana Liskovich, Didem Foss, Dingkan Wang, Duc Le, Dustin Holland, Edward Dowling, Eissa Jamil, Elaine Montgomery, Eleonora Presani, Emily

Hahn, Emily Wood, Eric-Tuan Le, Erik Brinkman, Esteban Arcaute, Evan Dunbar, Evan Smothers, Fei Sun, Felix Kreuk, Feng Tian, Filippos Kokkinos, Firat Ozgenel, Francesco Caggioni, Frank Kanayet, Frank Seide, Gabriela Medina Florez, Gabriella Schwarz, Gada Badeer, Georgia Swee, Gil Halpern, Grant Herman, Grigory Sizov, Guangyi, Zhang, Guna Lakshminarayanan, Hakan Inan, Hamid Shojanazeri, Han Zou, Hannah Wang, Hanwen Zha, Haroun Habeeb, Harrison Rudolph, Helen Suk, Henry Aspegren, Hunter Goldman, Hongyuan Zhan, Ibrahim Damlaj, Igor Molybog, Igor Tufanov, Ilias Leontiadis, Irina-Elena Veliche, Itai Gat, Jake Weissman, James Geboski, James Kohli, Janice Lam, Japhet Asher, Jean-Baptiste Gaya, Jeff Marcus, Jeff Tang, Jennifer Chan, Jenny Zhen, Jeremy Reizenstein, Jeremy Teboul, Jessica Zhong, Jian Jin, Jingyi Yang, Joe Cummings, Jon Carvill, Jon Shepard, Jonathan McPhie, Jonathan Torres, Josh Ginsburg, Junjie Wang, Kai Wu, Kam Hou U, Karan Saxena, Kartikay Khandelwal, Katayoun Zand, Kathy Matosich, Kaushik Veeraraghavan, Kelly Michelena, Keqian Li, Kiran Jagadeesh, Kun Huang, Kunal Chawla, Kyle Huang, Lailin Chen, Lakshya Garg, Lavender A, Leandro Silva, Lee Bell, Lei Zhang, Liangpeng Guo, Licheng Yu, Liron Moshkovich, Luca Wehrstedt, Madian Khabsa, Manav Avalani, Manish Bhatt, Martynas Mankus, Matan Hasson, Matthew Lennie, Matthias Reso, Maxim Groshev, Maxim Naumov, Maya Lathi, Meghan Keneally, Miao Liu, Michael L. Seltzer, Michal Valko, Michelle Restrepo, Mihir Patel, Mik Vyatskov, Mikayel Samvelyan, Mike Clark, Mike Macey, Mike Wang, Miquel Jubert Hermoso, Mo Metanat, Mohammad Rastegari, Munish Bansal, Nandhini Santhanam, Natascha Parks, Natasha White, Navyata Bawa, Nayan Singhal, Nick Egebo, Nicolas Usunier, Nikhil Mehta, Nikolay Pavlovich Laptev, Ning Dong, Norman Cheng, Oleg Chernoguz, Olivia Hart, Omkar Salpekar, Ozlem Kalinli, Parkin Kent, Parth Parekh, Paul Saab, Pavan Balaji, Pedro Rittner, Philip Bontrager, Pierre Roux, Piotr Dollar, Polina Zvyagina, Prashant Ratanchandani, Pritish Yuvraj, Qian Liang, Rachad Alao, Rachel Rodriguez, Rafi Ayub, Raghotham Murthy, Raghu Nayani, Rahul Mitra, Rangaprabhu Parthasarathy, Raymond Li, Rebekkah Hogan, Robin Battey, Rocky Wang, Russ Howes, Ruty Rinott, Sachin Mehta, Sachin Siby, Sai Jayesh Bondu, Samyak Datta, Sara Chugh, Sara Hunt, Sargun Dhillon, Sasha Sidorov, Satadru Pan, Saurabh Mahajan, Saurabh Verma, Seiji Yamamoto, Sharadh Ramaswamy, Shaun Lindsay, Shaun Lindsay, Sheng Feng, Shenghao Lin, Shengxin Cindy Zha, Shishir Patil, Shiva Shankar, Shuqiang Zhang, Shuqiang Zhang, Sinong Wang, Sneha Agarwal, Soji Sajuyigbe, Soumith Chintala, Stephanie Max, Stephen Chen, Steve Kehoe, Steve Satterfield, Sudarshan Govindaprasad, Sumit Gupta, Summer Deng, Sungmin Cho, Sunny Virk, Suraj Subramanian, Sy Choudhury, Sydney Goldman, Tal Remez, Tamar Glaser, Tamara Best, Thilo Koehler, Thomas Robinson, Tianhe Li, Tianjun Zhang, Tim Matthews, Timothy Chou, Tzook Shaked, Varun Vontimitta, Victoria Ajayi, Victoria Montanez, Vijai Mohan, Vinay Satish Kumar, Vishal Mangla, Vlad Ionescu, Vlad Poenaru, Vlad Tiberiu Mihailescu, Vladimir Ivanov, Wei Li, Wenchen Wang, Wenwen Jiang, Wes Bouaziz, Will Constable, Xiaocheng Tang, Xiaojian Wu, Xiaolan Wang, Xilun Wu, Xinbo Gao, Yaniv Kleinman, Yanjun Chen, Ye Hu, Ye Jia, Ye Qi, Yenda Li, Yilin Zhang, Ying Zhang, Yossi Adi, Youngjin Nam, Yu, Wang, Yu Zhao, Yuchen Hao, Yundi Qian, Yunlu Li, Yuzi He, Zach Rait, Zachary DeVito, Zef Rosnbrick, Zhaoduo Wen, Zhenyu Yang, Zhiwei Zhao, and Zhiyu Ma. The llama 3 herd of models, 2024. URL <https://arxiv.org/abs/2407.21783>.

Wes Gurnee, Theo Horsley, Zifan Carl Guo, Tara Rezaei Kheirkhah, Qinyi Sun, Will Hathaway, Neel Nanda, and Dimitris Bertsimas. Universal neurons in GPT2 language models. *Trans. Mach. Learn. Res.*, 2024, 2024. URL <https://openreview.net/forum?id=ZeI104QZ8I>.

Michael Hanna, Sandro Pezzelle, and Yonatan Belinkov. Have faith in faithfulness: Going beyond circuit overlap when finding model mechanisms. In *First Conference on Language Modeling*, 2024. URL <https://openreview.net/forum?id=TZ0CCGDcuT>.

Zhengfu He, Wentao Shu, Xuyang Ge, Lingjie Chen, Junxuan Wang, Yunhua Zhou, Frances Liu, Qipeng Guo, Xuanjing Huang, Zuxuan Wu, Yu-Gang Jiang, and Xipeng Qiu. Llama scope: Extracting millions of features from llama-3.1-8b with sparse autoencoders. *CoRR*, abs/2410.20526, 2024. doi: 10.48550/ARXIV.2410.20526. URL <https://doi.org/10.48550/arXiv.2410.20526>.

Thomas Heap, Tim Lawson, Lucy Farnik, and Laurence Aitchison. Sparse autoencoders can interpret randomly initialized transformers, 2025. URL <https://arxiv.org/abs/2501.17727>.

- Joel Hestness, Sharan Narang, Newsha Ardalani, Gregory F. Diamos, Heewoo Jun, Hassan Kianinejad, Md. Mostofa Ali Patwary, Yang Yang, and Yanqi Zhou. Deep learning scaling is predictable, empirically. *CoRR*, abs/1712.00409, 2017. URL <http://arxiv.org/abs/1712.00409>.
- Robert Huben, Hoagy Cunningham, Logan Riggs Smith, Aidan Ewart, and Lee Sharkey. Sparse autoencoders find highly interpretable features in language models. In *The Twelfth International Conference on Learning Representations, ICLR 2024, Vienna, Austria, May 7-11, 2024*. OpenReview.net, 2024. URL <https://openreview.net/forum?id=F76bWRSLeK>.
- Arthur Jacot, Clément Hongler, and Franck Gabriel. Neural tangent kernel: Convergence and generalization in neural networks. In Samy Bengio, Hanna M. Wallach, Hugo Larochelle, Kristen Grauman, Nicolò Cesa-Bianchi, and Roman Garnett (eds.), *Advances in Neural Information Processing Systems 31: Annual Conference on Neural Information Processing Systems 2018, NeurIPS 2018, December 3-8, 2018, Montréal, Canada*, pp. 8580–8589, 2018. URL <https://proceedings.neurips.cc/paper/2018/hash/5a4belfa34e62bb8a6ec6b91d2462f5a-Abstract.html>.
- Jared Kaplan, Sam McCandlish, Tom Henighan, Tom B. Brown, Benjamin Chess, Rewon Child, Scott Gray, Alec Radford, Jeffrey Wu, and Dario Amodei. Scaling laws for neural language models. *CoRR*, abs/2001.08361, 2020. URL <https://arxiv.org/abs/2001.08361>.
- Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization, 2017. URL <https://arxiv.org/abs/1412.6980>.
- Tanishq Kumar, Blake Bordelon, Samuel J. Gershman, and Cengiz Pehlevan. Grokking as the transition from lazy to rich training dynamics. In *The Twelfth International Conference on Learning Representations, ICLR 2024, Vienna, Austria, May 7-11, 2024*. OpenReview.net, 2024. URL <https://openreview.net/forum?id=vt5mnLVIVo>.
- Edmund Lau, Zach Furman, George Wang, Daniel Mufet, and Susan Wei. The local learning coefficient: A singularity-aware complexity measure, 2024. URL <https://arxiv.org/abs/2308.12108>.
- Jack Lindsey, Adly Templeton, Jonathan Marcus, Thomas Conerly, Joshua Batson, and Christopher Olah. Sparse crosscoders for cross-layer features and model diffing. *Transformer Circuits Thread*, 2024. <https://transformer-circuits.pub/2024/crosscoders/index.html>.
- Samuel Marks, Can Rager, Eric J. Michaud, Yonatan Belinkov, David Bau, and Aaron Mueller. Sparse feature circuits: Discovering and editing interpretable causal graphs in language models. In *The Thirteenth International Conference on Learning Representations, ICLR 2025, Singapore, April 24-28, 2025*. OpenReview.net, 2025. URL <https://openreview.net/forum?id=I4e82CIDxv>.
- Julian Minder, Clément Dumas, Caden Juang, Bilal Chughtai, and Neel Nanda. Overcoming sparsity artifacts in crosscoders to interpret chat-tuning, 2025. URL <https://arxiv.org/abs/2504.02922>.
- Siddharth Mishra-Sharma, Trenton Bricken, Jack Lindsey, Adam Jermyn, Jonathan Marcus, Kelley Rivoire, Christopher Olah, and Thomas Henighan. Sparse crosscoders for cross-layer features and model diffing. *Transformer Circuits Thread*, 2024. <https://transformer-circuits.pub/2025/crosscoder-diffing-update/index.html>.
- Neel Nanda, Lawrence Chan, Tom Lieberum, Jess Smith, and Jacob Steinhardt. Progress measures for grokking via mechanistic interpretability, 2023. URL <https://arxiv.org/abs/2301.05217>.
- Timothy Nguyen. Understanding transformers via n-gram statistics. In Amir Globersons, Lester Mackey, Danielle Belgrave, Angela Fan, Ulrich Paquet, Jakub M. Tomczak, and Cheng Zhang (eds.), *Advances in Neural Information Processing Systems 38: Annual Conference on Neural Information Processing Systems 2024, NeurIPS 2024, Vancouver, BC, Canada, December 10 - 15, 2024*, 2024. URL [http://papers.nips.cc/paper\\_files/paper/2024/hash/b1c446eebd9a317dd0e96b16908c821a-Abstract-Conference.html](http://papers.nips.cc/paper_files/paper/2024/hash/b1c446eebd9a317dd0e96b16908c821a-Abstract-Conference.html).

