

PRIORS IN TIME: MISSING INDUCTIVE BIASES FOR LANGUAGE MODEL INTERPRETABILITY

000
001
002
003
004
005 **Anonymous authors**
006 Paper under double-blind review
007
008
009
010

ABSTRACT

011 A central aim of interpretability tools applied to language models is to recover
012 meaningful concepts from model activations. Existing feature extraction methods
013 focus on single tokens regardless of the context, implicitly assuming independence
014 (and therefore stationarity). This leaves open whether they can capture the rich
015 temporal and context-sensitive structure in the activations of language models
016 (LMs). Adopting a Bayesian view, we demonstrate that standard Sparse Autoen-
017 coders (SAEs) impose priors that assume independence of concepts across time.
018 We then show that LM representations exhibit rich temporal dynamics, including
019 systematic growth in conceptual dimensionality, context-dependent correlations,
020 and pronounced non-stationarity, in direct conflict with the priors of SAEs. This
021 mismatch casts doubt on existing SAEs' ability to reflect temporal structures of
022 interest in the data. We introduce a novel SAE architecture—Temporal SAE—with
023 a temporal inductive bias that decomposes representations at a given time into two
024 parts: a predictable component, which can be inferred from the context, and a
025 residual component, which captures novel information unexplained by the context.
026 Experiments on LLM activations with Temporal SAE demonstrate its ability to
027 correctly parse garden path sentences, identify event boundaries, and more broadly
028 delineate abstract, slow-moving information from novel, fast-moving information,
029 while existing SAEs show significant pitfalls in all the above tasks. Our results
030 underscore the need for inductive biases that match the data in designing robust
031 interpretability tools.

1 INTRODUCTION

032 Given the success of Language Models (LMs) (Bubeck et al., 2023; Deepmind, 2025), there is
033 growing interest in understanding how such models incrementally update over sequences of tokens to
034 exhibit complex behaviors (Murthy et al., 2025; Lindsey, 2025; Lindsey et al., 2025; Lepori et al.,
035 2025; Bigelow et al., 2025; Tuckute et al., 2024; Klindt et al., 2025). Interpretability research aims to
036 make such analyses tractable, offering tools for hypothesis design, testing, and intervention based on
037 evaluation of intermediate activations (Geiger et al., 2025; Sharkey et al., 2025; Bereska and Gavves,
038 2024). Often, such work builds on hypothesized computational models of how concepts are encoded
039 in a neural network's representations, e.g., the linear representation hypothesis (LRH) (Elhage et al.,
040 2022; Arora et al., 2018), correspondingly motivating tools such as sparse autoencoders (SAEs) (Gao
041 et al., 2024; Cunningham et al., 2023) for unsupervised extraction of a dictionary of vectors that
042 (ideally) mediate human-interpretable concepts (Mueller et al., 2025).

043 A central challenge in “bottom-up” approaches to interpretability, like SAEs, is the mismatch between
044 the assumptions of their underlying implementational account and the precise behavior or computation
045 they intend to explain (Jonas and Kording, 2017; Geiger et al., 2025; Costa et al., 2025) (see Fig. 1).
046 For instance, since LRH posits that different concepts correspond to directions in activation space
047 that can be independently manipulated, it implicitly claims the data distribution can be factorized
048 into independently varying latent variables (Allen, 2024). This mismatch between the structure
049 of the data distribution and strong priors codified in LRH can lead to misleading or pathological
050 explanations when using SAEs to understand neural networks (Chanin et al., 2025; Bricken et al.,
051 2023; Hindupur et al., 2025). This raises a set of critical questions for using SAEs to interpret models
052 trained on sequential data like language. Specifically, since language exhibits rich temporal structure
053 at *multiple scales* (Marslen-Wilson and Tyler, 1980; Thompson, 1999)—e.g., sentences contain

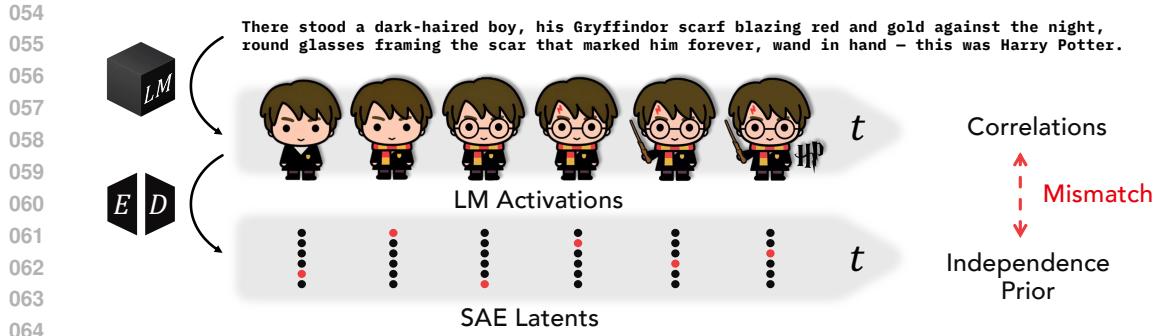


Figure 1: **The mismatch between SAE assumptions and temporal structure of language.** An illustrative sentence describing attributes of Harry Potter is shown. When passed into a language model (LM), it leads to activations x_t that include concepts within them (possibly entangled): note the presence of large numbers of shared attributes over time, which manifest as correlations across time of activations. Sparse Autoencoders (SAEs) implicitly have an independence (i.i.d.) prior across time t over their latents and thereby over concepts, which clashes with the true structure of language.

dependencies that link words across time (Gibson et al., 2000; McElree et al., 2003), upcoming words can be anticipated from context (Hale, 2001; Levy, 2008), and discourse imposes structure over longer timescales through phenomena like event boundaries (Zacks et al., 2007; Baldassano et al., 2017)—one can ask *what assumptions about temporal structures do SAEs make? How do these assumptions align with the actual temporal structure present in a LM’s activations?*

This work. Building on the Bayesian interpretation of sparse coding (Olshausen and Field, 1996; 1997)—the framework that motivates SAEs—we rephrase the optimization objective of SAEs as a MAP (maximum a posteriori) estimation problem. This allows us to make explicit prior assumptions about temporal structure embedded in SAEs, showing they implicitly assume concepts are *uncorrelated across time* and the number of concepts necessary to explain an activation is *time-invariant*—that is, the information present at each token position is independent of information at other positions and uniformly distributed. As we empirically show, these assumptions stand in stark contrast to the actual temporal structure present in language and language model representations, and can result in empirically observed pathologies in SAEs, such as feature splitting (Bricken et al., 2023; Chanin et al., 2025; Bussmann et al., 2025).

These results then motivate us to draw a broader parallel between SAEs and computational neuroscience approaches for understanding neural data. Specifically, population-level analyses of neural recordings have revealed that representations often lie on structured manifolds (Khona and Fiete, 2022; Nogueira et al., 2023; Sohn et al., 2019), challenging the reductionist assumption in sparse coding that computations occur via independently firing, monosemantic features (Eichenbaum, 2018; Saxena and Cunningham, 2019; Barack and Krakauer, 2021). This motivated a paradigm shift towards more structured analysis protocols—methods designed around the generative process of the behavior one aims to explain (Schneider et al., 2023; Chen et al., 2018). Motivated by this and similar findings of intricate geometrical structure in neural network representations (Fel et al., 2025; Gurnee et al., 2025; Modell et al., 2025), we propose **Temporal SAEs**, a new protocol for interpreting language model activations that incorporates explicit inductive biases about temporal structure. Our approach decomposes activations at each timestep into two orthogonal components: a *predictable component*, obtained by projecting current representations onto past context using a learned attention mechanism, and a *novel component*, representing residual information orthogonal to the predictable component. That is, we assume the *novel* component—not the total representation—is uncorrelated over time. This allows correlations between total codes and hence enables our method to capture the temporal structure of LM activations. Overall, we argue **interpretability methods should be driven by the behavior one is trying to explain**.

2 PRELIMINARIES

Notations. Let bold, lowercase letters represent vectors (e.g., z). Subscripts on vectors denote different samples (e.g., z_i), while superscripts denote the index within the vector, leading to a scalar

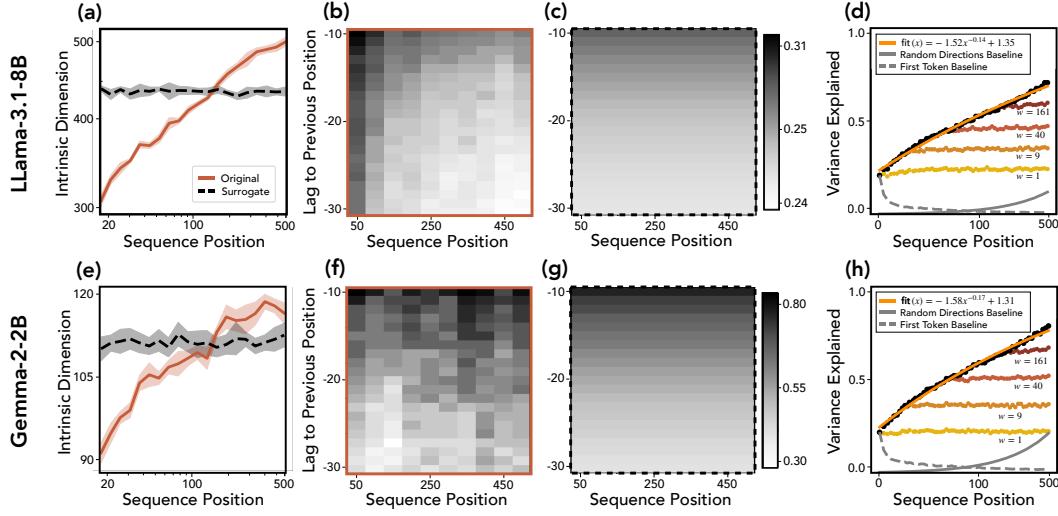


Figure 2: **Temporal structure of LLM activations reveals nonstationarity.** We use Pile samples (Monolog, 2021) to analyze temporal structure from activations of two pretrained LMs, comparing it to a surrogate signal that is stationary in nature (see App. J.3). **(a, e)** Intrinsic dimension of model activations and stationary surrogate. **(b, f)** Autocorrelations $A(x_t, x_{t-\tau})$ as a function of sequence position (t) and lag (τ). **(c, g)** Autocorrelation of the stationary surrogate. **(d, h)** Variance explained by projecting current representation x_t onto past context window $\{x_{t-1}, \dots, x_{t-w}\}$ with different sizes w , along with a baseline. Results consistently show representations getting ‘denser’ over time and being significantly more structured than a stationary surrogate.

(e.g., z^k). We denote model activations by $\mathbf{x} \in \mathbb{R}^n$, SAE latents (sparse code) by $\mathbf{z} \in \mathbb{R}^M$, and the dictionary by $\mathbf{D} \in \mathbb{R}^{n \times M}$ (M is the dictionary size).

Sparse Coding. Sparse dictionary learning (Olshausen and Field, 1996; 1997) expresses data as a sparse linear combination of dictionary elements, where both the weights and dictionary are learned from data. Intuitively, the dictionary behaves as a data-adaptive overcomplete basis; i.e., it typically has more elements than the dimension of ambient space. The optimization problem involved in this framework is $\arg \min_{\mathbf{D}, \mathbf{z}} \frac{1}{N} \sum_{i=1}^N \|\mathbf{x}_i - \mathbf{D}\mathbf{z}_i\|_2^2 + \lambda \mathcal{R}(\mathbf{z}_i)$, where $\mathcal{R}(\cdot)$, typically chosen to be the ℓ_1 -norm, is a sparsity-inducing regularizer. Sparsity assists in picking the fewest most relevant dictionary atoms to explain a given data point.

Sparse Autoencoders (SAEs). SAEs (Shu et al., 2025) aim to disentangle (Bengio et al., 2013; Higgins et al., 2018; Olah, 2023) neural network activations into human-interpretable concepts (Cunningham et al., 2023; Bricken et al., 2023). Specifically, SAEs transform their inputs (i.e., neural network activations) into a latent representation which is encouraged to be sparse. As shown by Hindupur et al. (2025), this is achieved by solving the sparse coding problem using a specific parametric form for the sparse codes:

$$\arg \min_{\mathbf{D}, \mathbf{z}} \frac{1}{T} \sum_{i=1}^T \|\mathbf{x}_i - \mathbf{D}\mathbf{z}_i\|_2^2 + \lambda \mathcal{R}(\mathbf{z}_i), \quad \text{s.t.} \quad \mathbf{z}_k = f_{\text{SAE}}(\mathbf{x}_k) \quad \forall k, \quad \tilde{g}(\mathbf{z}_1, \dots, \mathbf{z}_T) = 0, \quad (1)$$

where $\mathcal{R}(\cdot)$ is the regularizer (typically the L_1 norm), f_{SAE} is the SAE encoder architecture, and $\tilde{g}(\cdot)$ captures SAE-specific sparsity constraints on \mathbf{z} . f_{SAE} is typically a single hidden layer as in the ReLU SAE (Bricken et al., 2023; Cunningham et al., 2023), TopK SAE (Gao et al., 2024; Makhzani and Frey, 2013), JumpReLU SAE (Rajamanoharan et al., 2024) and BatchTopK SAE (Bussmann et al., 2024), though recent work has also explored alternative architectures inspired by sparse coding algorithms to capture specific structures, e.g., hierarchies (Muchane et al., 2025; Costa et al., 2025).

3 TEMPORAL STRUCTURE IN LANGUAGE MODEL ACTIVATIONS

To contextualize the prior assumptions made by SAEs about temporal structure in LMs’ activations, we first perform an empirical characterization of such temporal structure in pretrained LMs.

162 Specifically, since LMs are trained to generate coherent text by learning the distribution of natural
 163 language, one can expect their representations capture the rich phenomenology of its sequential
 164 structure (Elman, 1990); indeed, recent work has in fact found LM representations to be predictive
 165 of human neural recordings during language comprehension (Hosseini et al., 2024; Schrimpf et al.,
 166 2021; Tuckute et al., 2024; Hong et al., 2024; Georgiou et al., 2023). Motivated by this, we perform
 167 two experiments relevant to our discussion (see Fig. 2): (i) *measuring intrinsic dimensionality*—an
 168 approximation of the number of concepts necessary to explain the data, which can be expected to
 169 increase in a monotonic manner with time (Zhong et al., 2024; Can, 2025; Barak and Tsodyks, 2014;
 170 Meister et al., 2021)—and (ii) *signal nonstationarity*—which assesses whether model activations
 171 reflect the contextual relations between phrases of a passage (Zacks et al., 2007).

172 **Increasing intrinsic dimensionality.** Fig. 2 (a,e) show the dimensionality of the underlying manifold
 173 structure (intrinsic dimension) in model activations. We estimate the intrinsic dimensionality at a
 174 fixed position across a set of sequences with the U-statistic (App. Sec. J.2). For language model activations,
 175 this metric increases steadily with sequence position. On the other hand, a stationary surrogate
 176 of the data (see App. Sec. J.3) shows nearly constant intrinsic dimension over time. This indicates
 177 that model activations get ‘denser’, i.e., they possess more information over time. Correspondingly,
 178 the number of concepts needed to explain them varies with context.

180 **Non-stationarity: Context explains bulk of signal variance.** Subplots (b), (f) show the auto-
 181 correlation of model activations (App. Sec. J.1), which is noticeably different at different sequence
 182 positions (x-axis), while the stationary surrogate, as expected, shows nearly position-invariant auto-
 183 correlation values (subplots (c), (g)). This finding is a clear signature of time-dependent correlation
 184 structure, and therefore of non-stationarity. We quantify the similarity of a representation with its
 185 context in subplots (d), (h). Specifically, we project representations of token x_t at a given time t
 186 onto the subspace spanned by preceding representations in the context $\{x_{<t}\}$. These subplots show
 187 that up to 80% variance of x_t is explained by a context of 500 tokens, further highlighting strong
 188 cross-temporal correlations. Significant variance in the representation at time t , x_t , can be predicted
 189 (expressed) using representations from the past context.

4 TEMPORAL PRIORS OF SPARSE AUTOENCODERS

193 We now state the prior assumptions made by existing SAEs regards sequential structure in an input,
 194 contrasting these assumptions with the empirical results shown in Sec. 3. Specifically, building on
 195 the arguments used by Olshausen and Field (1997) to formalize the problem of sparse coding, we
 196 note that the SAE training objective (Eq. 1) can be interpreted from a Bayesian lens: minimize the
 197 negative log posterior $\arg \min_{\{z_t\}} -\log P(z_1, \dots, z_T | x_1, \dots, x_T)$ of the data, which, by Bayes’
 198 rule, can be written as the sum of log likelihood (MSE) and log prior (the regularizer \mathcal{R}). From this
 199 lens, SAEs’ prior assumptions on sequential structure in LM activations can be described as follows.

200 **Proposition 4.1** (Independence prior over time). *Consider the SAE maximum a posteriori (MAP)
 201 objective from Eq. 1. Since the sparsity constraints are additive over time, this objective has an
 202 independent and identically distributed (i.i.d.) prior over time:*

$$P(z_1, \dots, z_T) \propto \prod_{t=1}^T \exp \left(-\lambda \mathcal{R}(z_t) - \tilde{\lambda} \tilde{g}(z_t) \right) = \prod_i P(z_i). \quad (2)$$

208 A more precise version of the claim for specific SAE architectures is provided in Appendix I.1.
 209 Intuitively, the claim above says that SAEs assume an independence of latents, and hence the
 210 concepts underlying the generative process of language, over time. This directly conflicts with the
 211 rich contextual structure of LM activations we empirically observed in Fig. 2b–d, f–h. Crucially, this
 212 also implies that SAEs assume the sparsity of latent codes necessary to explain model activations to
 213 be time-invariant, as stated formally in the corollary below.

214 **Corollary 4.1.1** (Assumptions of time-invariant sparsity). *As a consequence of the i.i.d. priors over
 215 time from Prop. 4.1, standard SAEs assume that sparsity of representations emerges from a fixed
 216 distribution independently over time (i.i.d.), i.e., $P(\|z_1\|_0, \dots, \|z_T\|_0) = \prod_t P(\|z_t\|_0)$.*

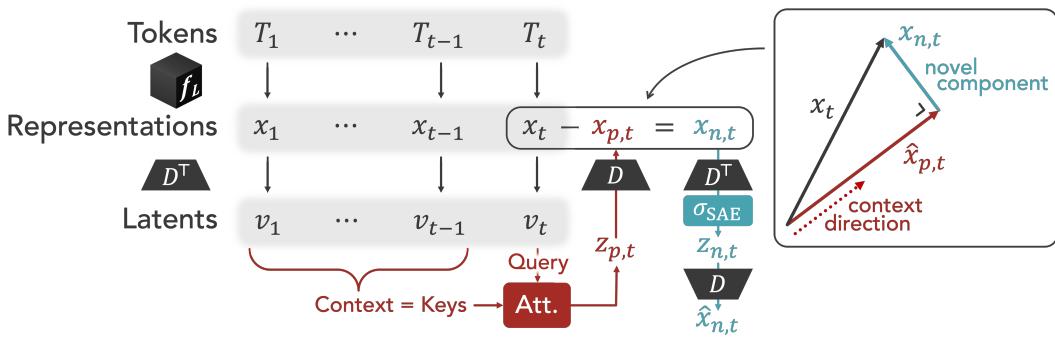


Figure 3: **Schematic of Temporal SAEs.** Temporal SAEs decompose activations x_t into two components: a predictable component, obtained by projecting x_t onto a context direction (derived from the past $x_{<t}$ using attention), and a sparse, novel component orthogonal to the predictable component that captures new information seen at time t .