- Chris Olah, Nick Cammarata, Ludwig Schubert, Gabriel Goh, Michael Petrov, and Shan Carter. Zoom in: An introduction to circuits. *Distill*, 2020. doi: 10.23915/distill.00024.001. <https://distill.pub/2020/circuits/zoom-in>.
- Team OLMo, Pete Walsh, Luca Soldaini, Dirk Groeneveld, Kyle Lo, Shane Arora, Akshita Bhagia, Yuling Gu, Shengyi Huang, Matt Jordan, Nathan Lambert, Dustin Schwenk, Oyvind Tafjord, Taira Anderson, David Atkinson, Faeze Brahman, Christopher Clark, Pradeep Dasigi, Nouha Dziri, Michal Guerquin, Hamish Ivison, Pang Wei Koh, Jiacheng Liu, Saumya Malik, William Merrill, Lester James V. Miranda, Jacob Morrison, Tyler Murray, Crystal Nam, Valentina Pyatkin, Aman Rangapur, Michael Schmitz, Sam Skjonsberg, David Wadden, Christopher Wilhelm, Michael Wilson, Luke Zettlemoyer, Ali Farhadi, Noah A. Smith, and Hannaneh Hajishirzi. 2 olmo 2 furious, 2024. URL <https://arxiv.org/abs/2501.00656>.
- 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>.
- Catherine Olsson, Nelson Elhage, Neel Nanda, Nicholas Joseph, Nova DasSarma, Tom Henighan, Ben Mann, Amanda Askell, Yuntao Bai, Anna Chen, Tom Conerly, Dawn Drain, Deep Ganguli, Zac Hatfield-Dodds, Danny Hernandez, Scott Johnston, Andy Jones, Jackson Kernion, Liane Lovitt, Kamal Ndousse, Dario Amodei, Tom Brown, Jack Clark, Jared Kaplan, Sam McCandlish, and Chris Olah. In-context learning and induction heads. *Transformer Circuits Thread*, 2022. <https://transformer-circuits.pub/2022/in-context-learning-and-induction-heads/index.html>.
- OpenAI. GPT-4 technical report, 2023. URL <https://arxiv.org/abs/2303.08774>.
- Guilherme Penedo, Hynek Kydlíček, Loubna Ben allal, Anton Lozhkov, Margaret Mitchell, Colin Raffel, Leandro Von Werra, and Thomas Wolf. The fineweb datasets: Decanting the web for the finest text data at scale, 2024. URL <https://arxiv.org/abs/2406.17557>.
- Steven T. Piantadosi. Zipf’s word frequency law in natural language: A critical review and future directions. *Psychonomic Bulletin & Review*, 21:1112 – 1130, 2014. URL <https://api.semanticscholar.org/CorpusID:14264582>.
- Alethea Power, Yuri Burda, Harri Edwards, Igor Babuschkin, and Vedant Misra. Grokking: Generalization beyond overfitting on small algorithmic datasets. *CoRR*, abs/2201.02177, 2022. URL <https://arxiv.org/abs/2201.02177>.
- Alec Radford, Karthik Narasimhan, Tim Salimans, and Ilya Sutskever. Improving language understanding by generative pre-training. *Technical Report*, 2018.
- Senthooran Rajamanoharan, Tom Lieberum, Nicolas Sonnerat, Arthur Conmy, Vikrant Varma, János Kramár, and Neel Nanda. Jumping ahead: Improving reconstruction fidelity with jumprelu sparse autoencoders. *CoRR*, abs/2407.14435, 2024. doi: 10.48550/ARXIV.2407.14435. URL <https://doi.org/10.48550/arXiv.2407.14435>.
- Peter Sanders, Kurt Mehlhorn, Martin Dietzfelbinger, and Roman Dementiev. *Sequential and Parallel Algorithms and Data Structures: The Basic Toolbox*. Springer Publishing Company, Incorporated, 1st edition, 2019. ISBN 3030252086.
- Teven Le Scao, Angela Fan, Christopher Akiki, Ellie Pavlick, Suzana Ilic, Daniel Hesslow, Roman Castagné, Alexandra Sasha Luccioni, François Yvon, Matthias Gallé, Jonathan Tow, Alexander M. Rush, Stella Biderman, Albert Webson, Pawan Sasanka Ammanamanchi, Thomas Wang, Benoît Sagot, Niklas Muennighoff, Albert Villanova del Moral, Olatunji Ruwase, Rachel Bawden, Stas Bekman, Angelina McMillan-Major, Iz Beltagy, Huu Nguyen, Lucile Saulnier, Samson Tan, Pedro Ortiz Suarez, Victor Sanh, Hugo Laurençon, Yacine Jernite, Julien Launay, Margaret Mitchell, Colin Raffel, Aaron Gokaslan, Adi Simhi, Aitor Soroa, Alham Fikri Aji, Amit Alfassy, Anna Rogers, Ariel Kreisberg Nitzav, Canwen Xu, Chenghao Mou, Chris Emezue, Christopher Klamn, Colin Leong, Daniel van Strien, David Ifeoluwa Adelani, and et al. BLOOM: A 176b-parameter open-access multilingual language model. *CoRR*, abs/2211.05100, 2022. doi: 10.48550/ARXIV.2211.05100. URL <https://doi.org/10.48550/arXiv.2211.05100>.

- Zhiqiang Shen, Tianhua Tao, Liqun Ma, Willie Neiswanger, Zhengzhong Liu, Hongyi Wang, Bowen Tan, Joel Hestness, Natalia Vassilieva, Daria Soboleva, and Eric P. Xing. Slimpajama-dc: Understanding data combinations for LLM training. *CoRR*, abs/2309.10818, 2023. doi: 10.48550/ARXIV.2309.10818. URL <https://doi.org/10.48550/arXiv.2309.10818>.
- Mohammad Shoeybi, Mostofa Patwary, Raul Puri, Patrick LeGresley, Jared Casper, and Bryan Catanzaro. Megatron-lm: Training multi-billion parameter language models using model parallelism. *CoRR*, abs/1909.08053, 2019. URL <http://arxiv.org/abs/1909.08053>.
- Ravid Shwartz-Ziv and Naftali Tishby. Opening the black box of deep neural networks via information. *CoRR*, abs/1703.00810, 2017. URL <http://arxiv.org/abs/1703.00810>.
- Lewis Smith, Sen Rajamanoharan, Arthur Conmy, Callum McDougall, Janos Kramar, Tom Lieberum, Rohin Shah, and Neel Nanda. Negative results for SAEs on downstream tasks and deprioritising SAE research (GDM mech interp team progress update #2), 2025. URL <https://www.lesswrong.com/posts/4uXCAJNuPKtKBsi28/negative-results-for-saes-on-downstream-tasks>. LessWrong.
- Mukund Sundararajan, Ankur Taly, and Qiqi Yan. Axiomatic attribution for deep networks. In Doina Precup and Yee Whye Teh (eds.), *Proceedings of the 34th International Conference on Machine Learning, ICML 2017, Sydney, NSW, Australia, 6-11 August 2017*, volume 70 of *Proceedings of Machine Learning Research*, pp. 3319–3328. PMLR, 2017. URL <http://proceedings.mlr.press/v70/sundararajan17a.html>.
- Anej Svete and Ryan Cotterell. Transformers can represent n-gram language models. In Kevin Duh, Helena Gómez-Adorno, and Steven Bethard (eds.), *Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers), NAACL 2024, Mexico City, Mexico, June 16-21, 2024*, pp. 6845–6881. Association for Computational Linguistics, 2024. doi: 10.18653/V1/2024.NAACL-LONG.381. URL <https://doi.org/10.18653/v1/2024.naacl-long.381>.
- Aaquib Syed, Can Rager, and Arthur Conmy. Attribution patching outperforms automated circuit discovery. *CoRR*, abs/2310.10348, 2023. doi: 10.48550/ARXIV.2310.10348. URL <https://doi.org/10.48550/arXiv.2310.10348>.
- Shuntaro Takahashi and Kumiko Tanaka-Ishii. Do neural nets learn statistical laws behind natural language? *CoRR*, abs/1707.04848, 2017. URL <http://arxiv.org/abs/1707.04848>.
- Adly Templeton, Tom Conerly, Jonathan Marcus, Jack Lindsey, Trenton Bricken, Brian Chen, Adam Pearce, Craig Citro, Emmanuel Ameisen, Andy Jones, Hoagy Cunningham, Nicholas L Turner, Callum McDougall, Monte MacDiarmid, C. Daniel Freeman, Theodore R. Summers, Edward Rees, Joshua Batson, Adam Jermyn, Shan Carter, Chris Olah, and Tom Henighan. Scaling monosemanticity: Extracting interpretable features from claude 3 sonnet. *Transformer Circuits Thread*, 2024. URL <https://transformer-circuits.pub/2024/scaling-monosemanticity/index.html>.
- Naftali Tishby and Noga Zaslavsky. Deep learning and the information bottleneck principle. In *2015 IEEE Information Theory Workshop, ITW 2015, Jerusalem, Israel, April 26 - May 1, 2015*, pp. 1–5. IEEE, 2015. doi: 10.1109/ITW.2015.7133169. URL <https://doi.org/10.1109/ITW.2015.7133169>.
- Oskar van der Wal, Pietro Lesci, Max Muller-Eberstein, Naomi Saphra, Hailey Schoelkopf, Willem Zuidema, and Stella Biderman. Polypythias: Stability and outliers across fifty language model pre-training runs, 2025. URL <https://arxiv.org/abs/2503.09543>.
- Francisco Vargas and Ryan Cotterell. Exploring the linear subspace hypothesis in gender bias mitigation. In Bonnie Webber, Trevor Cohn, Yulan He, and Yang Liu (eds.), *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing, EMNLP 2020, Online, November 16-20, 2020*, pp. 2902–2913. Association for Computational Linguistics, 2020. doi: 10.18653/V1/2020.EMNLP-MAIN.232. URL <https://doi.org/10.18653/v1/2020.emnlp-main.232>.

- George Wang, Jesse Hoogland, Stan van Wingerden, Zach Furman, and Daniel Murfet. Differentiation and specialization of attention heads via the refined local learning coefficient. In *The Thirteenth International Conference on Learning Representations, ICLR 2025, Singapore, April 24-28, 2025*. OpenReview.net, 2025a. URL <https://openreview.net/forum?id=SUclUOWndp>.
- Junxuan Wang, Xuyang Ge, Wentao Shu, Qiong Tang, Yunhua Zhou, Zhengfu He, and Xipeng Qiu. Towards universality: Studying mechanistic similarity across language model architectures. In *The Thirteenth International Conference on Learning Representations, 2025b*. URL <https://openreview.net/forum?id=2J18i8T0oI>.
- Kevin Ro Wang, Alexandre Variengien, Arthur Conmy, Buck Shlegeris, and Jacob Steinhardt. Interpretability in the wild: a circuit for indirect object identification in GPT-2 small. In *The Eleventh International Conference on Learning Representations, ICLR 2023, Kigali, Rwanda, May 1-5, 2023*. OpenReview.net, 2023. URL <https://openreview.net/forum?id=NpsVSN6o4u1>.
- Sumio Watanabe. Algebraic analysis for non-regular learning machines. In Sara A. Solla, Todd K. Leen, and Klaus-Robert Müller (eds.), *Advances in Neural Information Processing Systems 12, [NIPS Conference, Denver, Colorado, USA, November 29 - December 4, 1999]*, pp. 356–362. The MIT Press, 1999. URL <http://papers.nips.cc/paper/1739-algebraic-analysis-for-non-regular-learning-machines>.
- Sumio Watanabe. *Algebraic Geometry and Statistical Learning Theory*. Cambridge Monographs on Applied and Computational Mathematics. Cambridge University Press, 2009.
- Zhi-Qin John Xu, Yaoyu Zhang, and Yanyang Xiao. Training behavior of deep neural network in frequency domain. In Tom Gedeon, Kok Wai Wong, and Minho Lee (eds.), *Neural Information Processing - 26th International Conference, ICONIP 2019, Sydney, NSW, Australia, December 12-15, 2019, Proceedings, Part I*, volume 11953 of *Lecture Notes in Computer Science*, pp. 264–274. Springer, 2019. doi: 10.1007/978-3-030-36708-4\_22. URL [https://doi.org/10.1007/978-3-030-36708-4\\_22](https://doi.org/10.1007/978-3-030-36708-4_22).
- An Yang, Anfeng Li, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chang Gao, Chengen Huang, Chenxu Lv, Chujie Zheng, Dayiheng Liu, Fan Zhou, Fei Huang, Feng Hu, Hao Ge, Haoran Wei, Huan Lin, Jialong Tang, Jian Yang, Jianhong Tu, Jianwei Zhang, Jianxin Yang, Jiayi Yang, Jing Zhou, Jingren Zhou, Junyang Lin, Kai Dang, Keqin Bao, Kexin Yang, Le Yu, Lianghao Deng, Mei Li, Mingfeng Xue, Mingze Li, Pei Zhang, Peng Wang, Qin Zhu, Rui Men, Ruize Gao, Shixuan Liu, Shuang Luo, Tianhao Li, Tianyi Tang, Wenbiao Yin, Xingzhang Ren, Xinyu Wang, Xinyu Zhang, Xuancheng Ren, Yang Fan, Yang Su, Yichang Zhang, Yinger Zhang, Yu Wan, Yuqiong Liu, Zekun Wang, Zeyu Cui, Zhenru Zhang, Zhipeng Zhou, and Zihan Qiu. Qwen3 technical report, 2025. URL <https://arxiv.org/abs/2505.09388>.
- Greg Yang. Tensor programs III: neural matrix laws. *CoRR*, abs/2009.10685, 2020. URL <https://arxiv.org/abs/2009.10685>.
- Greg Yang and Etai Littwin. Tensor programs iib: Architectural universality of neural tangent kernel training dynamics. In Marina Meila and Tong Zhang (eds.), *Proceedings of the 38th International Conference on Machine Learning*, volume 139 of *Proceedings of Machine Learning Research*, pp. 11762–11772. PMLR, 18–24 Jul 2021. URL <https://proceedings.mlr.press/v139/yang21f.html>.
- Zihao Zhao, Eric Wallace, Shi Feng, Dan Klein, and Sameer Singh. Calibrate before use: Improving few-shot performance of language models. In Marina Meila and Tong Zhang (eds.), *Proceedings of the 38th International Conference on Machine Learning, ICML 2021, 18-24 July 2021, Virtual Event*, volume 139 of *Proceedings of Machine Learning Research*, pp. 12697–12706. PMLR, 2021. URL <http://proceedings.mlr.press/v139/zhao21c.html>.
- Zhanpeng Zhou, Yongyi Yang, Mahito Sugiyama, and Junchi Yan. New evidence of the two-phase learning dynamics of neural networks, 2025. URL <https://arxiv.org/abs/2505.13900>.
- George Kingsley Zipf. *The Psychobiology of Language*. Houghton-Mifflin, New York, NY, USA, 1935.

## A CROSSCODER TRAINING DETAILS

Section 3 presents the mathematical definition of our cross-snapshot crosscoder. However, in practical applications, additional architectural design and training techniques are required to advance the Pareto frontier of sparsity versus reconstruction fidelity. In this section, we detail our selection of activation function  $\sigma(\cdot)$  and regularization function  $\Omega(\cdot)$ , and present our training hyperparameters and their results. We also compare crosscoder performance to standard SAEs to evaluate how well crosscoders perform at sparse dictionary learning.

### A.1 SELECTION OF ACTIVATION FUNCTION AND REGULARIZATION FUNCTION

The sparsity of natural features is the fundamental hypothesis underlying superposition (Olshausen & Field, 1997; Elhage et al., 2022; Huben et al., 2024). To obtain crosscoder features with optimal sparsity for ideal interpretability and monosemanticity while maintaining reconstruction fidelity, we carefully select the activation function  $\sigma(\cdot)$  and the regularization function  $\Omega(\cdot)$ .

Previous SAE studies predominantly use ReLU activation with L1 regularization (Bricken et al., 2023; Huben et al., 2024). However, this configuration produces weak feature activations that are largely noise, compromising both interpretability and sparsity. We address this by adopting JumpReLU (Rajamanoharan et al., 2024) as the activation function, which eliminates activations below learned thresholds (trained via straight-through estimation (Bengio et al., 2013)).

To prevent features from becoming permanently inactive at certain snapshots, we incorporate decoder norms into the activation decision. Given pre-activation  $z(x)$ :

$$z(x) = \sum_{\theta \in \Theta} W_{\text{enc}}^{\theta} a^{\theta}(x) + b_{\text{enc}} \quad (7)$$

The  $i$ -th feature activation at snapshot  $\theta$  is defined as:

$$f_i^{\theta}(x) = z_i(x) \cdot H(z_i(x) \cdot \|W_{\text{dec},i}^{\theta}\| - t_i) \quad (8)$$

where  $H(\cdot)$  is the Heaviside step function and  $t_i \in \mathbb{R}$  is the JumpReLU threshold for feature  $i$ .

This design ensures that features with small decoder norms require stronger pre-activations to activate, preventing complete feature death while maintaining sparsity. Although this means truly inactive features retain small positive decoder norms rather than zero values, this architectural choice significantly improves crosscoder performance and enables better feature tracking across snapshots.

For regularization, we employ a combination of tanh and quadratic frequency penalties (Smith et al., 2025): the tanh component provides a superior L0 approximation by reducing penalties on strong activations, while the quadratic frequency penalty suppresses high-frequency features. This yields the following batched sparsity loss:

$$\begin{aligned} \omega_i^{\theta}(\mathcal{B}) &= \frac{1}{\mathcal{B}} \sum_{x \in \mathcal{B}} \tanh(f_i(x) \cdot \|W_{\text{dec},i}^{\theta}\|) \\ \mathcal{L}_{\text{sparsity}}(\mathcal{B}) &= \lambda_{\text{sparsity}} \sum_{\theta \in \Theta} \sum_{i=1}^{n_{\text{features}}} \omega_i^{\theta}(\mathcal{B}) \cdot \left(1 + \frac{\omega_i^{\theta}(\mathcal{B})}{\omega_0}\right) \end{aligned} \quad (9)$$

where  $\mathcal{B} \subset \mathcal{C} \times \mathbb{N}$  is an input batch, and  $\omega_i^{\theta}(\mathcal{B})$  is the single-batch differentiable estimate for the activation frequency on snapshot  $\theta$  of  $i$ -th crosscoder feature, using tanh as L0 approximation. A new hyperparameter  $\omega_0$  is introduced to quadratically penalize feature activation when  $\omega_i^{\theta}(\mathcal{B}) \gg \omega_0$ .

### A.2 SELECTION OF PRE-TRAINING SNAPSHOTS

To balance training cost with the granularity of feature evolution analysis, we train crosscoders using 32 source snapshots from the 154 open-source snapshots in the Pythia suite. Figure 10 shows the training steps of the selected snapshots.

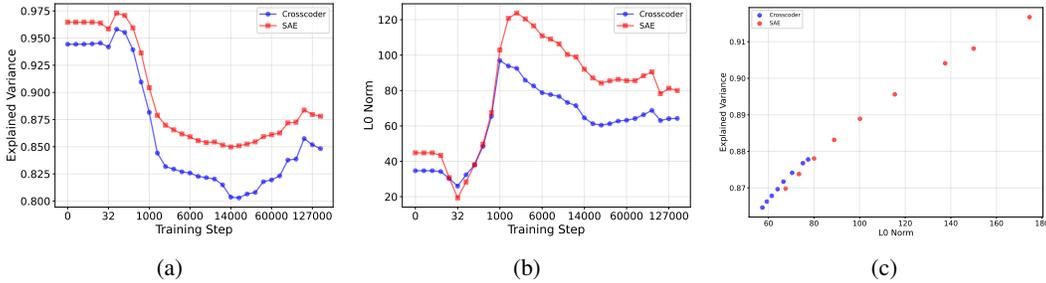


Figure 9: Comparison between crosscoders and per-snapshot SAEs. (a) The explained variance of crosscoders versus SAEs at each snapshot. (b) The L0 norm of crosscoders versus SAEs at each snapshot. (c) The Pareto frontier comparison of crosscoders and SAEs trained on the final snapshot.

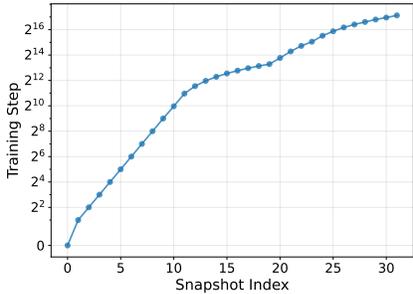


Figure 10: Selected snapshots from Pythia suite.

### A.3 DISTRIBUTED TRAINING STRATEGY FOR CROSSCODERS

Crosscoders require parameters that scale with the number of source snapshots  $n_{\text{snapshots}}$ , resulting in significantly higher memory and computational requirements compared to SAEs with equivalent feature counts, particularly when using many source snapshots.

To efficiently train crosscoders, we employ a **head parallelism** distributed training strategy, a variant of tensor parallelism (Shoeybi et al., 2019). With  $k$  processes where  $k$  divides  $n_{\text{snapshots}}$ , each process handles encoding and decoding for  $\frac{n_{\text{snapshots}}}{k}$  source snapshots. Pre-activations are computed via All-Reduce operations (Sanders et al., 2019). Unlike standard tensor parallelism, our approach processes activations from each snapshot separately on individual processes, reducing I/O overhead for reading activations from disk.

### A.4 EXPERIMENTS

We train crosscoders on Pythia-160M and Pythia-6.9B snapshots at various scales (6,144 to 98,304 features on Pythia-160M, and 32,768 features on Pythia-6.9B). We primarily focus on middle layers (Layer 6 in Pythia-160M, and Layer 16 in Pythia-6.9B)<sup>2</sup>, but also train crosscoders of 24,576 features on all layers of Pythia-160M for comprehensive analysis.

All crosscoders are trained on 800M tokens from the SlimPajama corpus using the Adam optimizer (Kingma & Ba, 2017) with  $\beta$  values of (0.9, 0.999). To prevent straight-through estimation from causing rapid threshold increases, we apply a reduced learning rate to JumpReLU threshold updates via a multiplier on the global learning rate. JumpReLU thresholds are initialized to 0.1 for all features. The learning rate schedule includes 10% warm-up steps followed by 20% decay steps. We initialize encoders as transposes of their corresponding decoders, with identical initialization matrices across all snapshots. We employ initialization search to identify optimal decoder norms that minimize loss at initialization. Additional hyperparameters are listed in Table 1.

<sup>2</sup>By “Layer  $i$ ”, we refer to the activations at the output of the  $i$ -th transformer layer.