SAEs thus assume that sparsity—which, in their underlying generative model of activations corresponds to the number of concepts necessary for explaining the data (Bricken et al., 2023; Elhage et al., 2022)—remains approximately constant over time. This again does not align with the increasing dimensionality of representations observed in LM activations (see Fig. 2a,e). Correspondingly, if enough concepts aggregate over context such that a model’s activations become ‘denser’ than the assumed sparsity budget, the assumption of time-invariant sparsity implies SAEs can fail to capture the temporal structure inherent in language—as is arguably already observed empirically with phenomena like feature splitting (Chanin et al., 2025; Bussmann et al., 2025; Bricken et al., 2023; Shu et al., 2025). We empirically validate this claim in Sec. 6.

5 TEMPORAL SAEs: EXPLICITLY MODELING TEMPORAL PRIORS

As stated in Sec. 2, sparse coding, a framework designed in computational neuroscience to understand neural representations in biological brains (Olshausen and Field, 1996; 1997), inspired SAEs as a framework for interpreting artificial neural networks (Bricken et al., 2023; Cunningham et al., 2023). In fact, the parallels between these communities can be made deeper: motivated by observations of intricate geometry of neural representations derived out of multi-dimensional population analyses (Khona and Fiete, 2022; Nogueira et al., 2023; Sohn et al., 2019), there were calls in computational neuroscience to discard the limiting reduction assumed in sparse coding that computations occur via a set of independently firing, monosemantic features (Eichenbaum, 2018; Saxena and Cunningham, 2019; Barack and Krakauer, 2021; Seung, 1996; Chung and Abbott, 2021)—similar to our arguments in Sec. 3, 4 (and results that follow in Sec. 6). Correspondingly, a need for more structured protocols was suggested (Eichenbaum, 2018; Barack and Krakauer, 2021), leading to methods that were motivated by the generative process of the behavior one is trying to explain (Schneider et al., 2023; Chen et al., 2018; Chen, 2019; Wiskott and Sejnowski, 2002; Linderman et al., 2017). We argue a similar paradigm shift is needed in language model interpretability: given that we train models to learn the distribution of highly structured data, we ought to embrace the fact that neural network activations can exhibit intricate geometrical organization. In what follows, as an attempt to qualify our arguments, we propose *one* such approach that focuses on the temporal structure of LM activations.

Temporal SAEs. In computational neuroscience, when analyzing data from dynamical domains (e.g., audio, language, or video), a commonly made assumption is that there is contextual information present in the recent history that informs the next state—this part of the signal is deemed *predictable* (Chen et al., 2025; Millidge et al., 2024), *slow-changing* (Berkes and Wiskott, 2005), *invariant* (Olshausen and Cadieu, 2007), or *dense* (Tasissa et al., 2022). Meanwhile, the remaining signal corresponds to new bits of information added by the observed state at the next timestep—this part can be deemed *novel / surprising*, *fast-changing*, *variant*, or *sparse* with respect to the context. We argue LM activations are amenable to a similar generative model. Specifically, our observations in Sec. 3 show that activations x_t at time t are strongly correlated with the context and can be

270 decomposed into two such parts. We thus propose the following generative model of LM activations:
 271

$$272 \quad \mathbf{x}_t = \mathbf{x}_{p,t} + \mathbf{x}_{n,t}, \quad \text{where} \quad \mathbf{x}_{p,t} = \mathbf{D}\mathbf{z}_{p,t} \text{ and } \mathbf{x}_{n,t} = \mathbf{D}\mathbf{z}_{n,t}. \quad (3)$$

273 In the above, $\mathbf{x}_{p,t}$ denotes a *predictable* component of the signal that captures the correlations of \mathbf{x}_t
 274 with past data $\{\mathbf{x}_{<t}\}$, while $\mathbf{x}_{n,t}$ denotes a *novel* component that represents new information added
 275 by the current token \mathbf{x}_t . To obtain $\mathbf{z}_{p,t}$, we project \mathbf{x}_t onto $\{\mathbf{x}_{<t}\}$ to explain the predictable variance
 276 in \mathbf{x}_t as a convex combination of past data. Specifically, we use a self-attention layer f on top of
 277 a single ReLU layer, yielding $\mathbf{z}_{p,t} = f(\{\mathbf{x}_1, \dots, \mathbf{x}_{t-1}\}, \mathbf{x}_t)$. Meanwhile, $\mathbf{z}_{n,t} = \tilde{f}(\mathbf{x}_t, \mathbf{z}_{p,t}) =$
 278 $\sigma(\mathbf{D}^T(\mathbf{x}_t - \mathbf{D}\mathbf{z}_{p,t}))$ captures the residual component of the code, which is not correlated with the
 279 past (note that σ is the nonlinearity). We use a standard SAE encoder (either TopK or BatchTopK)
 280 to instantiate σ , applying it to $\mathbf{x}_t - \mathbf{D}\mathbf{z}_{p,t}$ to derive $\mathbf{z}_{n,t}$. See Fig. 3 for an overall schematic of the
 281 encoding process. The learning objective in Temporal SAEs follows.

$$282 \quad \arg \min_{\mathbf{D}, \mathbf{z}} \frac{1}{T} \sum_{i=1}^T \|\mathbf{x}_i - \mathbf{D}(\mathbf{z}_{p,i} + \mathbf{z}_{n,i})\|_2^2 + \lambda \mathcal{R}(\mathbf{z}_{n,i}), \quad (4)$$

$$284 \quad \text{s.t. } \mathbf{z}_{p,k} = f_{\text{SAE}}(\{\mathbf{x}_{<k}\}, \mathbf{x}_k), \quad \mathbf{z}_{n,k} = \tilde{f}_{\text{SAE}}(\mathbf{x}_k, \mathbf{z}_{p,k}), \quad \mathbf{z}_k = \mathbf{z}_{p,k} + \mathbf{z}_{n,k} \quad \forall k.$$

$$285$$

286 Relating to Sec. 4, we note the prior assumption in Temporal SAEs is that the residual $\mathbf{z}_{n,t} = \mathbf{z}_t - \mathbf{z}_{p,t}$,
 287 which captures the novel information in \mathbf{x}_t remaining after removing the projections onto the past
 288 context, is *i.i.d.* over time. This prior allows temporal correlations between codes \mathbf{z} , and thereby allows
 289 correlations between concepts across time, instead of assuming them to be temporally independent.

290 **Sanity Checking Temporal SAEs.** Before
 291 analyzing how different approaches represent
 292 the temporal structure of language, we demon-
 293 strate that Temporal SAEs performs on par with
 294 SAEs on standard metrics such as recon-
 295 struction error. Specifically, we train a Temporal
 296 SAE and standard SAEs (ReLU, TopK, Batch-
 297 TopK) on 1B token activations extracted from
 298 Gemma-2-2B (Team et al., 2024) from the Pile-
 299 Uncopyrighted dataset (Monology, 2021). We
 300 also analyze a baseline of the prediction only
 301 module from Temporal SAEs, reported as ‘Pred.
 302 only’, which can be expected to underperform
 303 since predicting the next-token representation is
 304 likely to be more difficult than reconstructing
 305 it. Results are provided in Tab. 1, 2 and show
 306 competitive performance between all protocols,
 307 except Pred. only. One can also assess which
 308 part of a fully trained Temporal SAE is more
 309 salient in defining its performance, i.e., does the
 310 estimated predictive part $\hat{\mathbf{x}}_{p,t}$ contribute more to
 311 the reconstruction $\hat{\mathbf{x}}$ or does the estimated novel
 312 part $\hat{\mathbf{x}}_{n,t}$. Results are reported in Tab. 3. We see
 313 that the error vectors, i.e., $\mathbf{x} - \hat{\mathbf{x}}_p$ and $\mathbf{x} - \hat{\mathbf{x}}_n$, are
 314 approximately orthogonal, suggesting the mod-
 315 ules capture separate bits of information from
 316 the input—this is inline with results by Costa et al. (2025), who show optimizing to reconstruct
 317 residuals (as we do with the novel code) *can* lead to orthogonal codes at different stages of an SAE.
 318 Furthermore, we find that a bulk of the reconstructed signal $\hat{\mathbf{x}}_t$ (in the sense of norm) is captured
 319 by the predictive code—in fact, the percentage contribution of the predictive code is $\sim 80\%$, in line
 320 with numbers observed in Fig. 2d. However, analyzing the reconstruction performance, we see the
 321 predictive component primarily contributes to achieving a good NMSE, while the novel component
 322 is more responsible for explaining the input signal variance. These results align with the generative
 323 model assumed in Temporal SAEs (Eq. 3). Specifically, NMSE captures the average reconstruction,
 324 and hence a slower moving, contextual signal can expect to dominate its calculation. Meanwhile,
 325 variance assesses changes per dimension and timestamp in the signal, which better matches the
 326 inductive bias imposed on the novel part.

Table 1: Temporal SAEs achieves NMSE similar to standard SAEs across domains (Simple Stories, Webtext, Code).

	ReLU	TopK	BTOPK	Pred. Only	Temporal
Story	0.20	0.155	0.152	0.34	0.139
Web	0.19	0.144	0.139	0.36	0.139
Code	0.20	0.154	0.149	0.38	0.152

Table 2: Temporal SAEs explains similar amount of signal variance as standard SAEs.

	ReLU	TopK	BTOPK	Pred. Only	Temporal
Story	0.60	0.71	0.72	0.29	0.73
Web	0.69	0.78	0.79	0.40	0.79
Code	0.65	0.75	0.75	0.33	0.75

Table 3: predictive and novel codes explain different parts of the input signal across domains.

	Sim.	% Norm		NMSE		Var. Expl.	
		Pred.	Novel	Pred.	Novel	Pred.	Novel
Story	-0.02	76.2	23.5	0.53	4.03	0.11	0.64
Web	-0.02	80.5	19.5	0.49	4.28	0.17	0.66
Code	-0.02	74.2	26.0	0.57	3.84	0.14	0.65

6

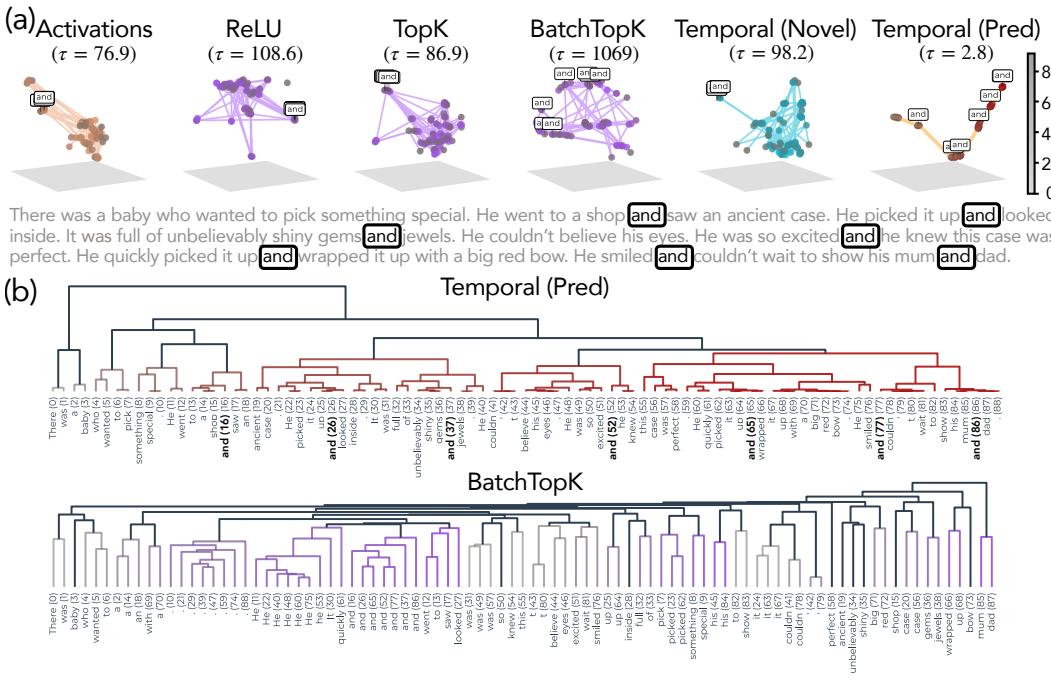


Figure 4: **Temporal SAEs unroll stories, decomposing into events.** We consider model activations from a story and compute pairwise similarity of codes extracted from different interpretability protocols. (a) We see predictive codes from Temporal SAE organizes in hierarchical block structures that seem to align with (sub)event boundaries in the analyzed story, while the novel code primarily emphasizes sudden changes in the narrative; meanwhile, standard SAEs show a mixture of the two structures, with a stronger similarity to the structure exhibited by the novel codes. (b) We confirm the alignment of predictive codes with event boundaries by running an off-the-shelf hierarchical clustering algorithm, finding the token clusters indeed correspond to (sub)events occurring in the story as the narrative proceeds. Running this process on SAEs, we find this process yields temporally incoherent clusters that are primarily defined by lexical information.

6 CAPTURING TEMPORAL STRUCTURE WITH TEMPORAL SAE

We now evaluate the ability of different SAEs to capture temporal dynamics in language model representations. We first analyze a narrative setting where, locally in time, one can expect a lot of correlated structure as an event transpires, with strict event boundaries delineating events from each other. As we show, TemporalSAEs yield a clear delineation of the slow-moving local information from the fast-changing boundaries, while standard SAEs generally ignore the slow-moving information to optimize for the faster changes. The narrative evaluation assesses how different SAE codes represent local/global semantic information. We now investigate local/global syntactic information by analyzing SAE codes extracted from garden path sentences. Here, we find that unlike other SAEs, the use of a predictive module in TemporalSAEs yields codes that relate tokens from garden path sentences in a manner that would align with the ultimately correct (rather than garden-path) parses—an ability that language models are in fact known to possess, but that standard SAEs seem to not explain (Li et al., 2024; Hanna and Mueller, 2024).

6.1 THE GEOMETRY OF STORIES: A NARRATIVE-DRIVEN DOMAIN

UMAP of Latent Codes Suggests Models Temporally Straighten Activations. We consider the TinyStories datasets (Eldan and Li, 2023) for its relatively straightforward narrative structures, and qualitatively analyze the geometry of latent codes extracted from model activations when processing these stories. Visualizing the latent codes in a low-dimensional basis via a 3D UMAP projection (McInnes et al., 2018), we see SAEs yield a highly irregular and unstructured geometry (see Fig. 4a). Calculating Tortuosity (Bullitt et al., 2003), a measure of how aligned local arcs are

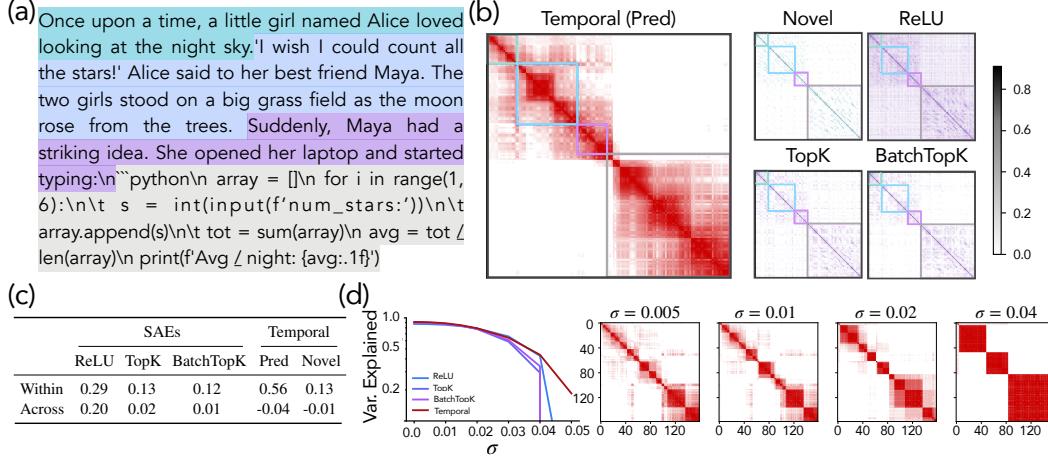


Figure 6: **Predictive codes decompose stories into events.** (a) We consider model representations from a synthetic story with well-defined event boundaries. (b) Computing pairwise cosine similarity of latent codes extracted using different protocols, we see the predictive code of Temporal SAE organizes in hierarchical block structures that seem to align with (sub)event boundaries in the analyzed story. (c) We confirm the alignment of predictive codes with event boundaries by computing average pairwise similarity of token latent codes for tokens that span the same event (‘within’) versus not (‘across’). Results clearly show high within-event similarity scores for the predictive code. (d) The results above are further corroborated by running a noising process on the latent codes: we add Gaussian noise of scale σ to the input before computing latent codes, defining the similarity maps and computing explained variance of un-noised data. This process elicits coarser grained clusters from the similarity maps for the predictive code, suggesting the multi-scale temporal structure of stories is reflected in predictive codes.

with respect to the global structure of a curve, we see very high values emerge for SAEs’ latent codes geometry, suggesting sudden changes in the local similarity as a story unravels. To further understand the results above, we highlight a specific token (‘and’) from the story, the UMAP analysis shows that standard SAEs generally just cluster tokens by lexical identity. This is further corroborated by running a hierarchical clustering algorithm on the latent codes (SciPy, 2025), finding temporally incoherent, but lexically related clusters (see Fig. 4b).

Quantifying Similarity to Slow vs. Fast Moving Signals. To further quantify the straightening claim, we compute the Fourier transform of the model activations and divide the frequency spectrum into two halves at a critical frequency f_c such that the energy (i.e., sum of squared value of phase information) in the frequencies below f_c equals that of the remaining ones. We call the first split “slow part” of a sequence, and latter the “fast part”. We then compute the correlation matrix defined by the slow and fast parts, compute their spectrum, and analyze how similar these spectrums are to the ones defined using different SAEs’ codes. Results are shown in Fig. 5. We clearly see the spectrum extracted from the predictive component of Temporal SAE approximates that of the slow part of the representation, while the novel component’s spectrum is similar to that of the fast moving part. Meanwhile, spectra of SAEs only exhibits similarity to the fast part, suggesting they do not capture longer range dependencies necessary for interpreting narrative-driven texts like most language domains. To quantify this further, we also measure kernel similarities (CKA) between the kernels respectively defined by the

Table 4: **Similarity between latent codes and model activations.** Predictive codes from Temporal SAEs align better with slow-changing part of activations; novel codes and standard SAEs’ codes align with fast-changing part.

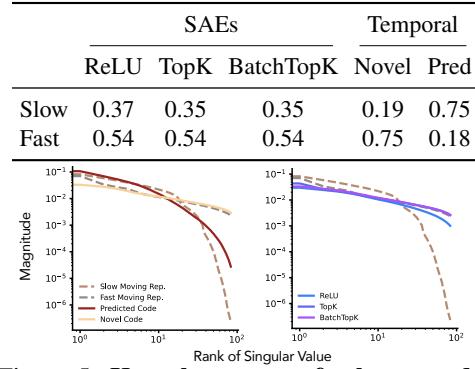
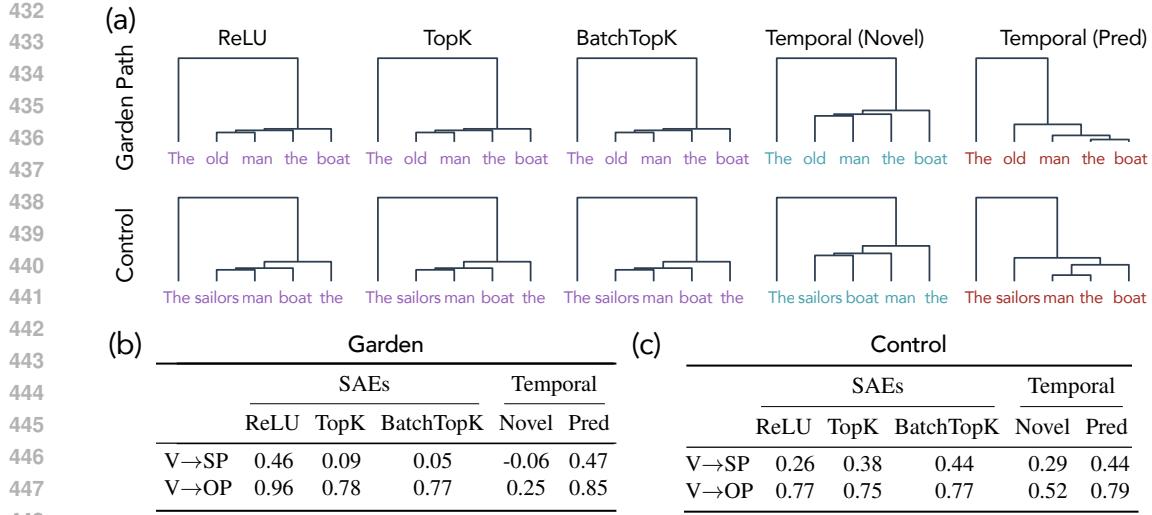


Figure 5: **Kernel spectrum for latent codes and model representations.** Kernels defined using novel code from Temporal SAE and standard SAEs both align well with the fast-changing part of model representations; meanwhile, only the predictive code shows strong similarity to the slow changing part.



449 Figure 7: **Hierarchically clustering SAE codes for garden path sentences.** (a) Pairwise similarity
 450 maps of the predictive code from Temporal SAE link long-distance heads and dependents that
 451 define the ultimately correct parse in garden path sentences, while standard SAEs (e.g., BatchTopK)
 452 emphasize only local, transient relations, falling for the misleading cues. (b, c) Comparing the cosine
 453 similarity of average latent code extracted from the subject phrase (SP), verb phrase (V), and object
 454 phrase (OP), we see across ambiguous garden path sentences and unambiguous control variants
 455 thereof, only the predictive component of Temporal Features Analyzers shows consistent similarity
 456 scores (as expected if the SP and V ambiguity is reflected in the latent codes).

457
 458 slow moving and fast moving signal to the different SAE codes. Results clearly show Standard SAEs
 459 are primarily similar to the fast moving signal.

460
 461 **Predictive Component Captures Local Event Boundaries.** The results above demonstrate Temporal
 462 SAEs’ predictive component qualitatively align with event boundaries in a story. To investigate this
 463 result more quantitatively, we use GPT-5 to create a synthetic dataset of 50 stories with well-defined
 464 event boundaries (see Fig. 6a for an example). We extract latent codes for these stories’ tokens, center
 465 them by subtracting the mean to remove any globally shared information, and compute the cosine
 466 similarity of token to token latent codes. If the latent codes reflect local event structure of a story,
 467 the cosine similarity (on average) will be high between token pairs that come from the same event
 468 and low (if not zero) between pairs sampled across events; see Fig. 6b for an example similarity map
 469 corresponding to the story shown in Fig. 6a. Results are shown in Fig. 6 (c) and corroborate our
 470 qualitative findings: we see predictive components from Temporal SAE show substantially higher
 471 similarity of codes if tokens are sampled from within an event, while the novel component and
 472 SAEs generally show low similarity between two tokens. These results are further supported by the
 473 robustness of Temporal SAEs to noise. Specifically, we see that when we add noise to the input data,
 474 which, on average, will lead to turning off of latents with small magnitudes (due to the encoding
 475 nonlinearity), the temporal structure of the data, if it is present, will be amplified. We see precisely
 476 this effect in Fig. 6d: the predictive components’ cosine similarity map under Gaussian noised input
 477 maps elicits coarser block structures with increasing noise scale; this is reminiscent of percolation
 478 or heat diffusion perspectives on graph clustering, wherein noise diffuses only within a connected
 479 component and hence community structure is elicited (Von Luxburg, 2007). Correspondingly, we see
 480 Temporal SAEs respond most gracefully to noise: variance explained reduces slower than SAEs’,
 481 which in fact drops to ~ 0 at some scale.

482 6.2 GARDEN PATH SENTENCES: AMBIGUITIES RESOLVED VIA TEMPORAL STRUCTURE

483
 484 Garden-path sentences—e.g., The old man the boat—initially cue an incorrect local parse
 485 before a later token forces reanalysis. Language models have been shown to be able to correctly parse
 such ambiguous sentences, offering in fact a predictive account of human per-token surprisals (Li

486 et al., 2024; Hanna and Mueller, 2024; Oh and Schuler, 2023). Interestingly, when using SAE codes to
 487 assess whether LLM representations offer a valid parse of the sentence, we find hierarchical clustering
 488 of SAE codes yields a parse that is suggestive of the misleading cue; meanwhile, Temporal SAE
 489 recovers the correct parse by separating the predictable, slow-moving component of the representation
 490 from the novel, fast-changing residual. Specifically, in Fig. 7, we see the predictive codes link long-
 491 distance heads and dependents that reflect the correct parse (e.g., man as verb), producing coherent
 492 similarity structure over the full span, whereas standard SAEs emphasize only transient, local changes
 493 and miss these cross-temporal constraints. These results suggest Temporal SAEs encode syntactic
 494 structure that unfolds over time when evidence to collapse the correct constituent parse emerges.
 495 To make these results more quantitative in nature, we use GPT-5 to synthetically generate a set
 496 of 50 garden path sentences where the subject is ambiguous. We create 50 control variants of
 497 these sentences such that the controls do not possess ambiguity with respect to typical parse of the
 498 sentence constituents: e.g., changing the subject in the sentence The old man the boat from
 499 old to sailors, yielding The sailors man the boat. We then divide all sentences into
 500 their respective subject phrase (SP), verb phrase (V), and object phrase (OP): e.g., The old (SP),
 501 man (V), and the road (OP). We compute the average latent code for tokens from these three
 502 constituent phrases, under the hypothesis that if the sentence ambiguity is reflected in the latent code,
 503 then until the OP shows up, the correct parse of prior words cannot be identified. Correspondingly, all
 504 valid parses must be stored in the same representation. This suggests the cosine similarity of latent
 505 codes of V and SP tokens should be of a similar order in both the garden path and control sentences;
 506 meanwhile, the similarity between V and OP should be much higher than V and SP. Results are
 507 reported in Fig. 7 (b,c). We clearly see extreme sensitivity in similarity values of SAE latent codes
 508 and the novel component of Temporal SAE, but the predictive component is essentially invariant
 509 across sentence type, suggesting it captures the temporal dynamics likely relevant for a LM to parse
 510 garden path sentences.

7 DISCUSSION

511 Our findings reinforce the broader lesson that interpretability tools must align their inductive biases
 512 with the statistical structure of the data they are applied to. We showed that standard SAEs impose
 513 independence priors across time, which are fundamentally misaligned with the nonstationary and
 514 context-dependent structure of language model activations. This mismatch explains why existing
 515 SAEs tend to underrepresent temporal dependencies, despite capturing other kinds of structure. By
 516 contrast, Temporal SAE incorporates empirically observed correlations across time as an inductive
 517 bias. Its decomposition of activations into predictable (slow-moving) and novel (fast-changing)
 518 components enables the recovery of temporal structure that standard SAEs fail to expose. In particular,
 519 we demonstrated that predictable codes align with stable, high-level information, while novel codes
 520 isolate transient or surprising information, allowing Temporal SAE to highlight event boundaries
 521 and syntactic reanalyses in garden-path sentences. Taken together, these results emphasize a general
 522 principle: interpretability methods should not be viewed as neutral feature extractors but as models
 523 with their own structural assumptions. When these assumptions mismatch the true data distribution,
 524 important aspects of representation may be obscured. Incorporating empirically motivated temporal
 525 priors offers one way to close this gap, suggesting that future progress in interpretability will require
 526 tailoring methods to the dynamics of the representations under study.

527
 528
 529
 530
 531
 532
 533
 534
 535
 536
 537
 538
 539

540 REFERENCES
541

542 Carl Allen. Unpicking data at the seams: Understanding disentanglement in vaes. *arXiv preprint*
543 *arXiv:2410.22559*, 2024.

544 Sanjeev Arora, Yuanzhi Li, Yingyu Liang, Tengyu Ma, and Andrej Risteski. Linear algebraic structure
545 of word senses, with applications to polysemy. *Transactions of the Association for Computational*
546 *Linguistics*, 6:483–495, 2018.

547 Christopher Baldassano, Janice Chen, Asieh Zadbood, Jonathan W Pillow, Uri Hasson, and Kenneth A
548 Norman. Discovering event structure in continuous narrative perception and memory. *Neuron*, 95
549 (3):709–721, 2017.

550 David L Barack and John W Krakauer. Two views on the cognitive brain. *Nature Reviews Neuroscience*,
551 22(6):359–371, 2021.

552 Omri Barak and Misha Tsodyks. Working models of working memory. *Current opinion in neurobiology*,
553 25:20–24, 2014.

554 Yoshua Bengio, Aaron Courville, and Pascal Vincent. Representation learning: A review and new
555 perspectives. *IEEE transactions on pattern analysis and machine intelligence*, 35(8):1798–1828,
556 2013.

557 Leonard Bereska and Efstratios Gavves. Mechanistic interpretability for ai safety—a review. *arXiv*
558 *preprint arXiv:2404.14082*, 2024.

559 Pietro Berkes and Laurenz Wiskott. Slow feature analysis yields a rich repertoire of complex cell
560 properties. *Journal of vision*, 5(6):9–9, 2005.

561 Eric Bigelow, Daniel Wurgافت, YingQiao Wang, Noah Goodman, Tomer Ullman, Hidenori Tanaka,
562 and Ekdeep Singh Lubana. Belief dynamics reveal the dual nature of in-context learning and
563 activation steering. *arXiv preprint arXiv:2511.00617*, 2025.

564 Trenton Bricken, Adly Templeton, Joshua Batson, Brian Chen, Adam Jermyn, Tom Conerly, Nick
565 Turner, Cem Anil, Carson Denison, Amanda Askell, Robert Lasenby, Yifan Wu, Shauna Kravec,
566 Nicholas Schiefer, Tim Maxwell, Nicholas Joseph, Zac Hatfield-Dodds, Alex Tamkin, Karina
567 Nguyen, Brayden McLean, Josiah E Burke, Tristan Hume, Shan Carter, Tom Henighan, and
568 Christopher Olah. Towards monosemanticity: Decomposing language models with dictionary
569 learning. *Transformer Circuits Thread*, 2023. <https://transformer-circuits.pub/2023/monosemantic-features>.

570 Sébastien Bubeck, Varun Chandrasekaran, Ronen Eldan, Johannes Gehrke, Eric Horvitz, Ece Kamar,
571 Peter Lee, Yin Tat Lee, Yuanzhi Li, Scott Lundberg, et al. Sparks of artificial general intelligence:
572 Early experiments with gpt-4. *arXiv preprint arXiv:2303.12712*, 2023.

573 Elizabeth Bullitt, Guido Gerig, Stephen M Pizer, Weili Lin, and Stephen R Aylward. Measuring
574 tortuosity of the intracerebral vasculature from mra images. *IEEE transactions on medical imaging*,
575 22(9):1163–1171, 2003.

576 Bart Bussmann, Patrick Leask, and Neel Nanda. Batchtopk sparse autoencoders, 2024. URL
577 <https://arxiv.org/abs/2412.06410>.

578 Bart Bussmann, Noa Nabeshima, Adam Karvonen, and Neel Nanda. Learning multi-level features
579 with matryoshka sparse autoencoders. *arXiv preprint arXiv:2503.17547*, 2025.

580 Tankut Can. Statistical mechanics of semantic compression. *arXiv preprint arXiv:2503.00612*, 2025.

581 David Chanin, James Wilken-Smith, Tomas Dulka, Hardik Bhatnagar, Satvik Golechha, and Joseph
582 Bloom. A is for absorption: Studying feature splitting and absorption in sparse autoencoders, 2025.
583 URL <https://arxiv.org/abs/2409.14507>.

584 Yibe Chen, Zaid Zada, Samuel A Nastase, F Gregory Ashby, and Satrajit S Ghosh. Context modulates
585 brain state dynamics and behavioral responses during narrative comprehension. *bioRxiv*, pages
586 2025–04, 2025.

594 Yubei Chen. *The Sparse Manifold Transform and Unsupervised Learning for Signal Representation*.
 595 PhD thesis, University of California, Berkeley, 2019.
 596

597 Yubei Chen, Dylan Paiton, and Bruno Olshausen. The sparse manifold transform. *Advances in neural*
 598 *information processing systems*, 31, 2018.

599 SueYeon Chung and Larry F Abbott. Neural population geometry: An approach for understanding
 600 biological and artificial neural networks. *Current opinion in neurobiology*, 70:137–144, 2021.
 601

602 Valérie Costa, Thomas Fel, Ekdeep Singh Lubana, Bahareh Tolooshams, and Demba Ba. From
 603 flat to hierarchical: Extracting sparse representations with matching pursuit. *arXiv preprint*
 604 *arXiv:2506.03093*, 2025.

605 Hoagy Cunningham, Aidan Ewart, Logan Riggs, Robert Huben, and Lee Sharkey. Sparse autoen-
 606 coders find highly interpretable features in language models. *arXiv preprint arXiv:2309.08600*,
 607 2023.

608 Google Deepmind. Advanced version of gemini with deep think officially achieves gold-medal standard at the international mathematical olympiad. <https://deepmind.google/discover/blog/advanced-version-of-gemini-with-deep-think-officially-achieves-gold-medal-standard> 2025.

614 Howard Eichenbaum. Barlow versus hebb: When is it time to abandon the notion of feature detectors
 615 and adopt the cell assembly as the unit of cognition? *Neuroscience letters*, 680:88–93, 2018.
 616

617 Michael Elad. *Sparse and redundant representations: from theory to applications in signal and image*
 618 *processing*. Springer Science & Business Media, 2010.

619 Ronen Eldan and Yuanzhi Li. Tinystories: How small can language models be and still speak coherent
 620 english? *arXiv preprint arXiv:2305.07759*, 2023.

622 Nelson Elhage, Tristan Hume, Catherine Olsson, Nicholas Schiefer, Tom Henighan, Shauna Kravec,
 623 Zac Hatfield-Dodds, Robert Lasenby, Dawn Drain, Carol Chen, Roger Grosse, Sam McCandlish,
 624 Jared Kaplan, Dario Amodei, Martin Wattenberg, and Christopher Olah. Toy models of
 625 superposition. *Transformer Circuits Thread*, 2022.

626 Jeffrey L Elman. Finding structure in time. *Cognitive science*, 14(2):179–211, 1990.
 627

628 Thomas Fel, Binxu Wang, Michael A Lepori, Matthew Kowal, Andrew Lee, Randall Balestrieri,
 629 Sonia Joseph, Ekdeep S Lubana, Talia Konkle, Demba Ba, et al. Into the rabbit hull: From
 630 task-relevant concepts in dino to minkowski geometry. *arXiv preprint arXiv:2510.08638*, 2025.
 631

632 Leo Gao, Tom Dupré la Tour, Henk Tillman, Gabriel Goh, Rajan Troll, Alec Radford, Ilya
 633 Sutskever, Jan Leike, and Jeffrey Wu. Scaling and evaluating sparse autoencoders. *arXiv preprint*
 634 *arXiv:2406.04093*, 2024.

635 Atticus Geiger, Duligur Ibeling, Amir Zur, Maheep Chaudhary, Sonakshi Chauhan, Jing Huang,
 636 Aryaman Arora, Zhengxuan Wu, Noah Goodman, Christopher Potts, and Thomas Icard. Causal
 637 abstraction: A theoretical foundation for mechanistic interpretability. *Journal of Machine Learning*
 638 *Research*, 26(83):1–64, 2025. URL <http://jmlr.org/papers/v26/23-0058.html>.
 639

640 Antonios Georgiou, Tankut Can, Mikhail Katkov, and Misha Tsodyks. Using large language models
 641 to study human memory for meaningful narratives. *arXiv preprint arXiv:2311.04742*, 2023.

642 Edward Gibson et al. The dependency locality theory: A distance-based theory of linguistic complex-
 643 ity. *Image, language, brain*, 2000:95–126, 2000.
 644

645 Wes Gurnee, Emmanuel Ameisen, Isaac Kauvar, Julius Tarn, Adam Pearce, Chris Olah, and
 646 Joshua Batson. When models manipulate manifolds: The geometry of a counting task. *Trans-
 647 former Circuits Thread*, 2025. URL <https://transformer-circuits.pub/2025/linebreaks/index.html>.

648 John Hale. A probabilistic earley parser as a psycholinguistic model. In *Second meeting of the north*
 649 *american chapter of the association for computational linguistics*, 2001.

650

651 Michael Hanna and Aaron Mueller. Incremental sentence processing mechanisms in autoregressive
 652 transformer language models. *arXiv preprint arXiv:2412.05353*, 2024.

653

654 Irina Higgins, David Amos, David Pfau, Sebastien Racaniere, Loic Matthey, Danilo Rezende,
 655 and Alexander Lerchner. Towards a definition of disentangled representations. *arXiv preprint*
 656 *arXiv:1812.02230*, 2018.

657

658 Sai Sumedh R Hindupur, Ekdeep Singh Lubana, Thomas Fel, and Demba Ba. Projecting assumptions:
 659 The duality between sparse autoencoders and concept geometry. *arXiv preprint arXiv:2503.01822*,
 660 2025.