Table 1: Hyperparameters for crosscoder training

Parameter	Pythia-160M	Pythia-6.9B
Learning Rate	5e-5	1e-5
Batch Size	2048	2048
Feature Expansion Ratio	8x, 16x, 32x, 64x, 128x	8x
Sparsity Coefficient ( $\lambda_{\text{sparsity}}$ )	0.3	0.3
JumpReLU Threshold LR Multiplier	0.1	0.3

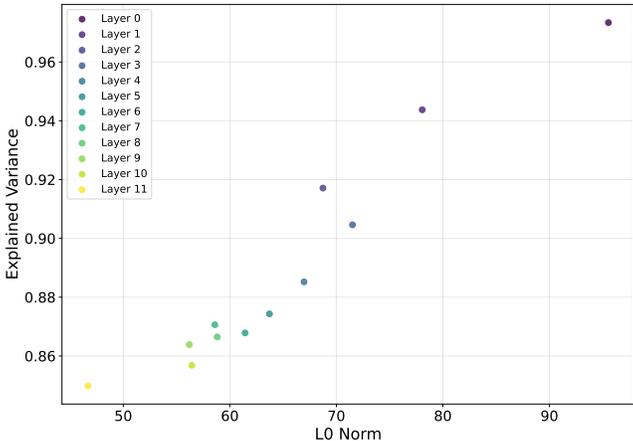


Figure 11: Explained variance versus L0 norm of crosscoders trained on all layers (0-11) of Pythia 160M

#### A.5 RESULTS ON OTHER LAYERS

We train additional crosscoders with 24,576 features on all layers (0-11) of Pythia-160M using the same hyperparameters as Layer 6 (Table 1). Figure 11 shows the explained variance and L0 norm for these crosscoders, which exhibit similar performance but different trade-offs between sparsity and reconstruction quality.

To demonstrate feature evolution across all layers, we apply the same visualization strategy from Section 4, plotting decoder norms for 50 randomly selected features per layer (Figures 12 and 13). Most middle layers exhibit the same evolutionary patterns as Layer 6, supporting our findings in Section 4. We would like to note that the majority of features in Layers 0 and 1 exist from initialization, which aligns with the observation that early layers implements more low-level, especially single-token features (He et al., 2024).

## B AUTOMATED SCORING OF COMPLEXITY

We leverage Claude Sonnet 4 for automated complexity score evaluation. For each feature, we select the top 10 activating samples with 100 surrounding tokens around the strongest activating tokens and apply the following prompt:

```
# Neural Network Feature Analysis Instructions
We're analyzing features in a neural network. Each feature activates on
specific words, substrings, or concepts within short documents. Activating
words are marked as '<<text, {activation}>>' where '{activation}' indicates
the strength of activation (higher values = stronger activation). You'll
receive documents containing highest activations tokens and tokens
surrounding them.
```

```
## Your task:
### 1. Summarize the Activation (<20 words)
```

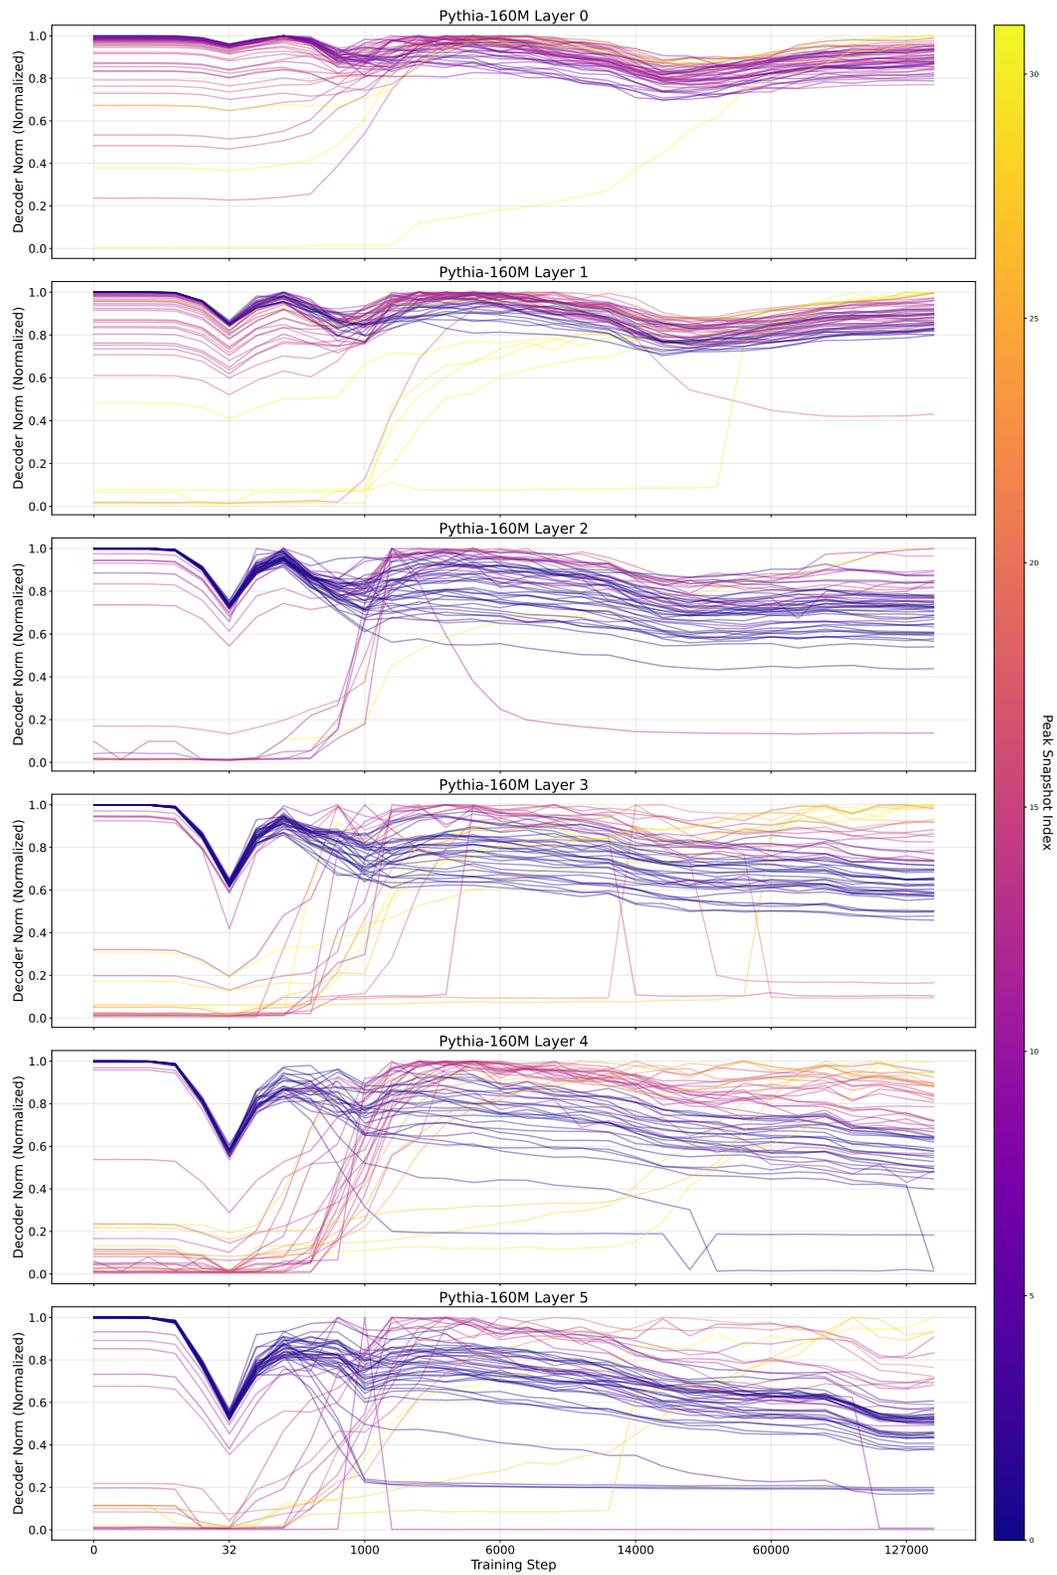


Figure 12: Feature decoder norm evolution of Layer 0 to Layer 5 in a 24,576-feature crosscoder trained on Pythia-160M.