661

662 Zhuoqiao Hong, Haocheng Wang, Zaid Zada, Harshvardhan Gazula, David Turner, Bobbi Aubrey,
 663 Leonard Niekerken, Werner Doyle, Sasha Devore, Patricia Dugan, et al. Scale matters: Large lan-
 664 guage models with billions (rather than millions) of parameters better match neural representations
 665 of natural language. *bioRxiv*, 2024.

666

667 Eghbal Hosseini, Colton Casto, Noga Zaslavsky, Colin Conwell, Mark Richardson, and Evelina
 668 Fedorenko. Universality of representation in biological and artificial neural networks. *bioRxiv*,
 669 2024.

670

671 Eric Jonas and Konrad Paul Kording. Could a neuroscientist understand a microprocessor? *PLoS*
 672 *computational biology*, 13(1):e1005268, 2017.

673

674 Mikail Khona and Ila R Fiete. Attractor and integrator networks in the brain. *Nature Reviews*
 675 *Neuroscience*, 23(12):744–766, 2022.

676

Connor Kissane, Robert Krzyzanowski, Arthur Conmy, and Neel Nanda. Saes (usually) transfer between base and chat models, 2024.
<https://www.alignmentforum.org/posts/fmwk6qxrW8d4jvbd/saes-usually-transfer-between-base-and-chat-models>.

677

678 David Klindt, Charles O’Neill, Patrik Reizinger, Harald Maurer, and Nina Miolane. From su-
 679 perposition to sparse codes: interpretable representations in neural networks. *arXiv preprint*
 680 *arXiv:2503.01824*, 2025.

681

682 Mikhail V Koroteev. Bert: a review of applications in natural language processing and understanding.
 683 *arXiv preprint arXiv:2103.11943*, 2021.

684

685 Michael A Lepori, Jennifer Hu, Ishita Dasgupta, Roma Patel, Thomas Serre, and Ellie Pavlick. Is this
 686 just fantasy? language model representations reflect human judgments of event plausibility. *arXiv*
 687 *preprint arXiv:2507.12553*, 2025.

688

Roger Levy. Expectation-based syntactic comprehension. *Cognition*, 106(3):1126–1177, 2008.

689

690 Andrew Li, Xianle Feng, Siddhant Narang, Austin Peng, Tianle Cai, Raj Sanjay Shah, and Sashank
 691 Varma. Incremental comprehension of garden-path sentences by large language models: Semantic
 692 interpretation, syntactic re-analysis, and attention. *arXiv preprint arXiv:2405.16042*, 2024.

693

694 Tom Lieberum, Senthooran Rajamanoharan, Arthur Conmy, Lewis Smith, Nicolas Sonnerat, Vikrant
 695 Varma, János Kramár, Anca Dragan, Rohin Shah, and Neel Nanda. Gemma scope: Open sparse
 696 autoencoders everywhere all at once on gemma 2. *arXiv preprint arXiv:2408.05147*, 2024.

697

698 Scott Linderman, Matthew Johnson, Andrew Miller, Ryan Adams, David Blei, and Liam Paninski.
 699 Bayesian learning and inference in recurrent switching linear dynamical systems. In *Artificial
 intelligence and statistics*, pages 914–922. PMLR, 2017.

700

701 Jack Lindsey. Emergent introspective awareness in large language models. *Trans-
 former Circuits Thread*, 2025. URL <https://transformer-circuits.pub/2025/introspection/index.html>.

702 Jack Lindsey, Wes Gurnee, Emmanuel Ameisen, Brian Chen, Adam Pearce, Nicholas L. Turner,
 703 Craig Citro, David Abrahams, Shan Carter, Basil Hosmer, Jonathan Marcus, Michael Sklar, Adly
 704 Templeton, Trenton Bricken, Callum McDougall, Hoagy Cunningham, Thomas Henighan, Adam
 705 Jermyn, Andy Jones, Andrew Persic, Zhenyi Qi, T. Ben Thompson, Sam Zimmerman, Kelley
 706 Rivoire, Thomas Conerly, Chris Olah, and Joshua Batson. On the biology of a large language
 707 model. *Transformer Circuits Thread*, 2025. URL <https://transformer-circuits.pub/2025/attribution-graphs/biology.html>.

708

709 Alireza Makhzani and Brendan Frey. K-sparse autoencoders. *arXiv preprint arXiv:1312.5663*, 2013.

710

711 William Marslen-Wilson and Lorraine Komisarjevsky Tyler. The temporal structure of spoken
 712 language understanding. *Cognition*, 8(1):1–71, 1980. ISSN 0010-0277. doi: [https://doi.org/10.1016/0010-0277\(80\)90015-3](https://doi.org/10.1016/0010-0277(80)90015-3). URL <https://www.sciencedirect.com/science/article/pii/0010027780900153>.

713

714

715 Brian McElree, Stephani Foraker, and Lisbeth Dyer. Memory structures that subserve sentence
 716 comprehension. *Journal of memory and language*, 48(1):67–91, 2003.

717

718 L. McInnes and J. Healy. Umap: Uniform manifold approximation and projection for dimension
 719 reduction, 2025. URL <https://umap-learn.readthedocs.io/en/latest/>.

720

721 Leland McInnes, John Healy, and James Melville. Umap: Uniform manifold approximation and
 722 projection for dimension reduction. *arXiv preprint arXiv:1802.03426*, 2018.

723

724 Clara Meister, Tiago Pimentel, Patrick Haller, Lena Jäger, Ryan Cotterell, and Roger Levy. Revisiting
 725 the uniform information density hypothesis. *arXiv preprint arXiv:2109.11635*, 2021.

726

727 Beren Millidge, Mufeng Tang, Mahyar Osanlouy, Nicol S Harper, and Rafal Bogacz. Predictive
 728 coding networks for temporal prediction. *PLOS Computational Biology*, 20(4):e1011183, 2024.

729

730 Alexander Modell, Patrick Rubin-Delanchy, and Nick Whiteley. The origins of representation
 731 manifolds in large language models. *arXiv preprint arXiv:2505.18235*, 2025.

732

733 Monology. The pile: Uncopyrighted subset. <https://huggingface.co/datasets/monology/pile-uncopyrighted>, 2021. Based on the original Pile dataset by Gao et al.

734

735 Mark Muchane, Sean Richardson, Kiho Park, and Victor Veitch. Incorporating hierarchical semantics
 736 in sparse autoencoder architectures, 2025. URL <https://arxiv.org/abs/2506.01197>.

737

738 Aaron Mueller, Jannik Brinkmann, Millicent Li, Samuel Marks, Koyena Pal, Nikhil Prakash, Can
 739 Rager, Aruna Sankaranarayanan, Arnab Sen Sharma, Jiuding Sun, Eric Todd, David Bau, and
 740 Yonatan Belinkov. The quest for the right mediator: Surveying mechanistic interpretability through
 741 the lens of causal mediation analysis, 2025. URL <https://arxiv.org/abs/2408.01416>.

742

743 Sonia K Murthy, Rosie Zhao, Jennifer Hu, Sham Kakade, Markus Wulfmeier, Peng Qian, and Tomer
 744 Ullman. Inside you are many wolves: Using cognitive models to interpret value trade-offs in llms.
 745 *arXiv preprint arXiv:2506.20666*, 2025.

746

747 Ramon Nogueira, Chris C Rodgers, Randy M Bruno, and Stefano Fusi. The geometry of cortical
 748 representations of touch in rodents. *Nature Neuroscience*, 26(2):239–250, 2023.

749

750 Byung-Doh Oh and William Schuler. Transformer-based language model surprisal predicts human
 751 reading times best with about two billion training tokens. *arXiv preprint arXiv:2304.11389*, 2023.

752

753 Chris Olah. Distributed Representations: Composition & Superposition, 2023. <https://transformer-circuits.pub/2023/superposition-composition/index.html>.

754

755 Bruno Olshausen and Charles Cadieu. Learning invariant and variant components of time-varying
 756 natural images. *Journal of Vision*, 7(9):964–964, 2007.

Bruno A Olshausen and David J Field. Emergence of simple-cell receptive field properties by learning
 757 a sparse code for natural images. *Nature*, 381(6583):607–609, 1996.

756 Bruno A Olshausen and David J Field. Sparse coding with an overcomplete basis set: A strategy
 757 employed by v1? *Vision research*, 37(23):3311–3325, 1997.
 758

759 Senthooran Rajamanoharan, Tom Lieberum, Nicolas Sonnerat, Arthur Conmy, Vikrant Varma, János
 760 Kramár, and Neel Nanda. Jumping ahead: Improving reconstruction fidelity with jumprelu sparse
 761 autoencoders, 2024. URL <https://arxiv.org/abs/2407.14435>.

762 Shreya Saxena and John P Cunningham. Towards the neural population doctrine. *Current opinion in*
 763 *neurobiology*, 55:103–111, 2019.

764 Steffen Schneider, Jin Hwa Lee, and Mackenzie Weygandt Mathis. Learnable latent embeddings for
 765 joint behavioural and neural analysis. *Nature*, 617(7960):360–368, 2023.

766 Martin Schrimpf, Idan Asher Blank, Greta Tuckute, Carina Kauf, Eghbal A Hosseini, Nancy Kan-
 767 wisher, Joshua B Tenenbaum, and Evelina Fedorenko. The neural architecture of language:
 768 Integrative modeling converges on predictive processing. *Proceedings of the National Academy of*
 769 *Sciences*, 118(45):e2105646118, 2021.

770 SciPy. Dendrogram, hierarchical clustering, *scipy*, 2025. URL <https://docs.scipy.org/doc/scipy/reference/generated/scipy.cluster.hierarchy.dendrogram.html>.

771 H Sebastian Seung. How the brain keeps the eyes still. *Proceedings of the National Academy of*
 772 *Sciences*, 93(23):13339–13344, 1996.

773 Lee Sharkey, Bilal Chughtai, Joshua Batson, Jack Lindsey, Jeff Wu, Lucius Bushnaq, Nicholas
 774 Goldowsky-Dill, Stefan Heimersheim, Alejandro Ortega, Joseph Bloom, et al. Open problems in
 775 mechanistic interpretability. *arXiv preprint arXiv:2501.16496*, 2025.

776 Dong Shu, Xuansheng Wu, Haiyan Zhao, Daking Rai, Ziyu Yao, Ninghao Liu, and Mengnan Du. A
 777 survey on sparse autoencoders: Interpreting the internal mechanisms of large language models.
 778 *arXiv preprint arXiv:2503.05613*, 2025.

779 Hansem Sohn, Devika Narain, Nicolas Meirhaeghe, and Mehrdad Jazayeri. Bayesian computation
 780 through cortical latent dynamics. *Neuron*, 103(5):934–947, 2019.

781 Abiy Tasissa, Emmanouil Theodosis, Bahareh Tolooshams, and Demba Ba. Towards improving
 782 discriminative reconstruction via simultaneous dense and sparse coding, 2022. URL <https://arxiv.org/abs/2006.09534>.

783 Gemma Team, Thomas Mesnard, Cassidy Hardin, Robert Dadashi, Surya Bhupatiraju, Shreya Pathak,
 784 Laurent Sifre, Morgane Rivière, Mihir Sanjay Kale, Juliette Love, et al. Gemma: Open models
 785 based on gemini research and technology. *arXiv preprint arXiv:2403.08295*, 2024.

786 Ellen Thompson. The temporal structure of discourse: the syntax and semantics of temporal then.
 787 *Natural Language & Linguistic Theory*, 17(1):123–160, 1999.

788 Greta Tuckute, Aalok Sathe, Shashank Srikant, Maya Taliaferro, Mingye Wang, Martin Schrimpf,
 789 Kendrick Kay, and Evelina Fedorenko. Driving and suppressing the human language network
 790 using large language models. *Nature Human Behaviour*, 8(3):544–561, 2024.

791 Ulrike Von Luxburg. A tutorial on spectral clustering. *Statistics and computing*, 17(4):395–416,
 792 2007.

793 Laurenz Wiskott and Terrence J Sejnowski. Slow feature analysis: Unsupervised learning of
 794 invariances. *Neural computation*, 14(4):715–770, 2002.

795 Jeffrey M Zacks, Nicole K Speer, Khena M Swallow, Todd S Braver, and Jeremy R Reynolds. Event
 796 perception: a mind-brain perspective. *Psychological bulletin*, 133(2):273, 2007.

797 Weishun Zhong, Tankut Can, Antonis Georgiou, Ilya Shnayderman, Mikhail Katkov, and Misha
 798 Tsodyks. Random tree model of meaningful memory. *bioRxiv*, pages 2024–12, 2024.

799

810 A EXPERIMENTAL DETAILS
811812 Below, we discuss broad details of the experimental setup, training protocol, and measures used for
813 evaluations.
814815 A.1 TRAINING, HARVESTING, AND INFERENCE FOR SAEs
816817 **Compute.** All experiments were performed using 1 H100, both for extracting representations,
818 analyzing them, and training Temporal and standard SAEs. To save training costs, we precache
819 activations on a local server and train in an “online manner”, i.e., sampling a new batch of activations
820 every training iteration.
821822 **SAEs’ training.** All analyzed SAEs were trained from scratch on 1B precached activations from
823 the Pile-Uncopyrighted dataset (Monology, 2021). While we did try to use existing, off-the-shelf
824 SAEs, we found the results to be wildly inconsistent depending on where we borrowed the SAE
825 from. To enable a consistent and fair evaluation, we thus preferred to train all SAEs from scratch.
826 For Gemma models, activations were extracted from Layer 12; for Llama models, from layer 15 (all
827 0-indexed).
828829 **Training procedure.** We use Adam optimizer with standard hyperparameters. We initialize training
830 with a warmup of 200 steps to a learning rate of 10^{-3} , following from thereon to a minimum of
831 9×10^{-4} , i.e., the learning rate remains essentially constant throughout training.
832833 **Activations normalization.** Following best practices, we normalize activations to unit *expected*
834 norm (Bricken et al., 2023; Costa et al., 2025). This helps put different SAEs on the same training
835 scale, and especially for ReLU SAEs, makes training substantially easier by reducing the need for
836 tuning regularization strength. On a related note, we emphasize our ReLU SAEs have a slightly
837 higher L_0 , i.e., are less sparse, than almost all other SAEs analyzed in this work.
838839 A.2 EXPERIMENTS WITH LLM REPRESENTATIONS: AUTOCORRELATION AND U-STATISTIC
840841 We primarily analyze Gemma-2-2B and Llama-3.1-8B models in this paper. Specifically, we precache
842 10K activations for the analyzed datasets for our experiments on dynamics on language model
843 representations.
844845 A.2.1 AUTOCORRELATION
846847 We compute autocorrelation by selecting evenly spaced tokens across the sequence and measuring the
848 cosine similarity between each token and tokens at various lags in the past. Specifically, for tokens at
849 position t , we compute similarities to tokens at $t - w$ where lag w ranges from 5 to 20. This creates a
850 heatmap where rows represent lag offsets and columns represent token positions.
851852 For a stationary process, we expect the autocorrelation pattern to remain consistent across time—that
853 is, the relationship between a token and its historical context should be similar regardless of position
854 in the sequence. This would manifest as similar autocorrelation patterns repeating horizontally across
855 token positions. In contrast, for a non-stationary process where representations evolve over time, we
856 expect the autocorrelation patterns to vary systematically across positions, with columns showing
857 different temporal dependency structures as the sequence progresses.
858859 A.2.2 U-STATISTIC
860861 We measure the effective dimensionality of LLM representations using a U-statistic (an unbiased
862 estimator) based on pairwise cosine similarities. Let $\mathbf{X}_t = [\mathbf{x}_t^{(1)}, \dots, \mathbf{x}_t^{(M)}]$ be M samples of
863 normalized activations at time t (from different timeseries). $\mathbf{G}_t = \mathbf{X}_t^T \mathbf{X}_t$ is the Gram matrix. Using
864 the above, a U-statistic of intrinsic dimension which we employ is defined below:
865

866
$$\text{U-stat}(t) = \frac{M^2 - M}{\|\mathbf{G}_t\|_F^2 - M} \quad (5)$$

867

864 where $\|\mathbf{G}_t\|_F^2$ is the squared Frobenius norm of the Gram matrix. This quantity estimates the effective
 865 rank $1/\text{tr}(\mathbf{C}_t^2)$, where $\mathbf{C}_t = \mathbb{E}[\mathbf{x}_t \mathbf{x}_t^T]$ is the second moment matrix and \mathbf{x}_t is the activation vector at
 866 time t . Under stationarity, U-stat remains constant. When representations evolve over time, U-stat
 867 increases systematically as more orthogonal directions become active.
 868

869 A.2.3 PROJECTION ANALYSIS

870 This analysis quantifies how much of the representation $\mathbf{x}_t \in \mathbb{R}^D$ at token position t can be
 871 reconstructed from its preceding context $\{\mathbf{x}_i\}_{i \in W}$, where the context window $W = [t-w, t-1]$
 872 contains the previous w tokens.
 873

874 For a population of B sentence samples, we compute the projection of each target representation onto
 875 the subspace spanned by its context. Let $\mathbf{X}_W^{(b)} = [\mathbf{x}_{t-w}^{(b)}, \dots, \mathbf{x}_{t-1}^{(b)}]^T \in \mathbb{R}^{w \times D}$ denote the matrix of
 876 context representations for sample b , where each row is a context vector. Prior to projection, we
 877 center all representations by subtracting the mean computed over the target positions across samples.
 878 The projection of $\mathbf{x}_t^{(b)}$ onto the span of the context is:
 879

$$880 \mathbf{c}_{t,w}^{(b)} = \text{span}(\mathbf{X}_W^{(b)}) \cdot \mathbf{x}_t^{(b)} \quad (6)$$

881 where $\text{span}(\mathbf{X}_W^{(b)})$ is the row space of $\mathbf{X}_W^{(b)}$, that can be computed via SVD. The variance explained
 882 by the context is:
 883

$$884 \text{expvar}(t, w) = \frac{\sum_{d=1}^D \text{var}(\mathbf{c}_{t,w,d})}{\sum_{d=1}^D \text{var}(\mathbf{x}_{t,d})} \quad (7)$$

885 where $\text{var}(\mathbf{c}_{t,w,d})$ and $\text{var}(\mathbf{x}_{t,d})$ denote the variance of the d -th dimension across the B samples. This
 886 ratio measures the fraction of total representational variance that can be linearly reconstructed from
 887 the preceding context. Values approaching 1 indicate high predictability from context; values near 0
 888 indicate representations largely orthogonal to their context subspace.
 889

890 A.2.4 SURROGATE

891 For the U-statistic and Autocorrelation metrics, we compare LLM activations to surrogate distributions
 892 that preserve certain statistical properties while removing temporal structure. We operate on
 893 representation vectors $\mathbf{X} \in \mathbb{R}^{B \times T \times d}$, where B denotes batch size, T denotes sequence length, and d
 894 denotes the model dimension.
 895

896 **U-statistic surrogate (Fig. 2 a, e):** For each sequence $i \in \{1, \dots, B\}$, we construct the surrogate
 897 $\tilde{\mathbf{X}}_i$ by applying a random permutation $\pi_i : \{1, \dots, T\} \rightarrow \{1, \dots, T\}$ to the temporal positions:
 898

$$900 \tilde{\mathbf{X}}_{i,t,:} = \mathbf{X}_{i,\pi_i(t),:} \quad \forall t \in \{1, \dots, T\}$$

901 This preserves the marginal distribution of activations within each sequence while destroying temporal
 902 dependencies.
 903

904 **Autocorrelation surrogate (Fig. 2 c, g):** Given the similarity matrix $\mathbf{S} \in \mathbb{R}^{T \times T}$ where $S_{ij} =$
 905 $\text{sim}(\mathbf{X}_{\cdot, i, \cdot}, \mathbf{X}_{\cdot, j, \cdot})$, we construct the surrogate similarity matrix $\tilde{\mathbf{S}}$ by replacing each diagonal with its
 906 mean:
 907

$$908 \tilde{S}_{ij} = \bar{S}_k \quad \text{where } k = |i - j|, \quad \bar{S}_k = \frac{1}{T-k} \sum_{t=1}^{T-k} S_{t,t+k}$$

910 This preserves the average correlation structure at each lag while removing position-specific temporal
 911 patterns.
 912

913 A.3 ANALYSIS OF STORIES: UMAP PROJECTIONS, DENDROGRAMS, SIMILARITY HEATMAPS, 914 AND EVENT BOUNDARY DETECTION

915 **Data.** For experiments in Sec. 6.1, we used either stories sampled from the TinyStories dataset (El-
 916 dan and Li, 2023) (for Fig. 4, 5) or synthetically sampled stories using GPT-5 (for Fig. 6). For the
 917 former, we sampled stories that were up to a 100 tokens long, allowing us to characterize compute

918 UMAP for both Gemma and Llama models without hitting memory bottlenecks. For consistency, we
 919 then used these stories for all experiments. For synthetically defined stories, we defined 3 in-context
 920 examples (one of which is shown in Fig. 6a) and prompted GPT-5 to produce 50 stories with similar
 921 such suddenly changing events. Stories varied in size, but were generally less than 150 tokens long.
 922 Following Georgiou et al. (2023), the generated stories’ events were labeled using GPT-5. These
 923 labels served as the ‘ground truth’ event boundaries. We qualitatively analyzed and confirmed the
 924 event boundaries overlap with our intuitively expected event structure in the stories.

925
 926 **Analysis Details.** To isolate the low-dimensional geometry of latent codes, we perform a 3D UMAP
 927 analysis using the open-source package (McInnes and Healy, 2025). Similarly, Dendograms in
 928 Fig. 4 are computed using the hierarchical clustering package in SciPy (SciPy (2025)). Branches are
 929 colored based on proximity, which in our case was between 0–1 (since we use cosine similarity as
 930 the clustering measure). We put a proximity bound of 0.2 as the distance under which branches are
 931 colored within a cluster. For Fourier Analysis, as mentioned in Sec. 6.1, we performed a Fourier
 932 transform of the model activations for a given story, isolated its lower frequency components—defined
 933 as $0.1 \times$ the Nyquist rate, which turns out to possess $\sim 50\%$ of the signal norm—and compute the
 934 cosine similarity kernel of the low / high frequency spectra. These kernels are then compared to
 935 kernels of latent codes derived using Temporal or standard SAEs.

936 A.4 ANALYZING GARDEN PATH SENTENCES: DENDROGRAMS AND PHRASE SIMILARITY

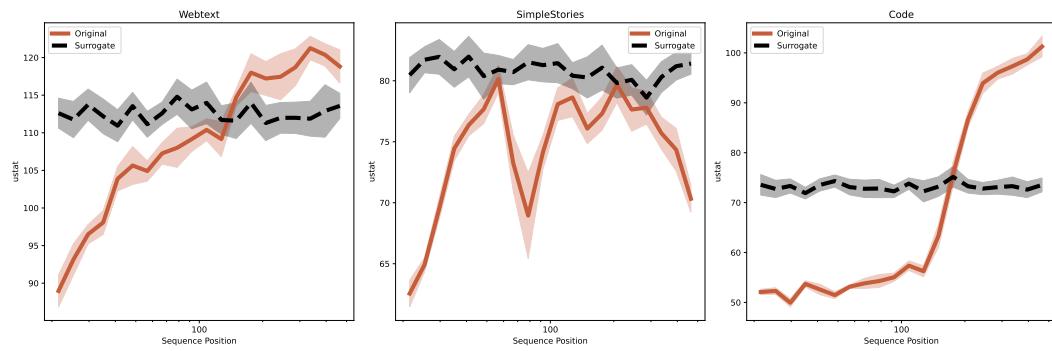
937
 938 **Data.** We used GPT-5 to sample 50 garden path sentences and control variants thereof (1 corre-
 939 sponding to each sentence). Sentences were 20–40 tokens long and had a structure such that the
 940 observation of the object phrase resolved ambiguity. The precise ambiguity structure was varied
 941 in type, i.e., the ambiguity could be resolved by altering the subject (e.g., old → sailors in The
 942 old man the road) or by other mechanisms such as adding punctuations to elicit a pause (e.g.,
 943 adding a comma, such as The old train the young fight → The old train, the
 944 young fight). Control variants had a mixture of such resolutions applied.

945
 946 **Analysis.** For assessing whether tokens in verb phrase relate more with the subject phrase, as
 947 would be expect by the typical parse in our used garden path sentences, or with the object phrase, as
 948 would be necessary for the correct parse, we computed Dendograms using the hierarchical clustering
 949 package in SciPy (SciPy, 2025). The similarity between subject phrase (SP), verb phrase (V), and
 950 object phrase (OP) was computed by taking the set of tokens T_P that belong to a phrase $P \in \{\text{SP}, \text{V},$
 951 $\text{OP}\}$, and computing the average latent code: $\bar{c}(P) = \frac{1}{|P|} \sum_{t \in T_P} c_t$, where c_t denotes the latent code
 952 extracted for token t using either SAEs or the predictive or novel component of Temporal SAEs. This
 953 protocol is similar to popularly used strategies for computing sentence embeddings (see, e.g., work
 954 using BERT embeddings (Koroteev, 2021)). Cosine similarity between these phrase-averaged garden
 955 path / control sentences yields the tables shown in Fig. 7.

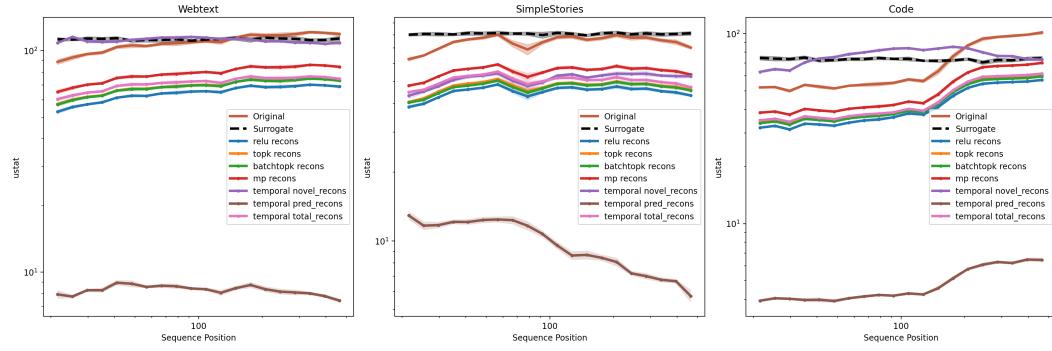
956
 957
 958
 959
 960
 961
 962
 963
 964
 965
 966
 967
 968
 969
 970
 971

972 B FURTHER EVALUATION RESULTS

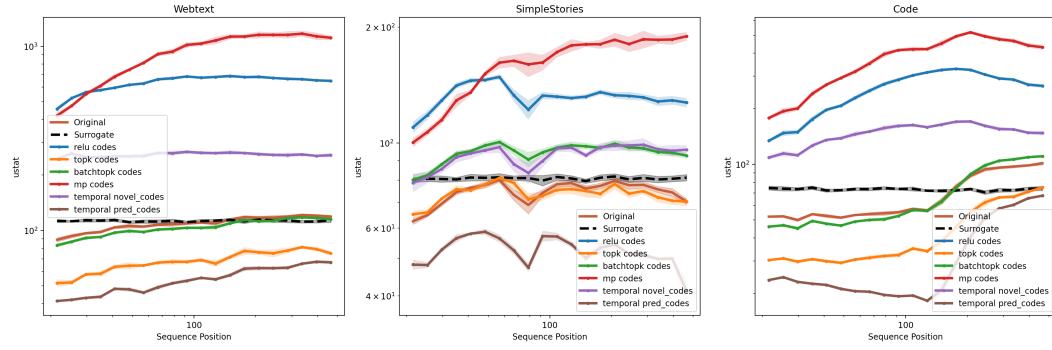
974 B.1 U-STATISTIC ACROSS DOMAINS FOR GEMMA-2-2B



987 Figure 8: U-Statistic across domains for LLM activations, surrogate



990 Figure 9: U-Statistic across domains for SAE reconstructions



1000 Figure 10: U-Statistic across domains for SAE codes

1015 B.2 AUTOCORRELATION ACROSS DOMAINS FOR GEMMA-2-2B

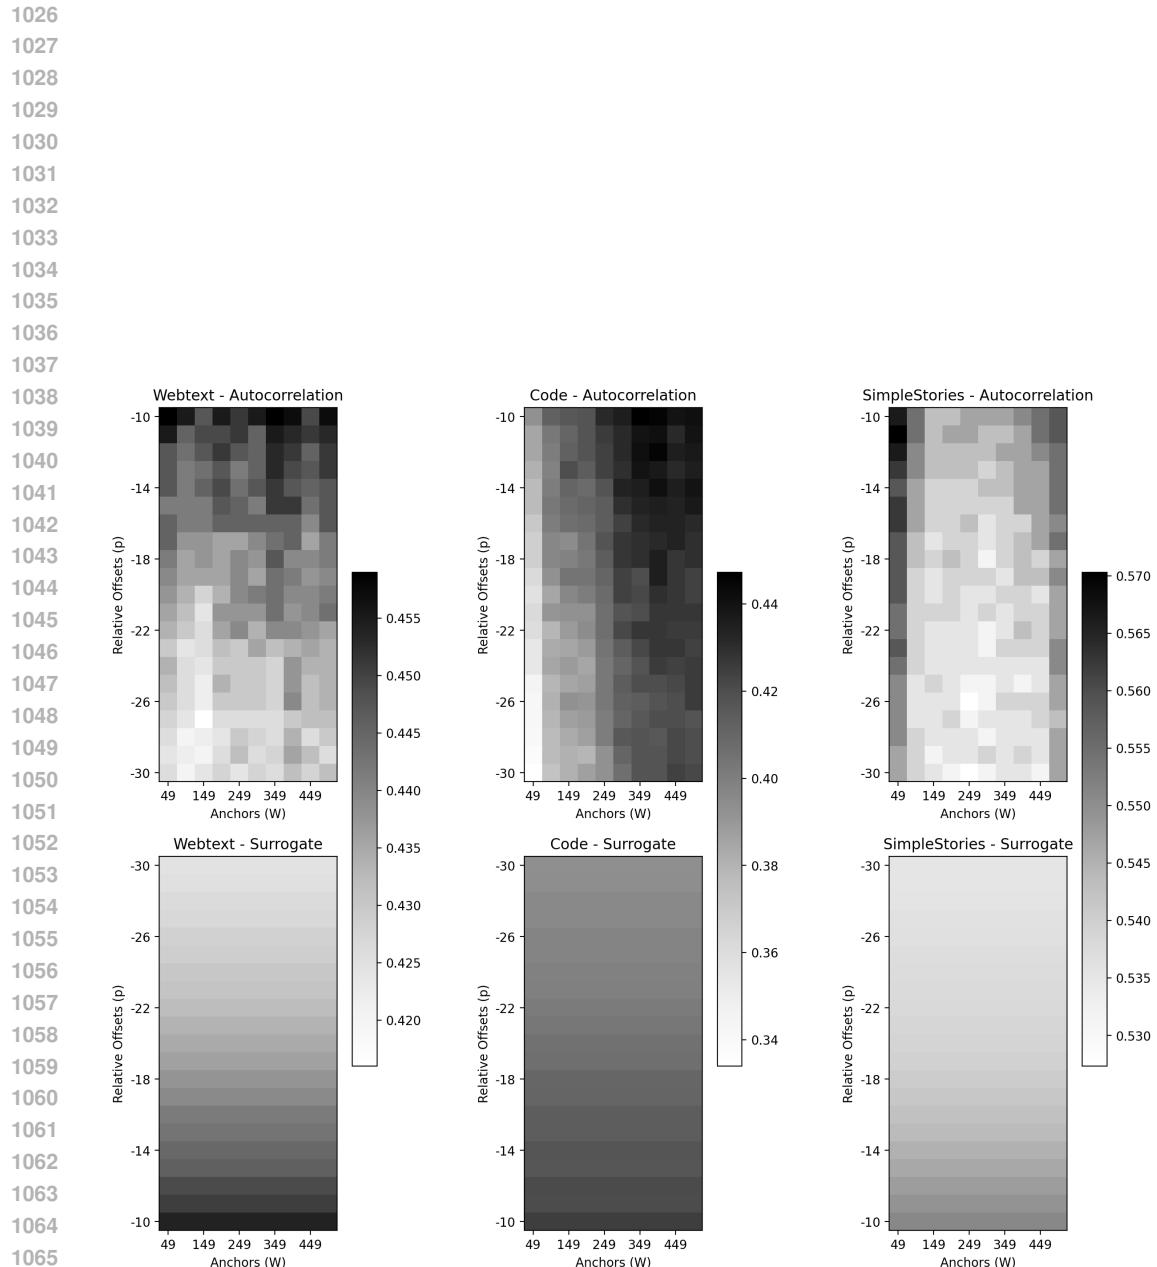


Figure 11: Autocorrelation of language model activations and a stationary surrogate across webtext, simple stories and code domains.

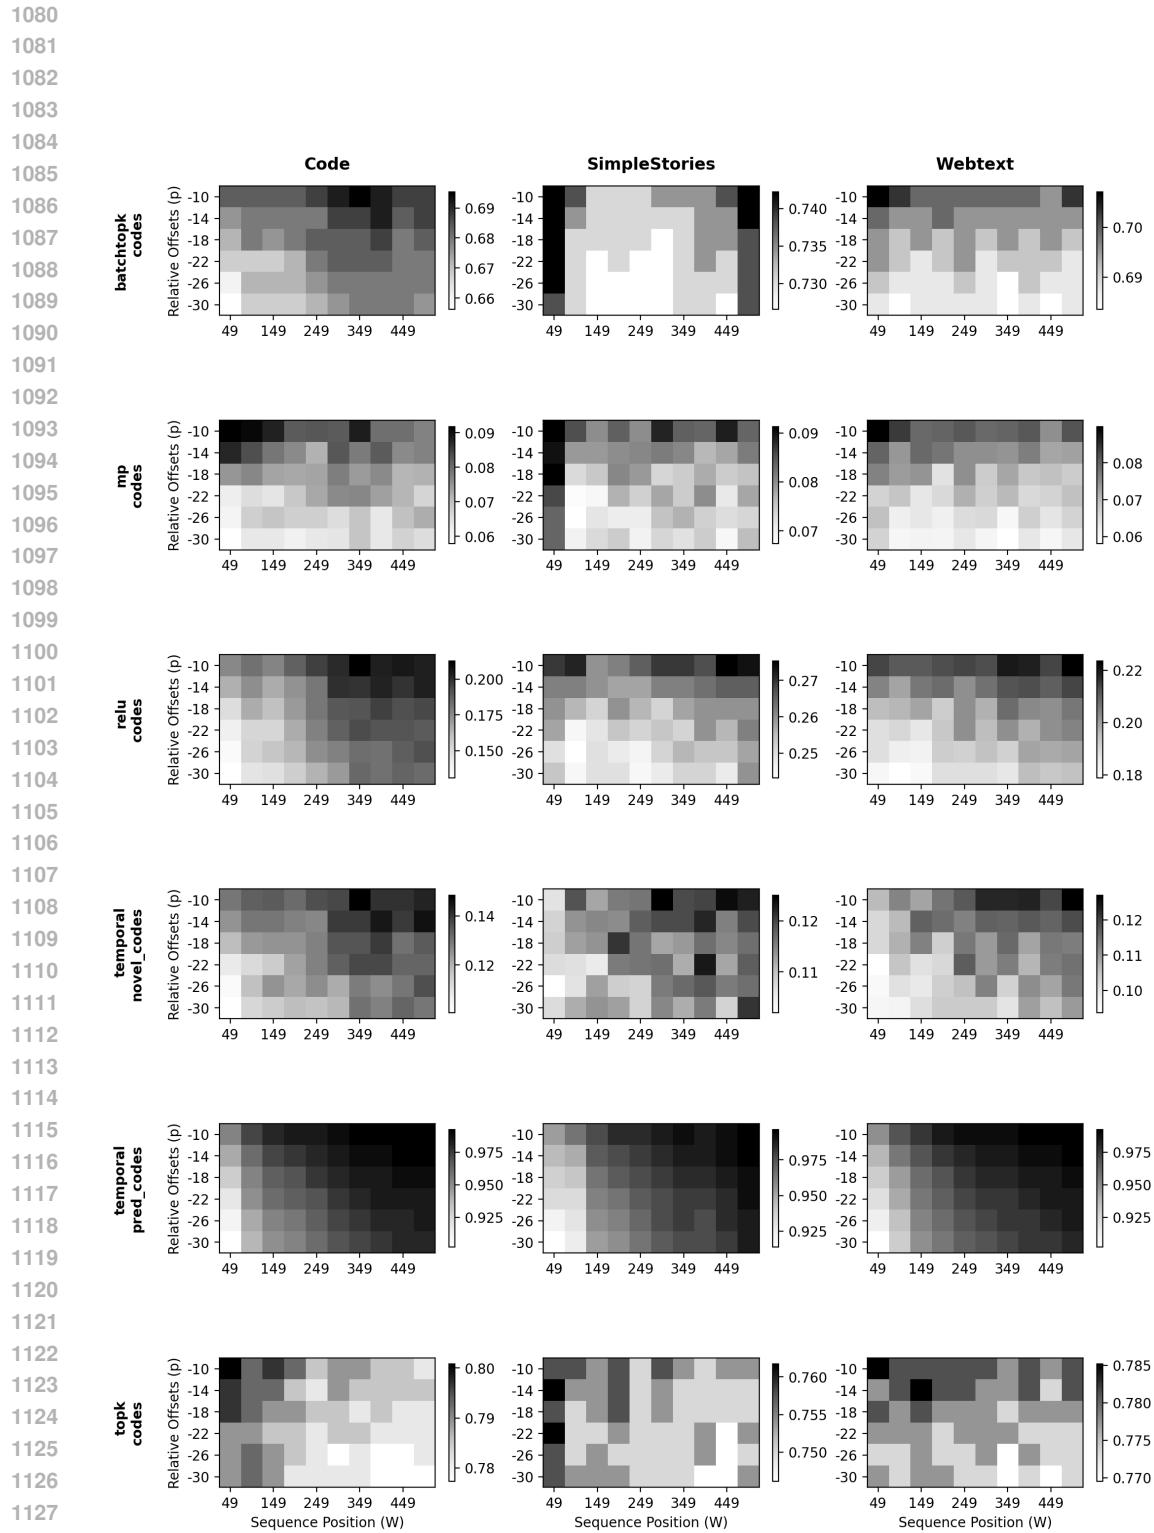
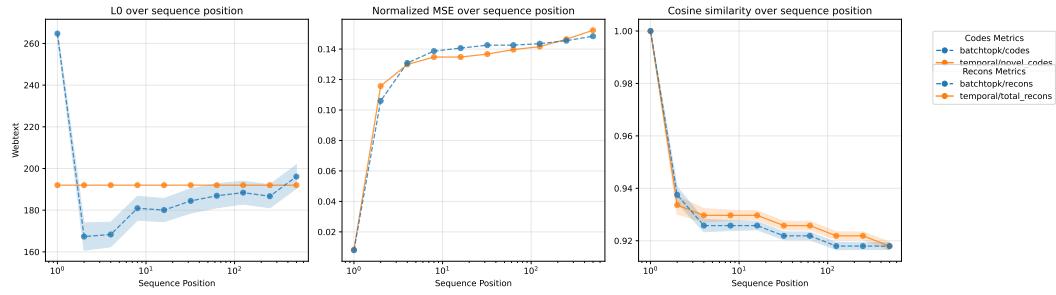
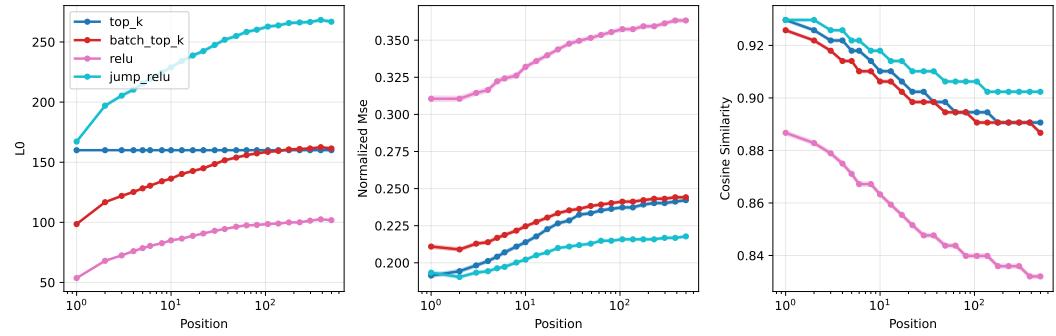
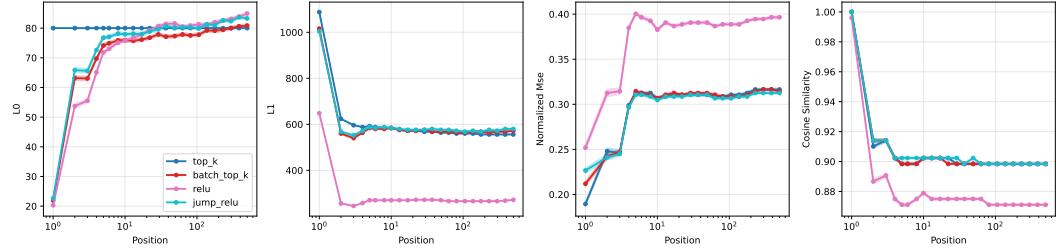