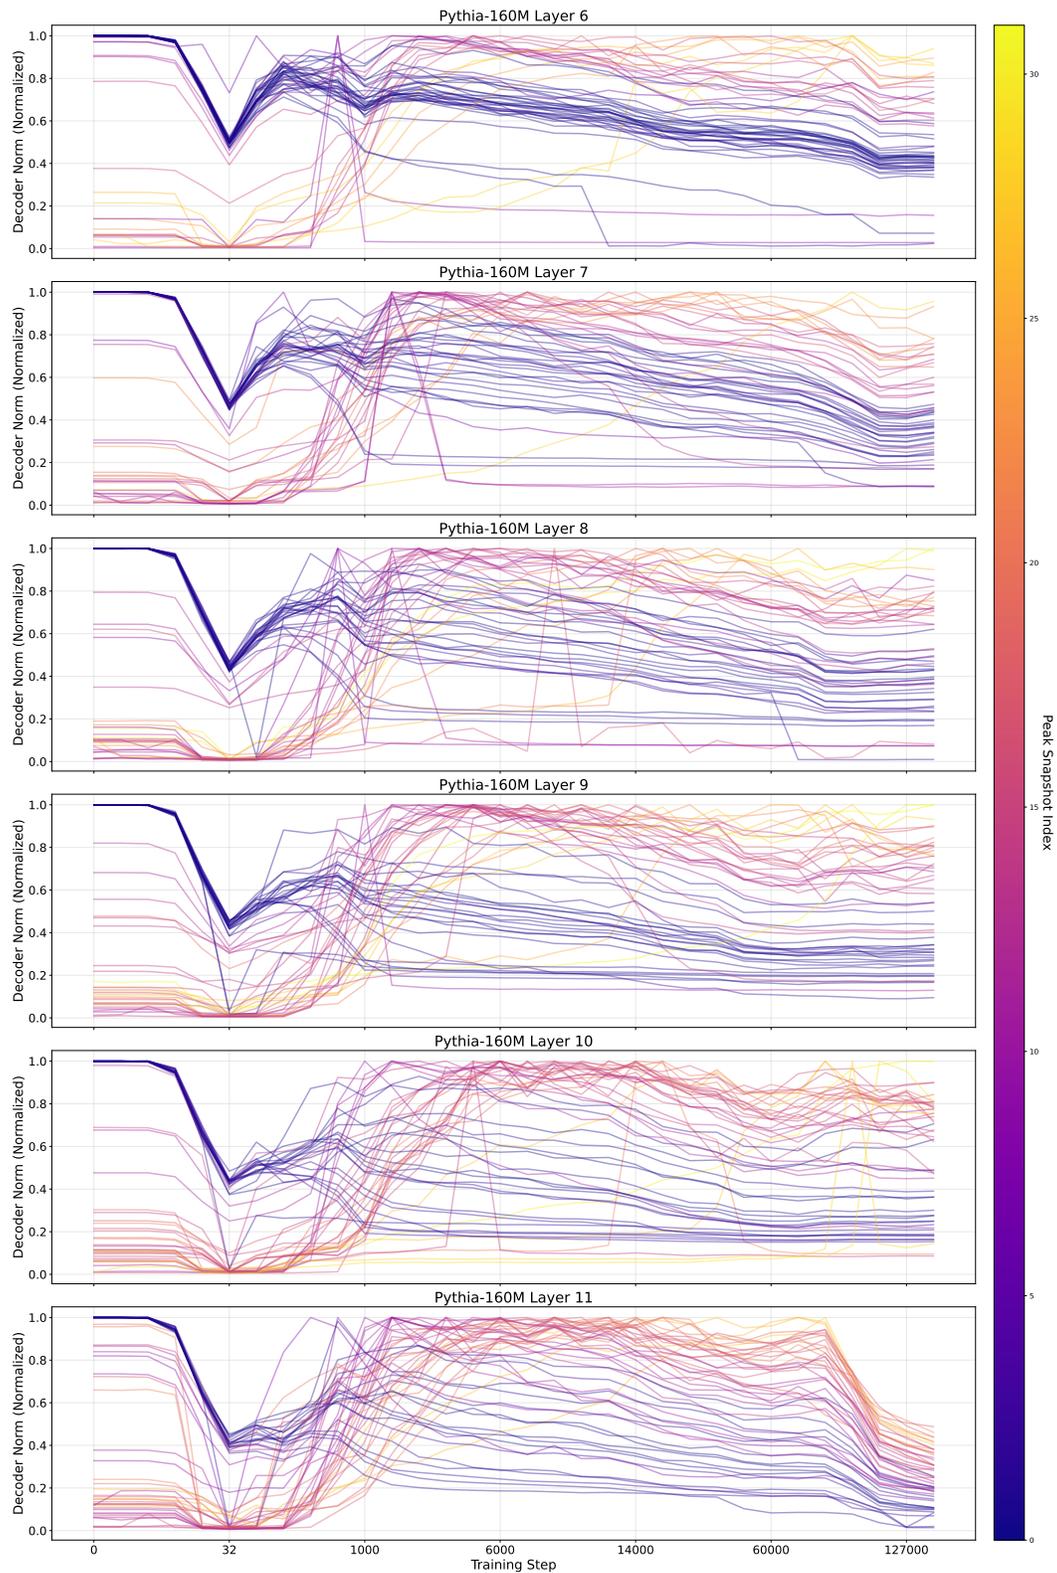


Figure 13: Feature decoder norm evolution of Layer 6 to Layer 11 in a 24,576-feature crosscoder trained on Pythia-160M.

Examine the marked activations and summarize what the feature detects in one sentence.

- Avoid being overly specific|your explanation should cover most/all activating words
- If all words in a sentence activate, focus on the sentence’s concept rather than individual words
- Note relevant patterns in capitalization, punctuation, or formatting
- Prioritize strongly activated tokens
- Keep explanations simple and concise
- Avoid long word lists

### 2. Assess Feature Complexity (1-5 scale, with decimal precision allowed)  
Rate the feature’s complexity:

- **\*\*5\*\***: Rich feature with diverse contexts unified by an interesting theme
- **\*\*4\*\***: High-level semantic structure with potentially dense activation
- **\*\*3\*\***: Moderate complexity|phrases, categories, or sentence structures
- **\*\*2\*\***: Synonyms or words at a same class
- **\*\*1\*\***: Single specific word or token

You may use decimal values (e.g., 3.7) for more precise assessment.

### Output Format

Your output should be in JSON format, with two fields: summarization and complexity. You should directly output the JSON object, without any other text.

## C ASSESSING CROSSCODER DECODER NORM

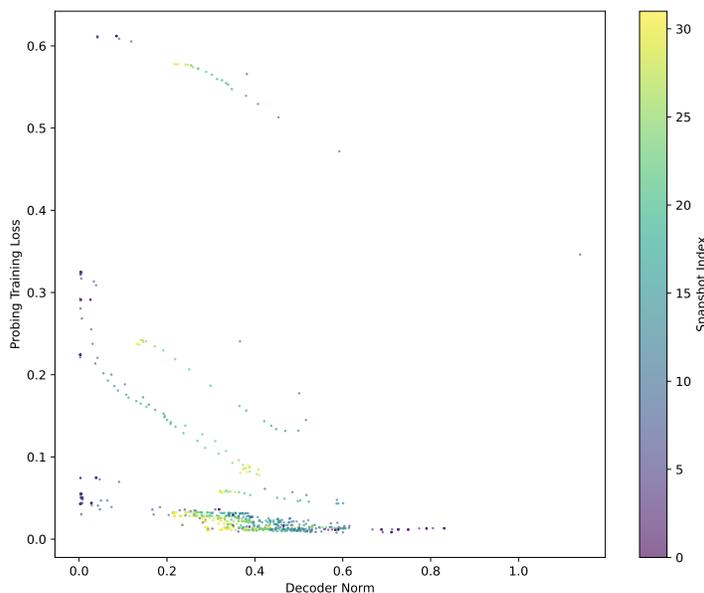


Figure 14: Linear probe training errors versus feature decoder norms of 20 randomly sampled features from a crosscoder with 6,144 features on Pythia-160M.

Section 4 takes advantage of the crosscoder decoder norms to study feature evolution. However, are the crosscoder decoder norms well quantitative indicators of the extent to which features evolved? Aside from theoretical statements, we conduct experiments to examine whether crosscoder decoder norms can reflect the intensity of features. We leverage linear probing to test the correlation between the feature decoder norm and the linear separability of the feature activations on each crosscoder feature.

We train linear probing classifiers (referred to as probes) separately on activations of each model snapshot to classify whether each crosscoder feature activates. For each feature  $i$  and each snapshot  $\theta$ , our probes map model activations to the predicted feature activating probability  $p_i^\theta(x)$  as:

$$p_i^\theta(x) = \text{sigmoid}(w_{\text{probe},i}^\theta \cdot a^\theta(x) + b_{\text{probe},i}^\theta) \quad (10)$$

where  $w_{\text{probe},i}^\theta \in \mathbb{R}^{d_{\text{model}}}$  and  $b_{\text{probe},i}^\theta \in \mathbb{R}$  is the weight and bias of the probe w.r.t. feature  $i$  and snapshot  $\theta$ . Each probe is trained to minimize the binary cross-entropy loss  $y_i \cdot \log p_i^\theta(x) + (1 - y_i) \cdot (\log(1 - p_i^\theta(x)))$ , given the label  $y_i = \text{sgn}(f_i(x))$ . The training loss of each probe should be a direct measure of the linear separability of the feature activations.

We train probes for 100M tokens for each feature in a 6,144-feature crosscoder on Pythia-160M. The probe errors of each feature show a mean Pearson correlation with the corresponding decoder norms by  $-0.867$ , with a standard deviation of  $0.153$ . We demonstrate example probe errors versus feature decoder norms of 20 randomly sampled features in Figure 14. This indicates a strong negative linear relationship between probe errors and crosscoder decoder norms, demonstrating the effectiveness of the crosscoder decoder norms as indicators of feature evolution.

## D RULES FOR FINDING TYPICAL CROSS-SNAPSHOT FEATURES

We define the rules used to identify previous token features, induction features, and context-sensitive features as follows:

1. **Previous Token Features:** We collect the directly preceding tokens of all activating tokens in the top 20 activating samples and assess their consistency. Token consistency is defined as the proportion of the largest single group when all tokens are normalized by stemming (removing leading and trailing spaces and ignoring case). To exclude bigram or multigram features, we also evaluate the consistency of the activating tokens themselves. A feature is classified as a previous token feature if it exhibits high consistency in previous tokens (above 0.8) and low consistency in activating tokens (below 0.3).
2. **Induction Features:** Induction features should activate on the second [A] in patterns [A][B]...[A][B]. For each activating token [A], we collect its following token [B] and search for previous occurrences of the bigram [A][B]. To distinguish induction features from simpler features that merely activate on any bigram [A][B], we require that the feature does not activate on the first appearance of [A][B]. A feature exhibiting this behavior in at least 20 instances within the top 20 activating samples is classified as an induction feature.
3. **Context-sensitive Features:** We identify context-sensitive features using a simpler rule based on activation density within specific contexts. Context-sensitive features should activate frequently in highly specific contexts, so we require features to have high activation counts within the top 20 activating samples (exceeding 4,000 activations). To exclude features that activate ubiquitously (such as positional or bias features), we filter out features with excessive total activations (below 2M total activations across 100M analyzed tokens).

## E DETAILS IN DOWNSTREAM TASK ATTRIBUTION

### E.1 FORMALIZATION OF IG ATTRIBUTION SCORE

To more accurately estimate the causal effect of features, we employ the integrated gradient (IG) version of the attribution score (Sundararajan et al., 2017; Hanna et al., 2024; Marks et al., 2025). IG attribution computes gradients along an interpolation path between baseline and target activations, providing more robust estimates than single-point gradients.

For the standard attribution score without clean/corrupted input pairs, we compute the IG version as:

$$\text{attr}_{\text{ig},i}^\theta(x) = \frac{1}{N} \sum_{\alpha} f_i(x) \cdot \frac{\partial m(a^\theta(x))}{\partial(\alpha f_i(x))} \quad (11)$$

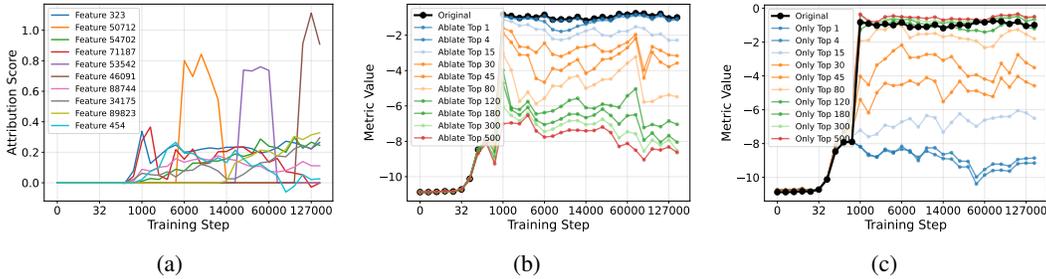


Figure 15: Crosscoder feature attribution on the induction task, using a crosscoder with 98,304 features trained on Pythia-160M.

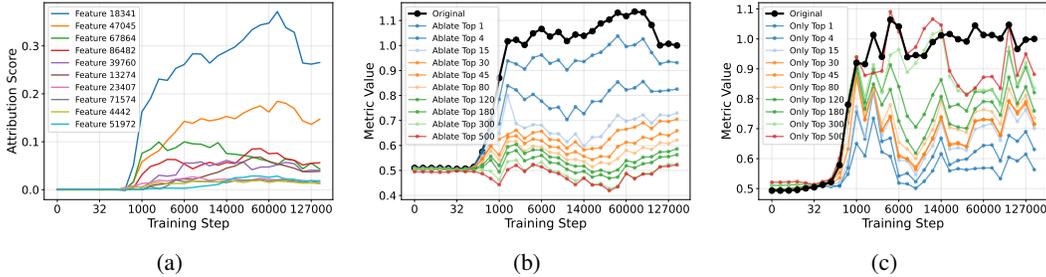


Figure 16: Crosscoder feature attribution on the Simple variant of the SVA task, using a crosscoder with 98,304 features trained on Pythia-160M.

For attribution patching with clean/corrupted input pairs, the IG version is:

$$\text{attr}_{\text{ig},i}^{\theta}(x, \tilde{x}) = \frac{1}{N} \sum_{\alpha} [f_i(x) - f_i(\tilde{x})] \cdot \frac{\partial m(a^{\theta}(x))}{\partial (\alpha f_i(x) + (1 - \alpha) f_i(\tilde{x}))} \quad (12)$$

where  $\alpha \in \{0, \frac{1}{N}, \dots, \frac{N-1}{N}\}$  linearly interpolates between the baseline and target feature activations. Consistent with Marks et al. (2025), we use  $N = 10$  interpolation steps for the IG attribution score.