Figure 12: Autocorrelation of SAE latent activations across webtext, simple stories and code domains.

1134
1135 C RECONSTRUCTION PERFORMANCE OF SAEs
11361145
1146 Figure 13: Attn SAE vs BatchTopK1147
1148
1149
1150
1151
1152
1153
1154
1155
1156
1157 Figure 14: Reconstruction performance of Llama-3.1-8B resid post Layer 16 (1/2 forward pass) SAEs on webtext.1158
1159
1160
1161
1162
1163
1164
1165
1166
1167
1168
1169
1170
1171
1172
1173
1174
1175
1176
1177
1178
1179
1180
1181
1182
1183
1184
1185
1186
1187 Figure 15: Reconstruction performance of Gemma-2-2B resid post Layer 12 (1/2 forward pass) SAEs on webtext.

D CONSTANT SPARSITY ARGUMENT: RANK AND U-STATISTIC

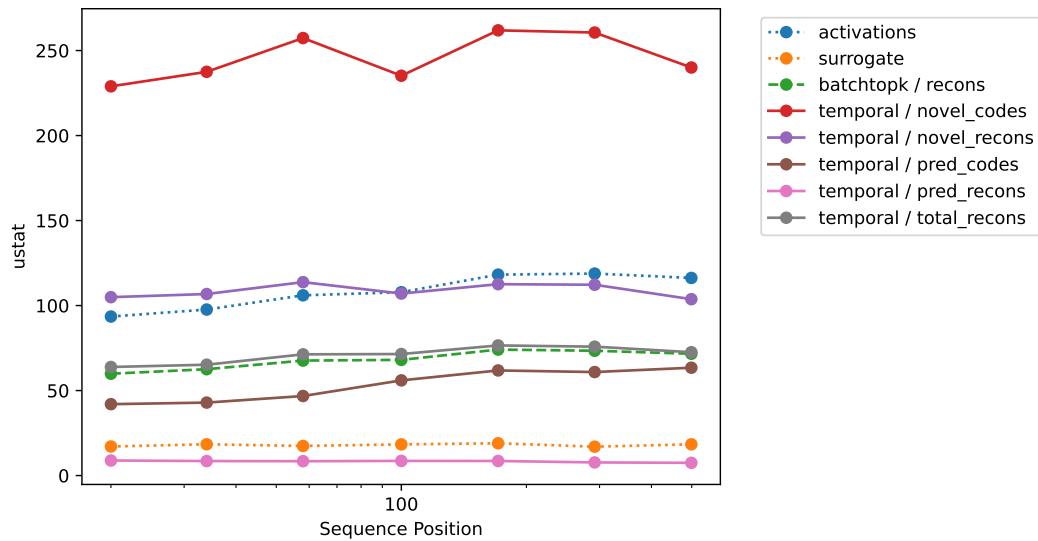


Figure 16: Gemma-2-2b: LLM activations vs Attention SAEs vs BatchTopK

D.1 LLAMA-3.1-8B

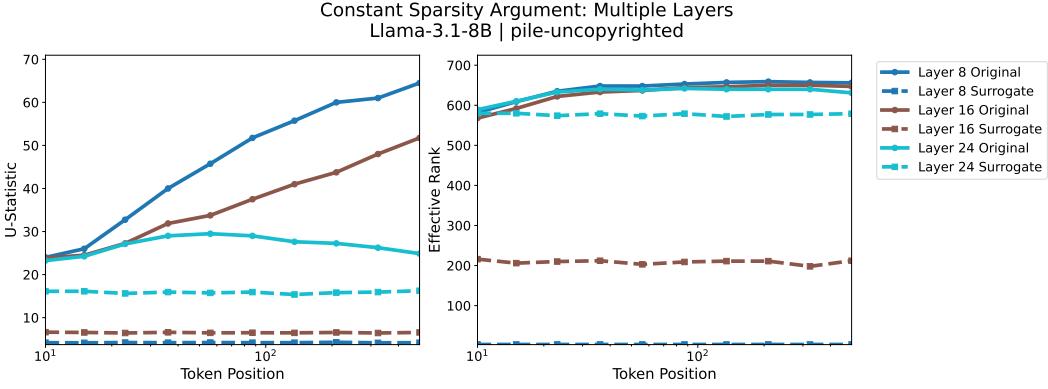


Figure 17: Rank and U-statistic, Llama-3.1-8B llm activations and surrogate, Multiple Layers

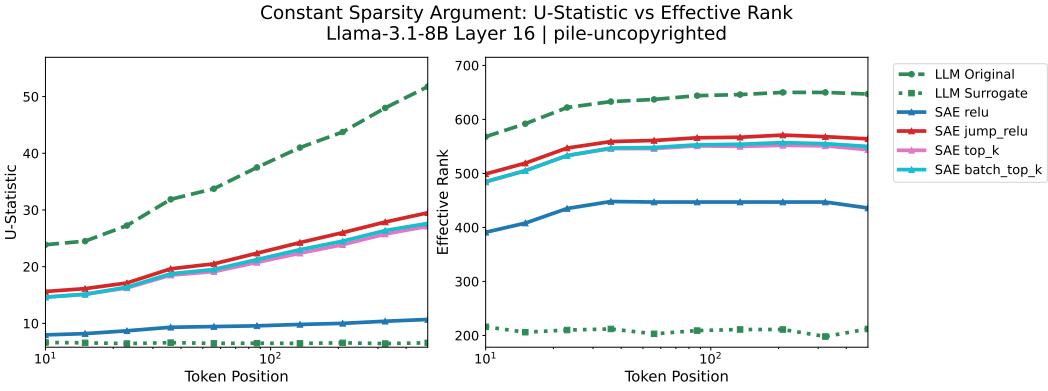


Figure 18: Rank vs U-statistic, Llama-3.1-8B, webtext

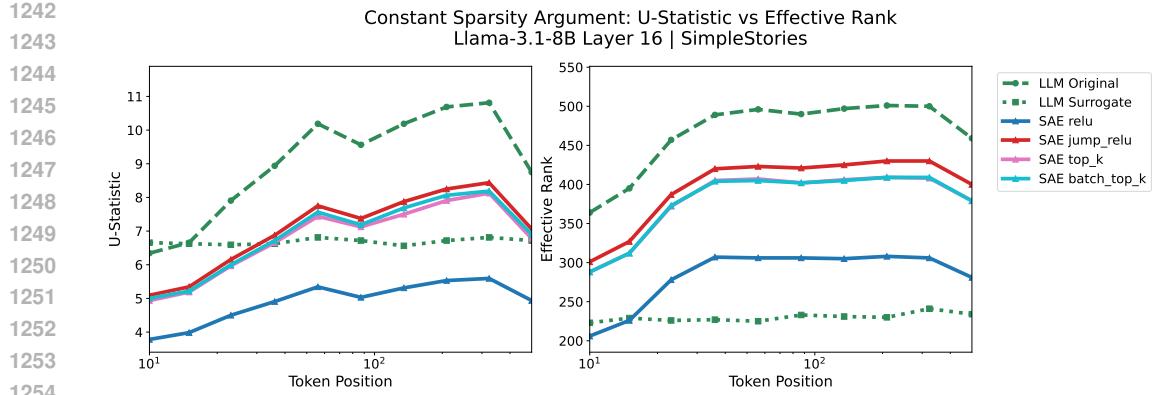


Figure 19: Rank vs U-statistic, llama-3.1-8B, simple stories

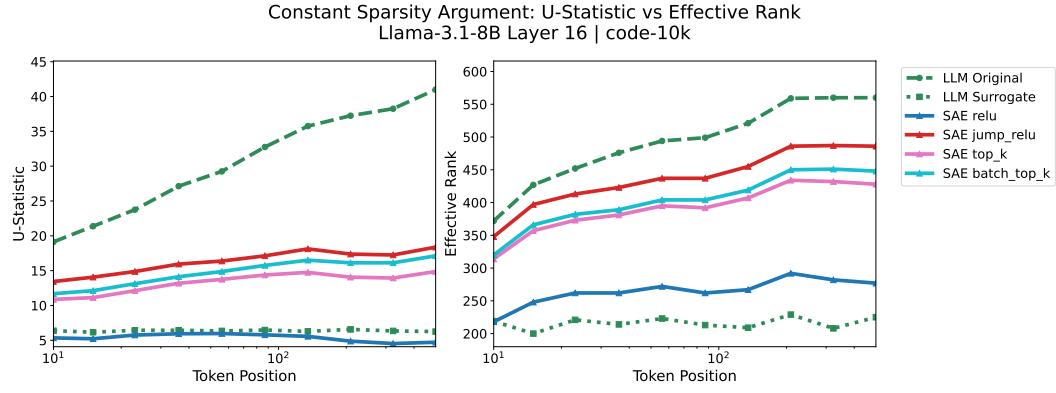


Figure 20: Rank vs U-statistic, llama-3.1-8B, Code

D.2 GEMMA-2-2B

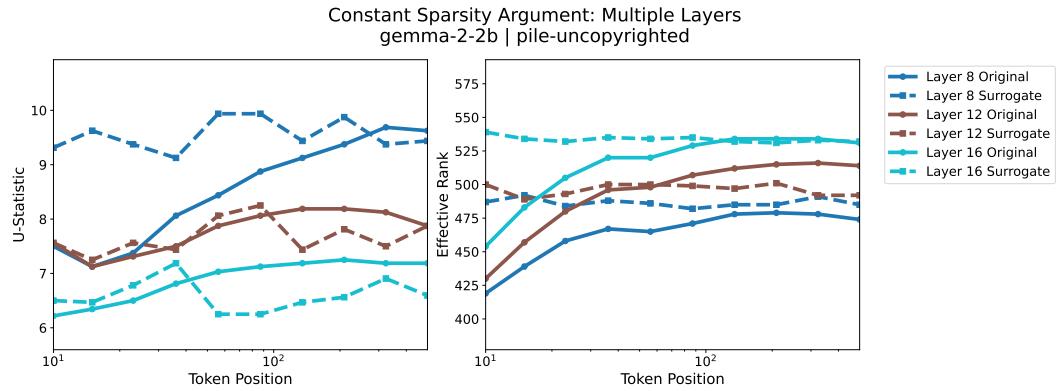


Figure 21: Rank and U-statistic, Gemma-2-2b ILM activations and surrogate, Multiple Layers

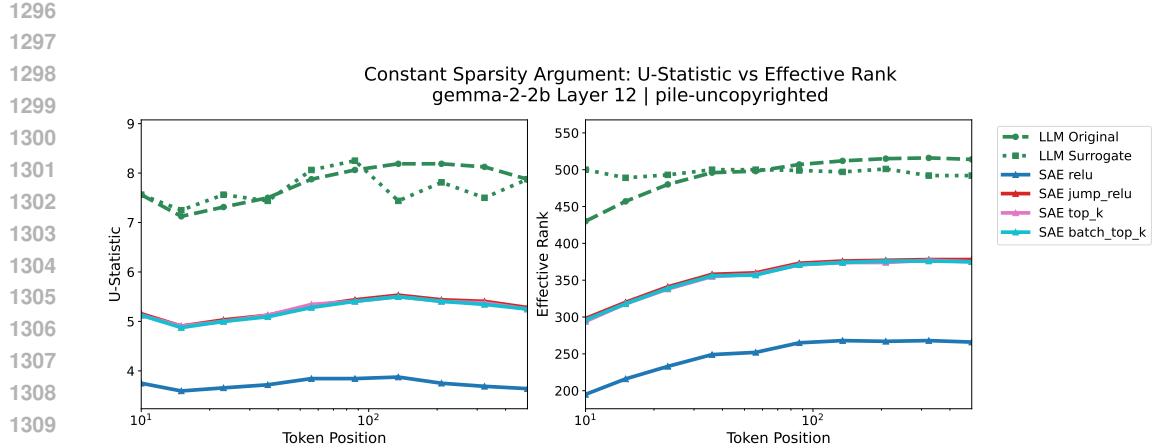


Figure 22: Rank vs U-statistic, Gemma-2-2b, webtext

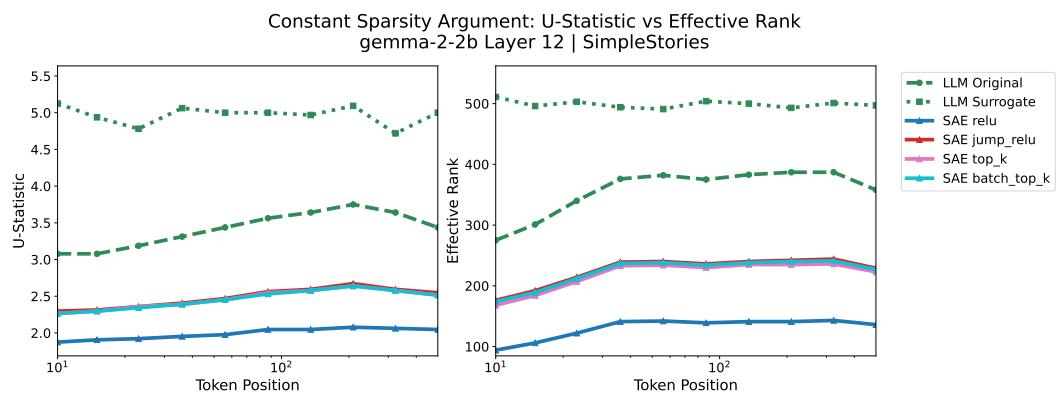


Figure 23: Rank vs U-statistic, Gemma-2-2b, simplestories

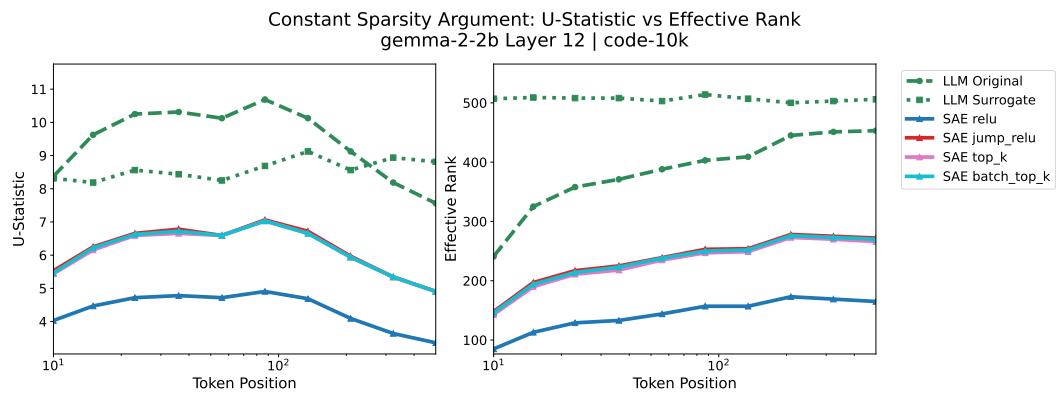
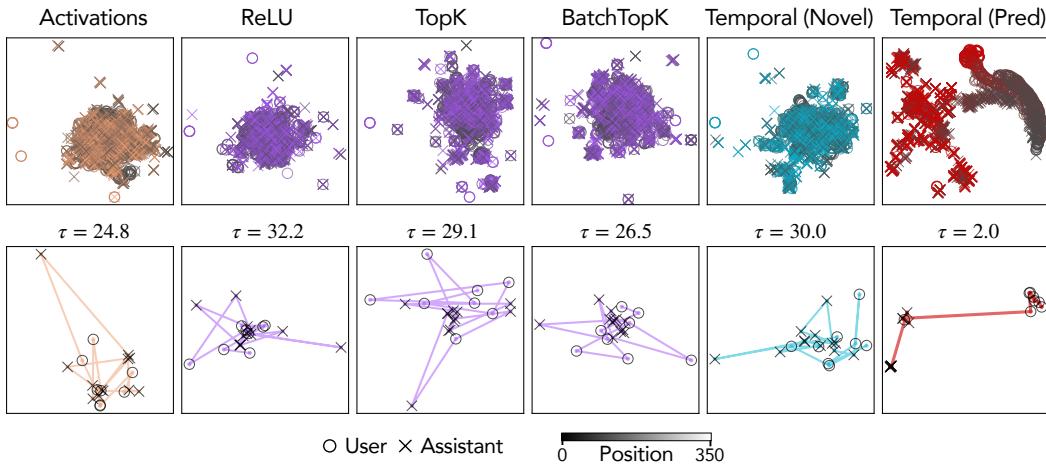


Figure 24: Rank vs U-statistic, Gemma-2-2b, code

1350
1351 E FURTHER INVESTIGATION OF TEMPORAL SAEs
1352

1353 In the results above, we showed across three domains that Temporal SAEs elicits intricate in-context
1354 geometry inline with our expectations based on experiments with LM activations in Sec. 3. Next, to
1355 make a case for following such more structured approaches for interpretability, we add to the results
1356 above to show that Temporal SAEs (i) offers a new way of interpreting temporally structured domains
1357 (e.g., user–model chats) via the context-dependent predictive component, and (ii) information that
1358 can be identified via SAEs continues to remain available within the novel component of Temporal
1359 SAEs.

1360 E.1 GEOMETRY OF LATENT CODES IN AN OOD DIALOGUE DOMAIN
1361

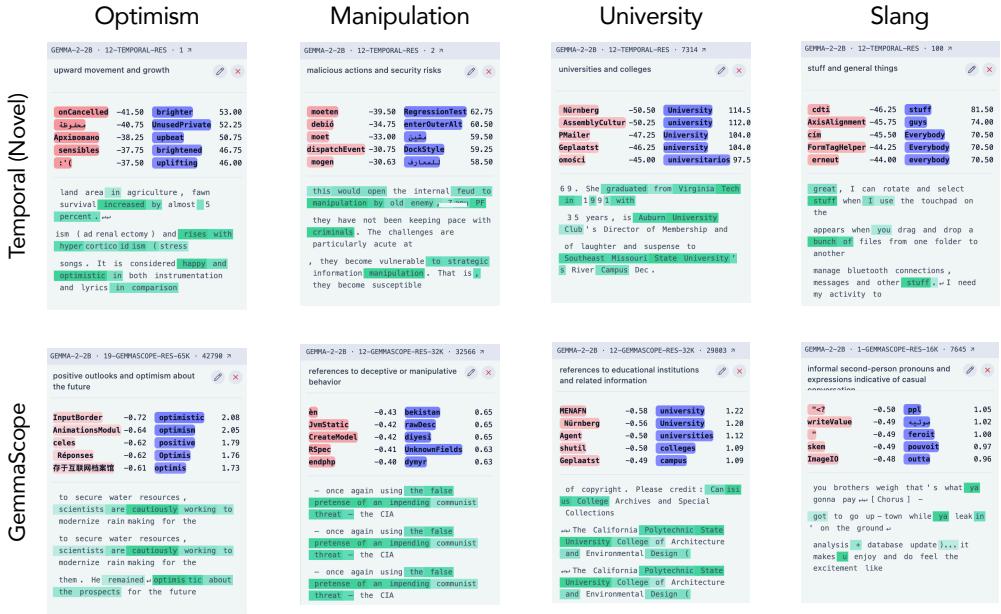
1377 **Figure 25: Population and trajectory geometry in OOD dialogue.** UMAPs of raw activations and
1378 SAE codes on Ultrachat (first 350 tokens per conversation). **Top:** population embeddings colored
1379 by sequence position (gray→color) and marked by role (user=○, assistant=×); only the *predictive*
1380 code cleanly separates speaker roles while preserving a smooth temporal gradient. **Bottom:** a single
1381 conversation overlaid on the same UMAP; the predictive code follows a smooth near-geodesic with
1382 low tortuosity τ , in contrast to jagged paths from standard SAEs/novel codes.

1383 We analyze a multi-turn chat domain to test whether Temporal SAEs continues to reveal slow-moving
1384 structure. We emphasize the activations in this experiment are sampled from a Gemma-2-2B-IT
1385 model, although all analyzed interpretability protocols were trained on the base (non-instruction
1386 tuned) model. Thus, in a sense, the experiments in this section also gauge the ability of different
1387 protocols to generalize out-of-distribution (OOD)—a property SAEs have been argued to partially
1388 possess when transferring between base and the corresponding instruction-tuned models (Kissane
1389 et al., 2024). Specifically, we use the Ultrachat dataset, where we take the first 350 tokens from 1,000
1390 conversations (mostly single-turn within this window, with up to four turns). For each token, we
1391 compute a shared 2D UMAP embedding from (i) raw model activations and (ii) latent codes from
1392 ReLU, TopK, and BatchTopK SAEs, as well as the novel and predictive components of the Temporal
1393 SAE. To visualize population structure clearly, we display only 100 conversations at a time to avoid
1394 clutter (Fig. 25), though fitting is applied to all 1,000.

1395 At the population level, the predictive component of Temporal SAEs is the only representation that
1396 cleanly organizes the data along *two* interpretable axes: (1) a *speaker-role separation*, with user and
1397 assistant tokens forming distinct yet smoothly adjoining manifolds, and (2) a *temporal gradient*, with
1398 positions flowing continuously from early to late tokens along the manifold. By contrast, standard
1399 SAEs and the novel code show diffuse clouds with weaker role separability and no coherent positional
1400 gradient (Fig. 10, top row). To assess within-sequence geometry, we overlay a single conversation’s
1401 trajectory onto the population embedding (Fig. 10, bottom row). Predictive codes trace a *smooth*
1402 *trajectory* that remains locally coherent within each role segment and transitions gracefully at speaker
1403 switches. In contrast, standard SAEs and the novel code produce jagged, zig–zagging paths. We again
1404 quantify this with *tortuosity* τ (Bullitt et al., 2003) (lower is straighter): predictive codes achieve
1405 $\tau \approx 2.0$, whereas SAEs and the novel code yield $\tau \approx 25\text{--}30$, indicating substantially more local

1404 turning (Fig. 10, panel headers). This straightening mirrors the story-domain results (Sec. 6.1), now
 1405 in an OOD conversational chat regime.
 1406

1407 E.2 AUTOMATIC INTERPRETABILITY OF NOVEL CODES

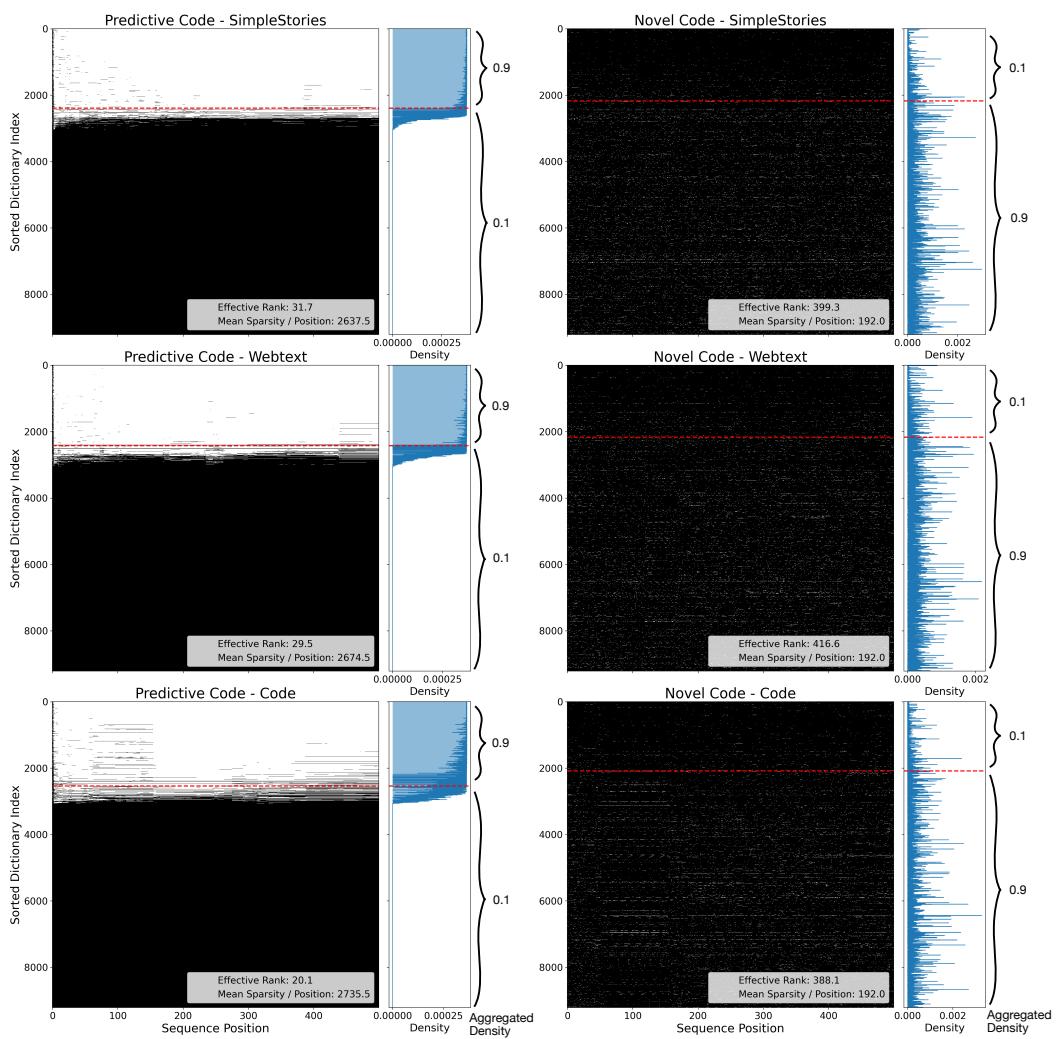


1429 **Figure 26: Interpretability of novel codes.** We qualitatively compare the interpretability of novel
 1430 codes extracted using Temporal SAEs with latent codes extracted using GemmaScope SAEs across
 1431 four diverse categories. Example feature cards are shown across four categories, produced with
 1432 the same activation-triggered examples plus direct logit attribution workflow used for standard
 1433 SAEs (Lieberum et al., 2024).
 1434
 1435

1436 A central benefit of SAEs is that their latents admit automatic interpretability pipelines (“autointerp”)
 1437 that attach human-readable semantics to features via activation-triggered examples, token/phrase
 1438 saliency, and lightweight dashboards (e.g., GemmaScope / Neuronpedia). Our Temporal SAE
 1439 adds a new, temporally aware predictive axis, but crucially *it does not sacrifice* these established
 1440 interpretability affordances: the novel code remains compatible with standard SAE autointerp
 1441 workflows. In other words, Temporal SAEs augment the toolbox rather than replacing it.

1442 To support this point, we apply a conventional autointerp loop to the novel codes: (i) collect high-
 1443 activation snippets and nearest-neighbor contexts for each latent; (ii) compute direct logit attribution
 1444 to inspect which vocabulary items a latent tends to support or suppress; and (iii) render compact,
 1445 browsable “feature cards” that aggregate examples and scores (as in GemmaScope / Neuronpedia).
 1446 While we do not perform an exhaustive causal-intervention study here, nothing in the pipeline is
 1447 specific to standard SAEs; the same probes and dashboards operate unchanged on the novel codes,
 1448 and we leave a systematic intervention suite to future work. Qualitative inspection confirms that the
 1449 novel codes recover the kinds of monosemantic, language-level features typically reported for SAEs.
 1450 Figure 11 (p. 13) shows representative examples across diverse categories—Optimism, Manipulation,
 1451 University, and Slang—where the novel latents consistently activate on intuitively relevant cue phrases
 1452 and contexts, and where their attribution profiles align with the expected lexical sets.

1453 Overall, when standard SAEs produce readable feature cards, the Temporal SAE’s novel stream does
 1454 as well. This complements our earlier results: the predictive stream captures slow, contextual structure
 1455 (events, roles, long-range constraints), whereas the novel stream concentrates the fast, stimulus-driven
 1456 information that matches inferences possible via existing SAEs, including autointerp pipelines.
 1457

1458 F EMERGENT SEPARATION OF PREDICTIVE AND NOVEL PARTS
1459

1493 **Figure 27: Separation of Predictive and Novel Dictionaries.** Predictive and novel codes for
1494 sequences drawn from three datasets—SimpleStories (top), Webtext (middle), and Code (bottom)—
1495 are binarized to zero (black) and non-zero (white). The dictionary elements are sorted by the sum
1496 of non-zero activations in the predictive code across the sequence. The ordering obtained from the
1497 predictive code is also applied to the novel code. The red dashed line marks the separation of 90% of
1498 non-zero activation counts; only 10% are overlapping into the other subset. Additionally, the effective
1499 rank of the predictive codes is two orders of magnitudes lower than the average number of non-zero
1500 dictionary elements (Mean L0).

1501 The Temporal SAEs studied in the main paper share a dictionary for both predictive and novel codes.
1502 We now investigate if there is any shared structure in the two dictionaries. In particular, given the
1503 predictive and novel codes play a different role, it is plausible the SAE learns to approximately split
1504 the dictionary into two parts: one responsible for computing the predictive code and the other for
1505 novel one. Specifically, we investigate a Temporal SAE trained on Gemma-2-2B Layer 12 residual
1506 stream activations. Results are shown in Fig. 27 and show a separation of dictionary elements for
1507 sequences drawn from three datasets: SimpleStories, Webtext, and Code. Specifically, we find $\sim 2K$
1508 dictionary elements participate in defining the predictive code, while the remaining $\sim 8K$ primarily
1509 participate in defining the novel code. However, the absolute count is merely indicative of the
1510 L0 sparsity of a code, which may not reflect how many *directions* are actually used to define the
1511 code—estimating latter requires computing the rank of the code (in fact, we note that rank-sparsity is
a well-known alternative to L0 sparsity in dictionary learning literature (Elad, 2010)). As we show

1512 in Fig. 27, the effective rank of the predictive codes ($\sim 20\text{--}30$) is substantially lower than both the
1513 effective rank ($\sim 390\text{--}410$) and absolute L0 sparsity of novel codes.

1514 Overall, the posthoc analysis above shows that predictive and novel codes largely use separate
1515 subsets of dictionary elements. This emergent disentanglement of the two components motivates
1516 one to preemptively split the dictionary into two components—one responsible for the predictive
1517 code and other for novel one. Our preliminary experiments show a Temporal SAE trained with
1518 this split-dictionary architecture is similarly performant as the tied dictionary one, resulting in
1519 predictive / novel codes with effective ranks $\sim 20\text{--}30 / \sim 390\text{--}420$, and a slightly better overall
1520 loss. It is possible optimizing hyperparameters (e.g., having different expansion factors for the two
1521 components) can make this architecture more performant than the tied dictionary one, but we leave a
1522 further characterization to future work.

1523

1524

1525

1526

1527

1528

1529

1530

1531

1532

1533

1534

1535

1536

1537

1538

1539

1540

1541

1542

1543

1544

1545

1546

1547

1548

1549

1550

1551

1552

1553

1554

1555

1556

1557

1558

1559

1560

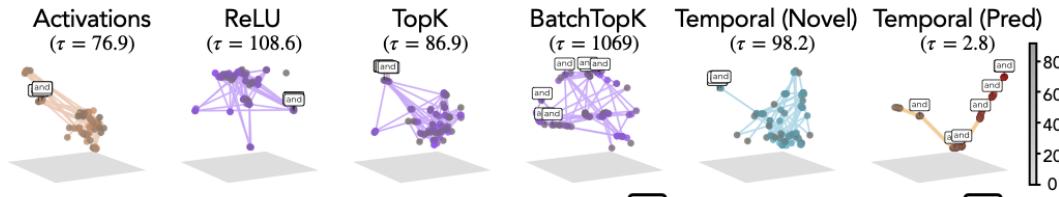
1561

1562

1563

1564

1565

1566 **G THE GEOMETRY OF STORIES: A NARRATIVE-DRIVEN DOMAIN**
15671568 **G.1 HIERARCHICAL CLUSTERING OF CODES FROM STORY TOKENS**
15691570 **G.1.1 STORY 1**
1571

1572
1573
1574
1575
1576
1577
1578 There was a baby who wanted to pick something special. He went to a shop **and** saw an ancient case. He picked it up **and** looked
1579 inside. It was full of unbelievably shiny gems **and** jewels. He couldn't believe his eyes. He was so excited **and** he knew this case was
1580 perfect. He quickly picked it up **and** wrapped it up with a big red bow. He smiled **and** couldn't wait to show his mum **and** dad.
1581
1582

Figure 28: **Geometry.** Repeating results of Fig. 4, we again find smooth trajectories in UMAP projections for Temporal SAEs.

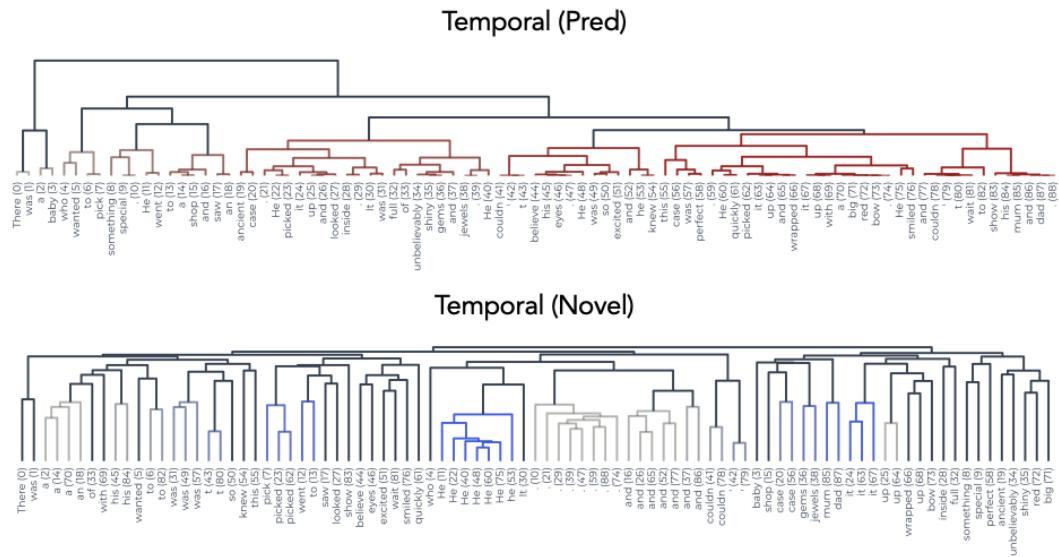


Figure 29: **Dendrograms of Predictive and Novel Components from Temporal SAEs.** Reproduction of the results from Fig. 4b, but including the novel component's dendrogram as well.