## E.2 INDUCTION TASK

Transformers exhibit in-context learning capabilities (Brown et al., 2020; Zhao et al., 2021; Gao et al., 2021) through induction heads (Olsson et al., 2022)—circuits that look back over the sequence for previous instances of the current token (A), identify the subsequent token (B), and predict the same completion will occur again (forming sequences [A][B] ... [A] → [B]).

To evaluate models’ induction abilities and trace them at the feature level, we construct samples with random tokens where identical patterns appear in the middle and at the end of sequences. We test whether models can correctly copy previous text as the next token. For precise feature-level analysis, we restrict next tokens to single capital letters with leading spaces. Since the induction task lacks natural corrupted counterparts, we use the log probability of the correct token as our evaluation metric. Results are shown in Figure 15.

## E.3 SUBJECT-VERB AGREEMENT TASKS

Subject-verb agreement tasks evaluate whether models can predict verbs in the appropriate grammatical form based on their subjects. We test four variants: (1) Simple—the verb directly follows the subject, e.g., “The parents are”; (2) Across-RC—a relative clause intervenes between subject and verb, e.g., “The athlete that the managers like does”; (3) Within-RC—both subject and verb appear within the relative clause, e.g., “The athlete that the managers like”; (4) Across-PP—a prepositional

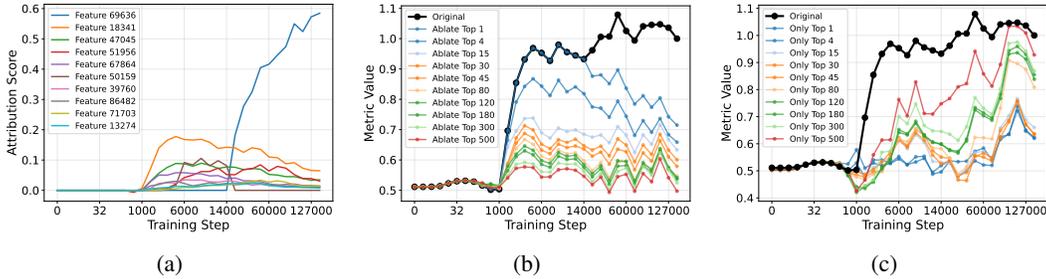


Figure 17: Crosscoder feature attribution on the Across-RC variant of the SVA task, using a cross-coder with 98,304 features trained on Pythia-160M.

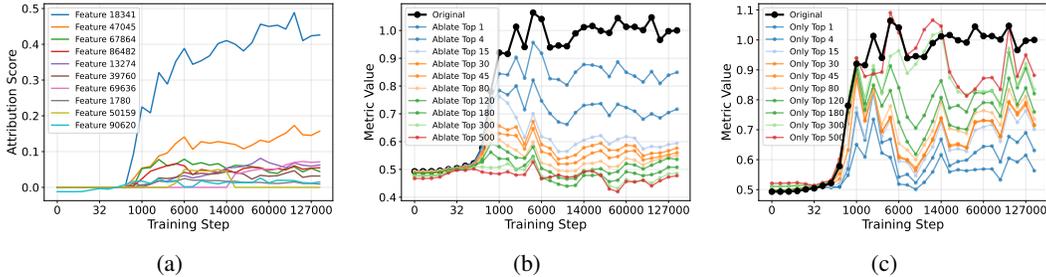


Figure 18: Crosscoder feature attribution on the Within-RC variant of the SVA task, using a cross-coder with 98,304 features trained on Pythia-160M.

phrase separates the subject and verb, e.g., “The secretaries near the cars have”. We use data provided by Marks et al. (2025), sampling 1000 examples from each variant. We use the verb in the wrong form as the counterpart in attribution patching. Results are shown in Figure 7, 16, 17, and 18.

#### E.4 INDIRECT OBJECT IDENTIFICATION TASK

Indirect Object Identification (IOI) evaluates whether models can correctly predict the indirect object (IO) based on repeated occurrence of subjects (Wang et al., 2023). In IOI tasks, sentences such as “When Mary and John went to the store, John gave a drink to” should be completed with “Mary.” We generate IOI samples following the same strategy as Wang et al. (2023), using template sentences with random person names. We use the subject name as the corrupted counterpart in attribution patching. Results are shown in Figure 19.

Note that while ablating top features significantly degrades performance, we cannot recover the original metric using only these top features. This likely occurs because IOI is a complex task requiring multiple feature interactions. The features that distinguish between indirect objects and subjects represent only part of the full computational requirements, making isolated feature sets insufficient for complete task execution.

#### E.5 RANDOM BASELINE OF ABLATION

In Figure 20, we add baselines to Figure 7 by ablating/preserving  $k$  random features. The results show that our feature selected by attribution patching is far more effective.

### F CROSSCODER FEATURE SPLITTING

Feature splitting is a well-known phenomenon in SAEs where increasing the dictionary size  $n_{\text{features}}$  causes features from smaller SAEs to fragment into multiple distinct features. Under feature splitting, a single concept may not be represented by one feature, but rather by multiple specialized features that activate on the same concept in different contexts.

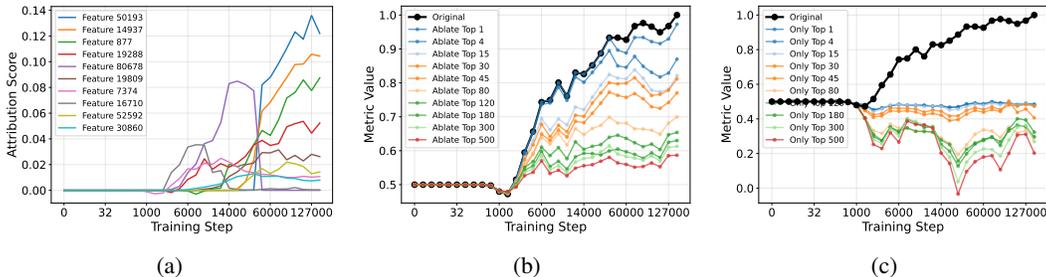


Figure 19: Crosscoder feature attribution on the IOI task, using a crosscoder with 98,304 features trained on Pythia-160M.

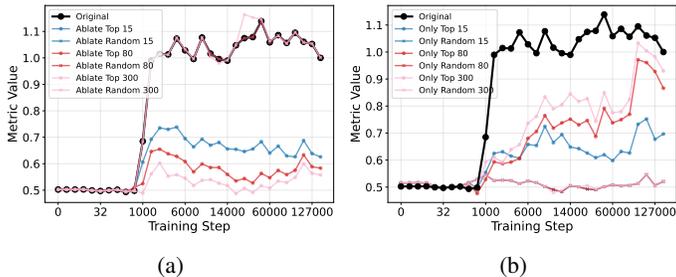


Figure 20: Crosscoder feature attribution on the Across-PP variant of the SVA task, with random baselines.

We observe similar feature splitting phenomenon in crosscoders, but across the temporal dimension. We find in rare cases, features with similar semanticity and direction in different snapshots may not be encoded in the same latent of crosscoder feature space, but result in separate features.

For example, we identify feature 66688, 53542, and 42307 that all activate on long sequences containing repeated text patterns, contributing to the induction task (Figures 21, 22, and 23). Each feature exists only during non-overlapping continuous periods across snapshots, with a drastic emergence and disappearance.

To examine whether feature splitting across snapshots relates to dictionary size, we search for feature decoder vectors (across all snapshots) with cosine similarity above 0.7—a high threshold in such high-dimensional space—in crosscoders with varying  $n_{\text{features}}$ . Across dictionary sizes ranging from 6,144 to 98,304 features, we observe that while each snapshot after step 2,000 activates almost exactly one feature representing this concept, the total number of distinct crosscoder features increases with dictionary size (Figure 24). This confirms that larger dictionaries lead to temporal feature splitting, where a single underlying concept splits into multiple features active at different training stages.

We also observe that temporal feature splitting predominantly occurs among densely activating features, i.e., features that activate frequently across many contexts. This crosscoder feature splitting likely arises from subtly different activation patterns that emerge over training. These findings suggest that language model features may evolve by refining their activation patterns, leading to more specialized representations that warrant separate feature assignments at different training stages.

## G GENERALIZABILITY STUDIES

In this section, we include more experiments on training and analyzing crosscoders on different base models/initialization/datasets and different numbers of snapshot selection, to examine the generalizability of our methods.

```

method is not threadsafe. An endpoint should be configured when the client is created and before any
service requests are made. Changing it afterwards creates inevitable race conditions for any
service requests in transit or retrying.
@param endpoint The endpoint (ex: "cognito-idp.us-east-1.amazonaws.com") or a full URL, including the protocol (ex: "https://cognito-idp.us-east-1.amazonaws.com") of the region specific AWS endpoint this client will
communicate with.
@deprecated use {@link AwsClientBuilder#setEndpointConfiguration(AwsClientBuilder.EndpointConfiguration)} for example:
{@code builder.setEndpointConfiguration(new EndpointConfiguration(endpoint, signingRegion));}
@Deprecated void setEndpoint(String endpoint);
/**
 * An alternative to {@link AWSCognitoIdentityProvider#setEndpoint(String)}, sets the regional endpoint for this client's service calls. Callers can use this method to control which AWS region they want to work with.
<p>
 * By default, all service endpoints in all regions use the https protocol. To use http instead, specify it in the {@link ClientConfiguration} supplied at construction.
<p>
 * <b>This method is not threadsafe. A region should be configured when the client is created and before any service requests are made. Changing it afterwards creates inevitable race conditions for any service requests in transit or retrying.</b>
 * @param region The region this client will communicate with. See {@link Region#getRegion(com.amazonaws.regions.Regions)}

```

Figure 21: Top activation of Feature 66688 in the 98,304-feature crosscoder on Pythia-160M.

Angry Birds Space Mod APK 2022 v2.2.15 (latest Version, Unlimited Money) This game is the ultimate in family fun. It's easy to learn, it's challenging enough for adults and kids alike, and it has hours of replay value. The best part? You can play Angry Birds Space Mod APK on your phone or tablet with no need for a computer download! So if you love this game as much as we do and who doesn't? -don't wait another moment before downloading Angry Birds Space Mod APK today. Rovio Entertainment Corporation Updated 2 days ago Size Android 4.1 Get it on 1 Angry Birds Space Mod APK Unlimited Everything 2 Amazing Features of Angry Birds Space Mod APK 2.1 Red Bird 2.2 Yellow Bird 2.3 Blue Bird 2.4 White Bird 2.5 Black Bird 2.6 Boomerang Bird 2.7 Pig Soldier 2.8 Bird Sling 2.9 Slow Motion 2.10 Hitting Obstacles with A Slingshot 2.11 Obstacles in Space 3 FAQs of Angry Birds Space Mod APK 3.1 Does Angry Birds Space Mod APK Works Offline? 3.2 Is It Ads-Free? 3.3 Is It Easy to Play? 3.4 Can I play offline without an internet connection on my laptop, PC, Mac, or smartphone? 3.5 Do I need to root my Android device? Angry Birds Space Mod APK Unlimited Everything Angry Birds Space Mod APK is the ultimate in family fun. It's easy to learn, it's challenging enough for adults and kids alike, and it has hours of replay value. The best part? You can play Angry Birds Space Mod APK on your phone or tablet with no need for a computer download. So if you love this game as much as we do and who doesn't? don't wait another moment before downloading Angry Birds Space Mod APK today. Download Angry Birds Rio Mod APK Download Adventure

Figure 22: Top activation of Feature 53542 in the 98,304-feature crosscoder on Pythia-160M.

## G.1 GENERALIZATION ACROSS MODEL FAMILIES/SEEDS

We conduct experiments on two additional Pythia 160M models with different random seeds (Pythia 160M Seed1 and Seed2) (van der Wal et al., 2025) and the Alias run of Stanford CRFM’s GPT-2 model<sup>3</sup>. All of these models have an activation space of 768 dimensions. We train crosscoders of 24,576 features (32x) on Layer 6 of them, and inspect their feature evolution.

**Different initializations (Pythia 160M variants).** We observe highly consistent results across differently initialized versions of Pythia 160M. Decoder norms of initialization features follow identical trajectories, and emergent features begin rising around step 1000, exhibiting a clear two-phase pattern with the same transition point (Figure 26 and 27). Attribution experiments confirm that we successfully identify sparse crosscoder features with significant contributions to downstream tasks (Figure 29 and 30).

**Different base model (GPT-2).** For Stanford CRFM’s GPT-2 model, we observe a clear two-phase pattern, with emergent features rising around step 1,000. Upon manual inspection of these features, we still find that initialization features exhibit token-level information, while more complex patterns emerge in emergent features.

However, we find a notable difference: due to absolute positional embeddings, initialization feature norms do not peak at the beginning but rather at later steps (with above-threshold norm at beginning). This difference arises because the absolute positional embeddings initially exhibit large norms (larger than the word embeddings) and dominate the activation norm. Since we normalize all activation sources to the same average activation norm  $\sqrt{d_{\text{model}}}$ , this large positional contribution substantially reduces the observed norms of initialization features. (Figure 28).

To sum up, the statistical-to-feature-learning phase transition still exists across models, despite small differences. Attribution experiments (Figure 31) again confirm the effectiveness of sparse crosscoder features, with generally similar feature evolution patterns for specific downstream tasks.

**Other models with accessible checkpoints.** We also investigated other model families, including BLOOM (Scao et al., 2022) and OLMo (OLMo et al., 2024). Unfortunately, these models either lack sufficient intermediate checkpoints for fine-grained analysis of pretraining dynamics (BLOOM), or lack sufficiently early checkpoints to capture the two-stage transition (BLOOM and OLMo).

- <strong><a data-lity="message-video" href="https://www.youtube.com/watch?v=Q-C4qqsgs8w" rel="noopener noreferrer nofollow">Power Rangers (2017 Movie) Official Teaser Trailer - 'Discover The Power' </a></strong> [YouTube] <br> <br><em>SABAN'S POWER RANGERS follows five ordinary teens who must become something extraordinary when they learn that their small town of Angel Grove — and the world — is on the verge of being obliterated by an alien threat. Chosen by destiny, our heroes quickly discover they are the only ones who can save the planet. But to do so, they will have to overcome their real-life issues and before it's too late, band together as the Power Rangers.</em> <br> <br> <em>Pretty typical post-Nolan franchise reboot on Lionsgate money. Doesn't really tap into Power Rangers in any enticing way. Not very impressive. Future trailers might excite more.</em> <em>Power Rangers (2017 Movie) Official Teaser Trailer - 'Discover The Power' [YouTube]</em><sup>3</sup>SABAN'S POWER RANGERS follows five ordinary teens who must become something extraordinary when they learn that their small town of Angel Grove — and the world — is on the verge of being obliterated by an alien threat. Chosen by destiny, our heroes quickly discover they are the only ones who can save the planet. But to do so, they will have to overcome their real-life issues and before it's too late, band together as the Power Rangers.<sup>3</sup>Pretty typical post-Nolan franchise reboot on Lionsgate money. Doesn't really tap into Power Rangers in any enticing way. Not very impressive. Future trailers might excite more.<sup>3</sup>Trailer looks awesome, I like the humor, I like the tone. It's breakfast club meets Chronicle, with an Iron

Figure 23: Top activation of Feature 42307 in the 98,304-feature crosscoder on Pythia-160M.

## G.2 GENERALIZABILITY ACROSS DATASETS

We conduct experiments on the original Pythia 160M snapshots, but train and analyze 24,576-feature crosscoders on the FineWeb-Edu dataset (Penedo et al., 2024), which consists of educational web-pages. Our results (Figure 32) show nearly identical patterns to those observed with Pythia 160M crosscoders trained on SlimPajama, exhibiting very similar feature evolution dynamics.

## G.3 SENSITIVITY TO SNAPSHOT SELECTION

All of our above experiments train crosscoders on 32 pre-training snapshots of the language model. To further examine whether snapshot selection affects the trained crosscoder and the captured feature evolution, we train a 24,576-feature crosscoder on 16 pre-training snapshots of Pythia 160M, downsampled from the original set by selecting every other snapshot. The result (Figure 33) shows that the overall feature evolution exhibits nearly identical trends, with only a reduction in temporal resolution.

## H MORE EXAMPLES OF FEATURE

We additionally list some emergent features and demonstrate their top activating samples and decoder norm evolution in Figure 34.

## I THE USE OF LARGE LANGUAGE MODELS

The use of Large Language Models as an assistive tool in this paper is limited to the following two aspects:

1. The automated generation of feature complexity scores, as detailed in Section 4.2 and Appendix B.
2. Grammar correction and stylistic refinement of the manuscript.

<sup>3</sup>Obtained from <https://huggingface.co/stanford-crfm/alias-gpt2-small-x21>.

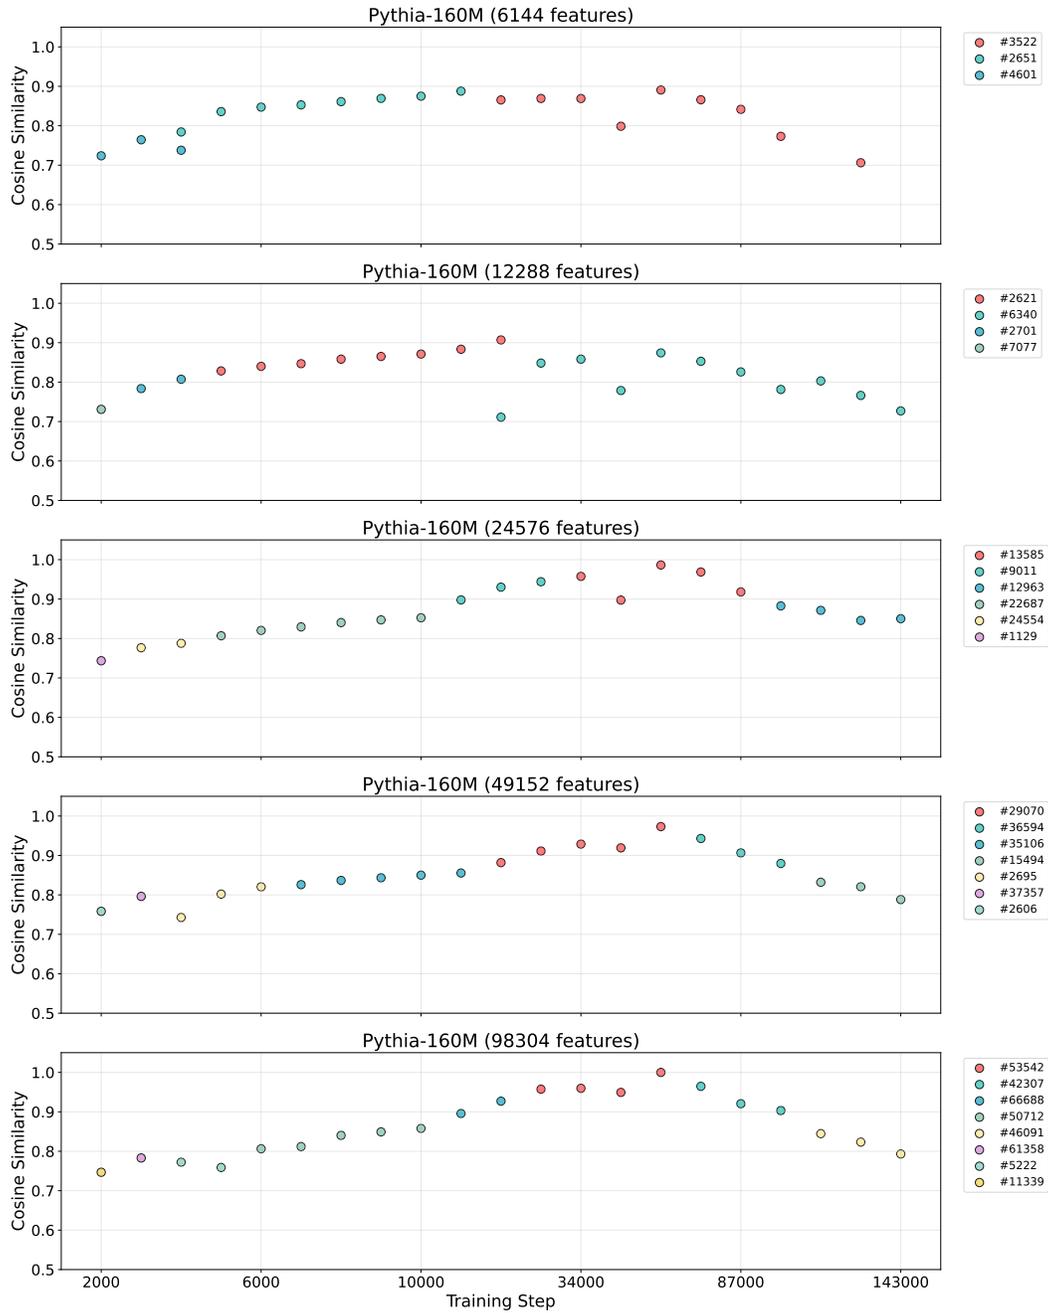


Figure 24: Crosscoder features similar to Feature 53542 of a 98,304-feature crosscoder. The number of features split increases with dictionary size.

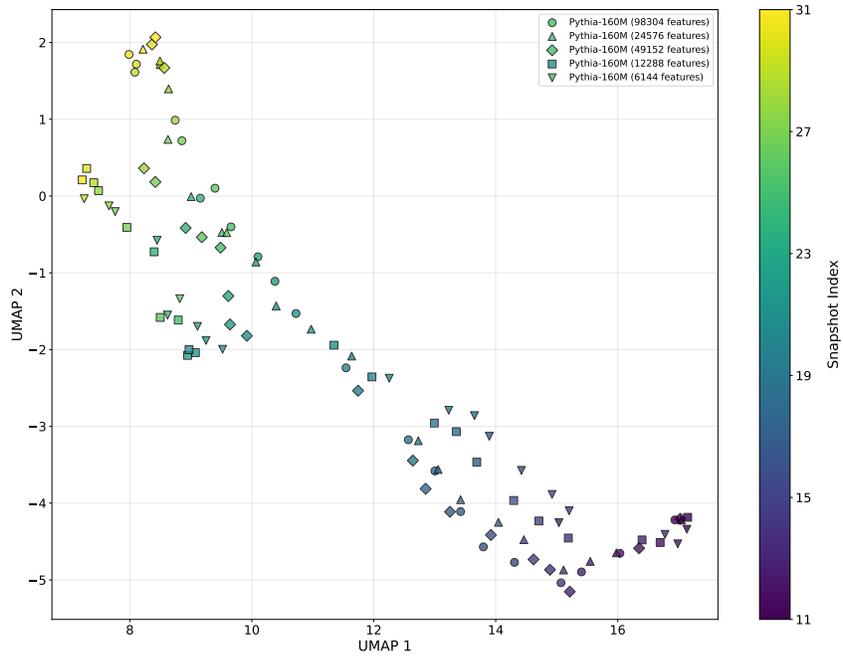


Figure 25: UMAP visualization of features shown in Figure 24.

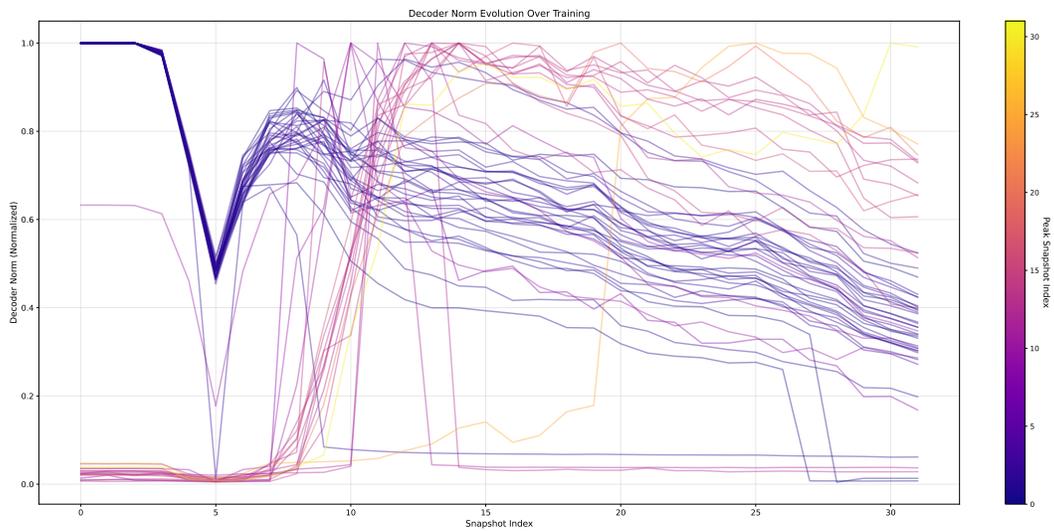


Figure 26: Overview of cross-snapshot feature decoder norm evolution in Pythia-160M Seed1.

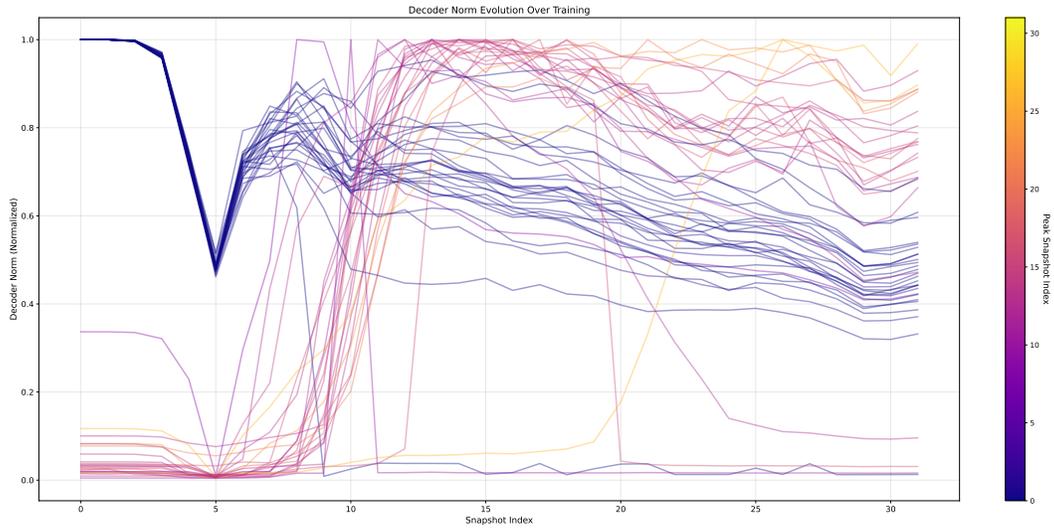


Figure 27: Overview of cross-snapshot feature decoder norm evolution in Pythia-160M Seed2.

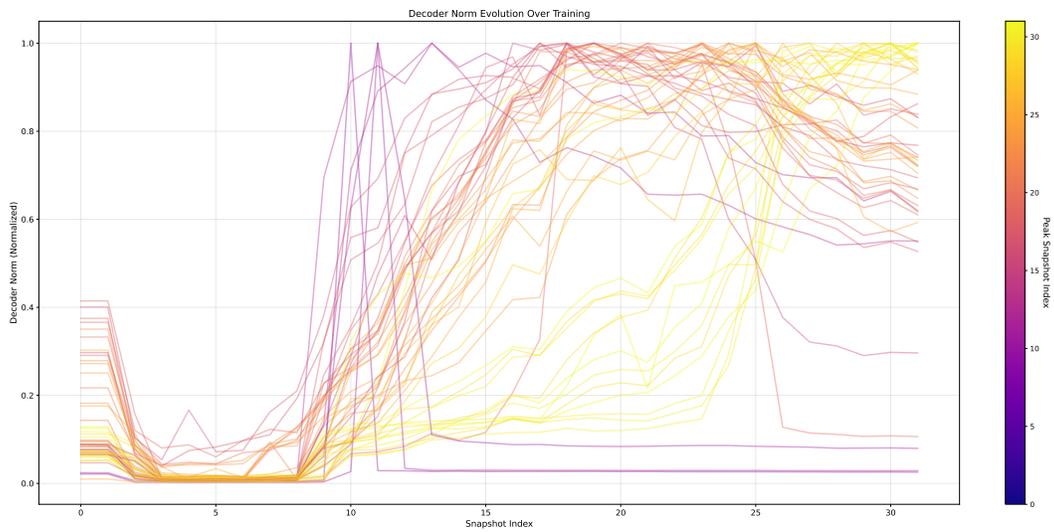


Figure 28: Overview of cross-snapshot feature decoder norm evolution in Alias run of Stanford CRFM GPT-2.

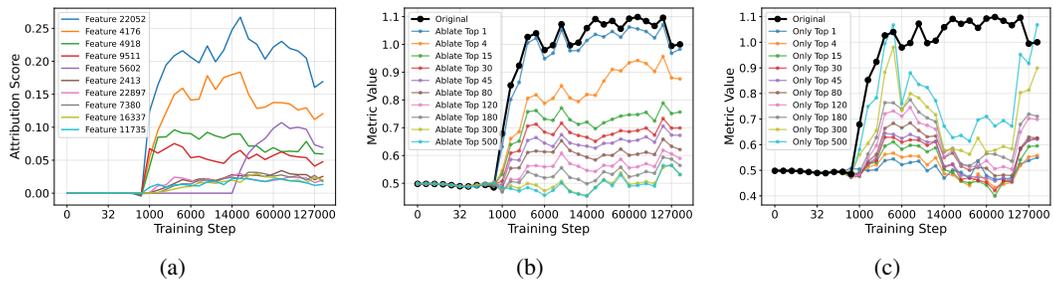


Figure 29: Crosscoder feature attribution on the Across-PP variant of the SVA task, using a cross-coder with 24,576 features trained on Pythia-160M Seed1.

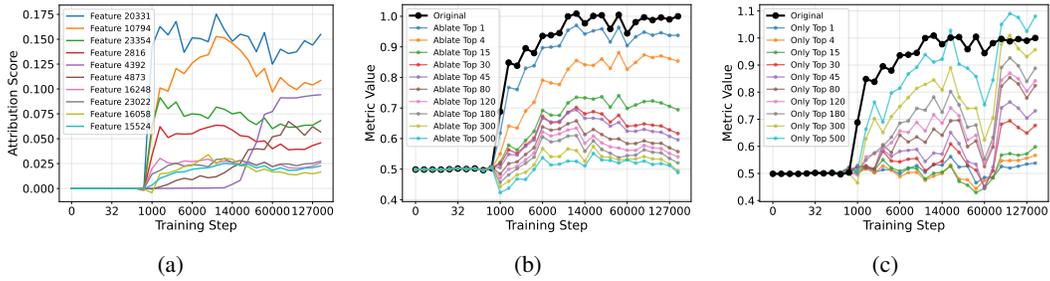


Figure 30: Crosscoder feature attribution on the Across-PP variant of the SVA task, using a cross-coder with 24,576 features trained on Pythia-160M Seed2.

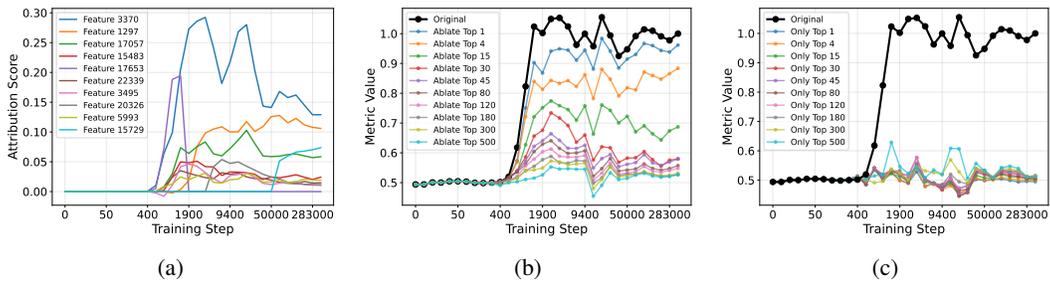


Figure 31: Crosscoder feature attribution on the Across-PP variant of the SVA task, using a cross-coder with 24,576 features trained on the Alias run of Stanford CRFM GPT-2.

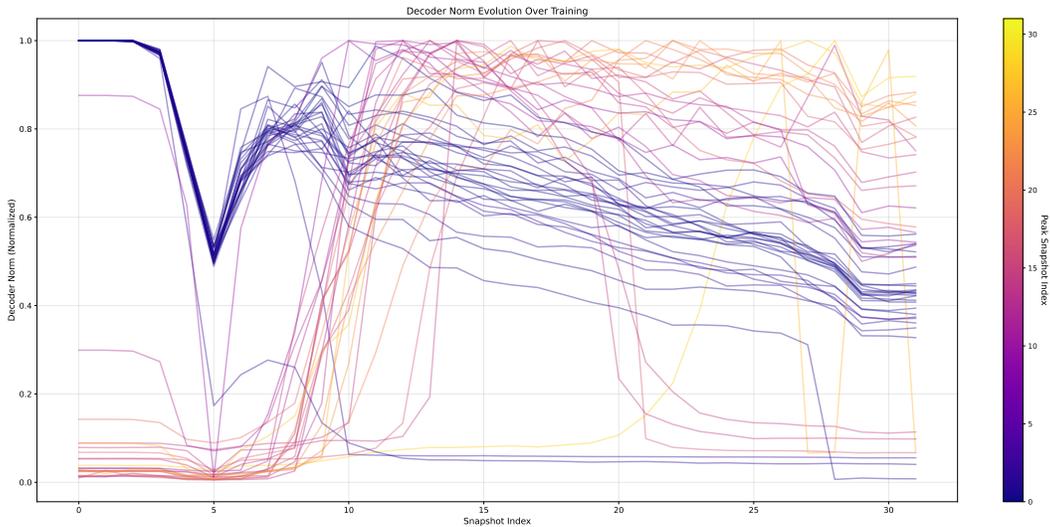


Figure 32: Overview of cross-snapshot feature decoder norm evolution in Pythia-160M, trained on FineWeb-Edu dataset.