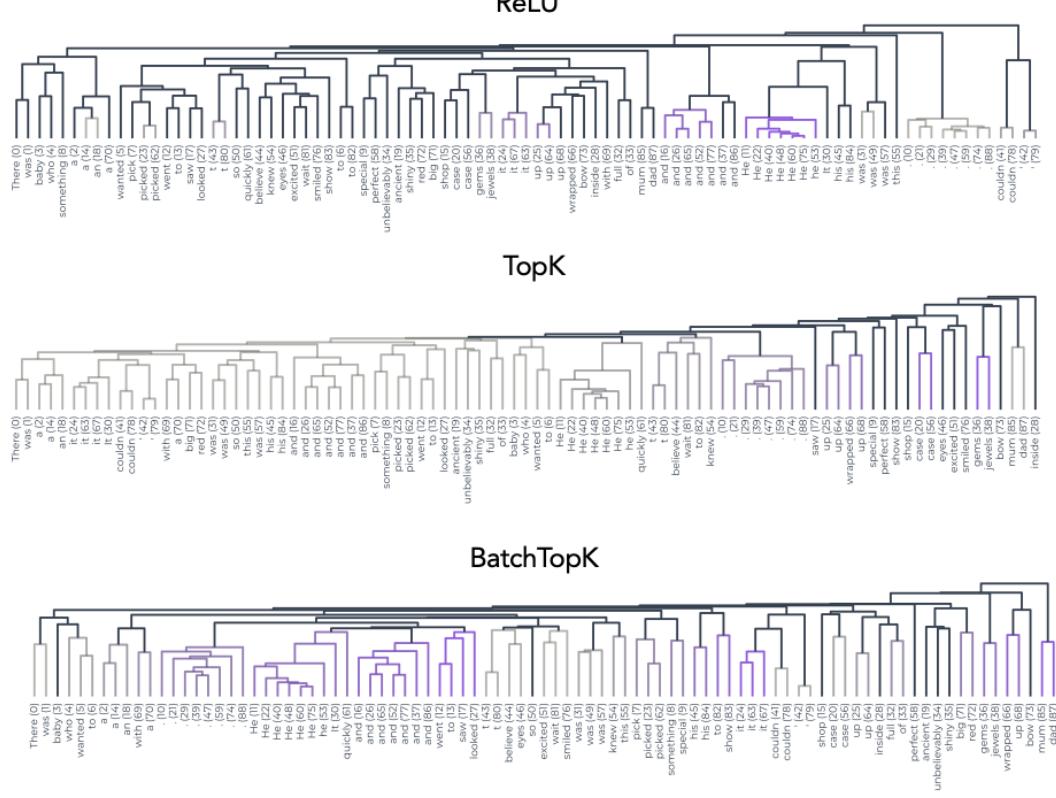
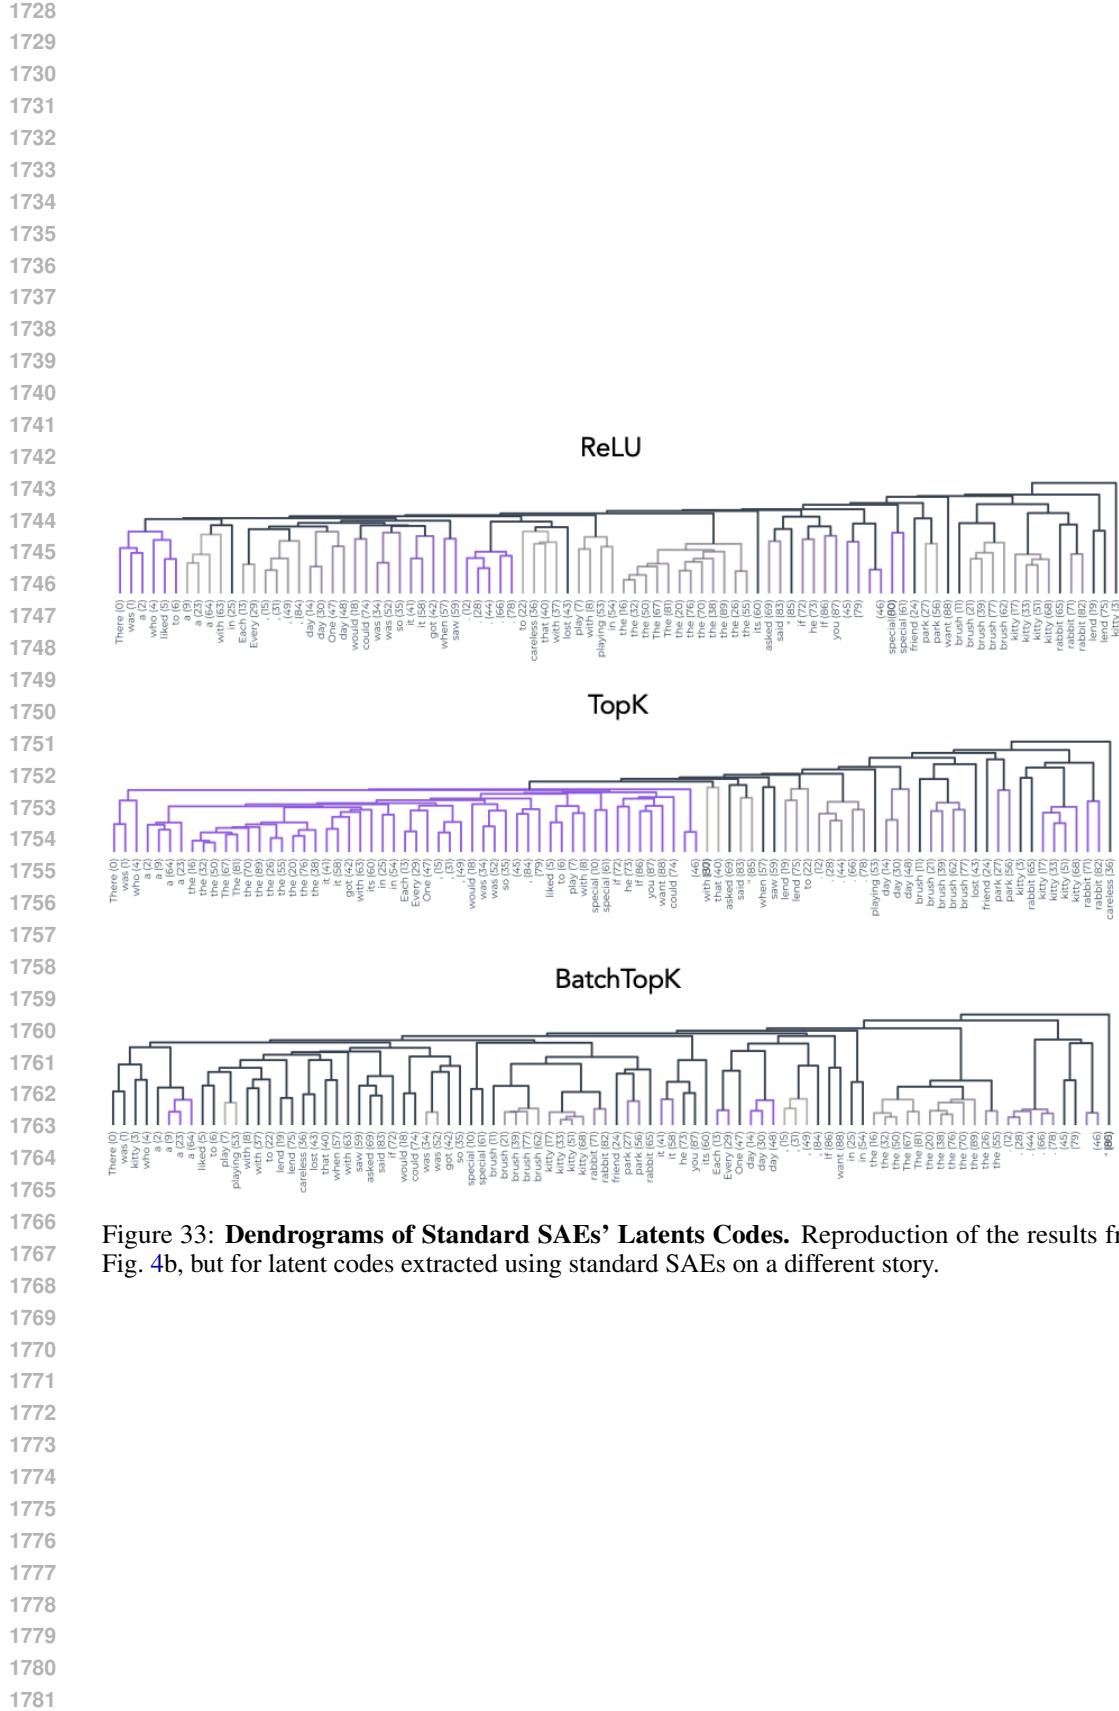
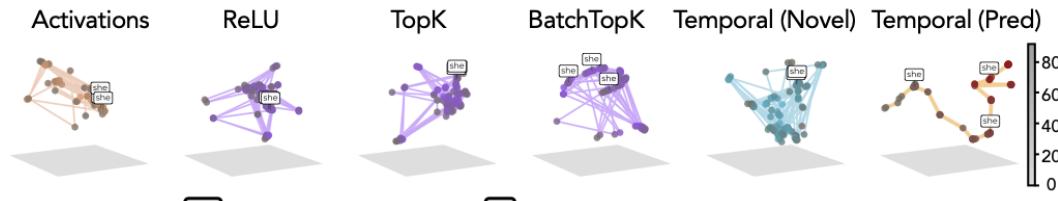


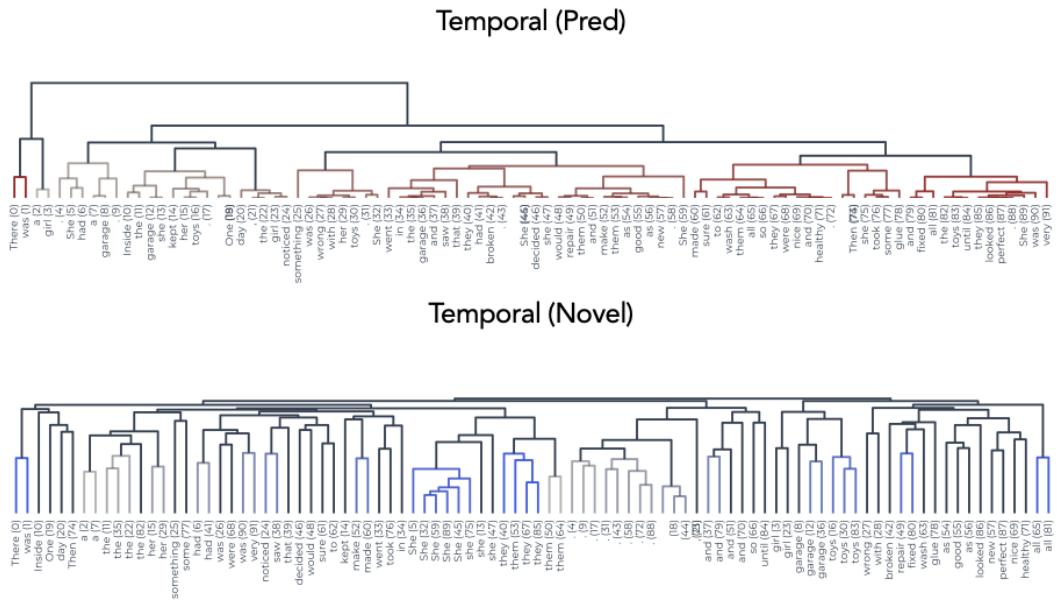
Figure 30: **Dendrograms of Standard SAEs’ Latents Codes.** Reproduction of the results from Fig. 4b, but for latent codes extracted using standard SAEs.



1782 G.1.3 STORY 3
1783

1790 Once there was a girl. **she** had a garage. Inside the garage **she** kept her toys. One day, the girl noticed something was wrong with
1791 her toys. **she** went in the garage and saw that they had broken. **she** decided she would repair them and make them as good as
1792 new. **she** made sure to wash them all so they were nice and healthy. Then **she** took some glue and fixed all the toys until they
1793 looked perfect.

1794 **Figure 34: Geometry.** Repeating results of Fig. 4 on a different story, we again find smooth
1795 trajectories in UMAP projections for Temporal SAEs.



1816 **Figure 35: Dendrograms of Predictive and Novel Components from Temporal SAEs.** Repro-
1817 duction of the results from Fig. 4b, but on a different story and with both the predictive and novel
1818 component.

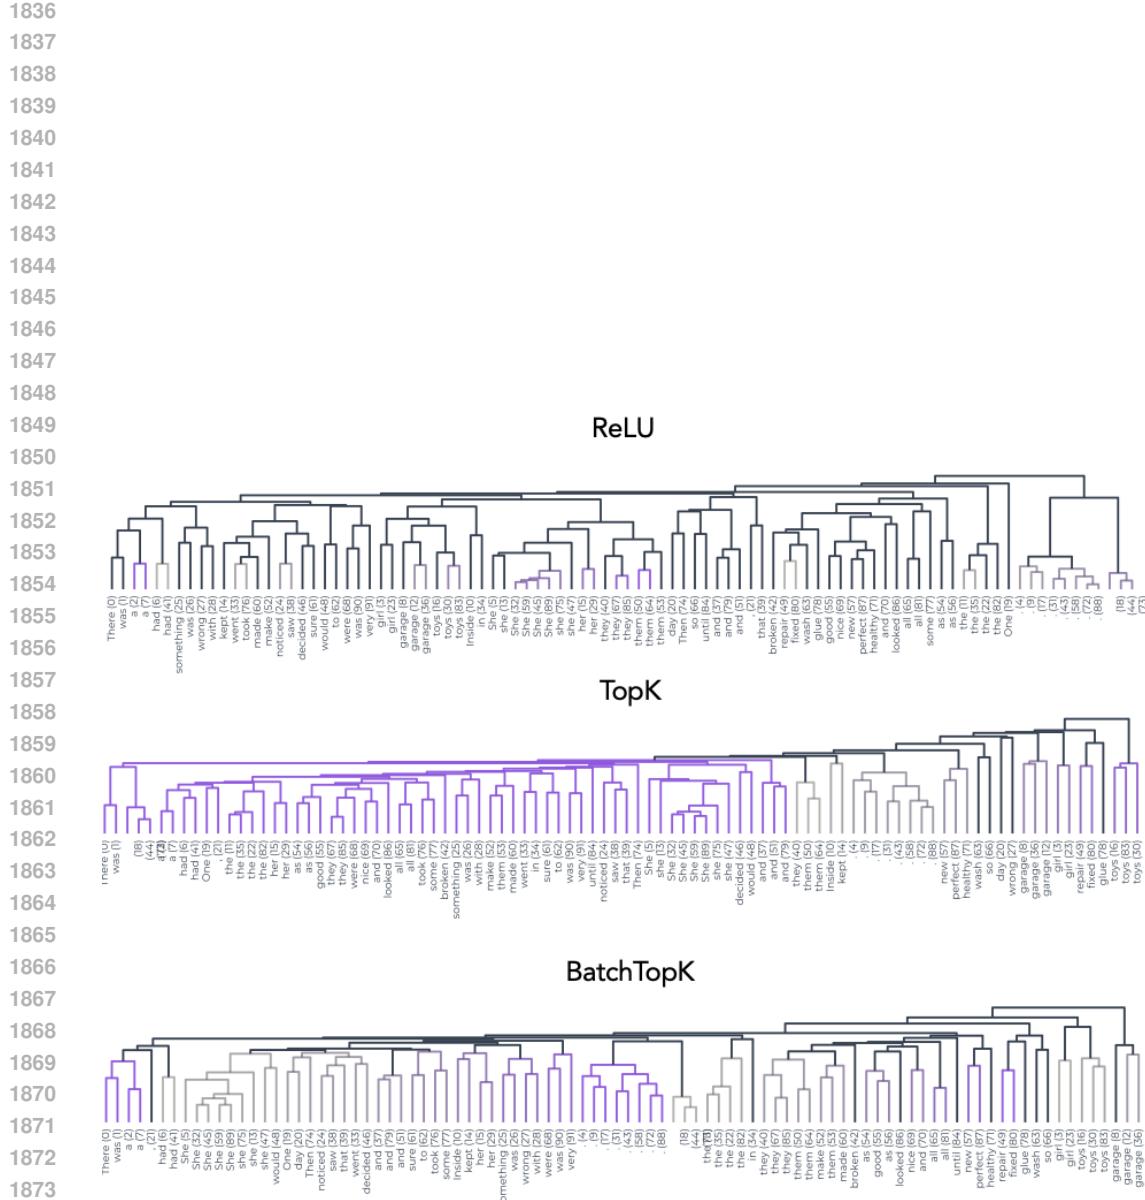
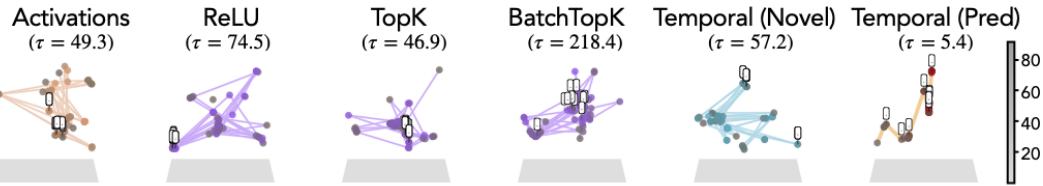


Figure 36: **Dendrograms of Standard SAEs’ Latents Codes.** Reproduction of the results from Fig. 4b, but for latent codes extracted using standard SAEs on a different story.

1890

1891 G.2 CODE: ANALYZING ANOTHER DOMAIN

1892



1898

1899 torch ; m = torch . nn . Linear (1 2 8 , 6 4) ; torch . nn . init . kaiming _ uniform _ (m . weight , non linear ity = ' relu ') ; torch . nn . init . zeros _ (m . bias) ; print (' init ' , m . weight)

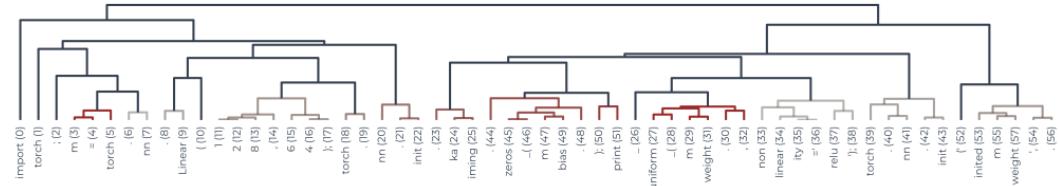
1900

1901 Figure 37: **Code.** Reproducing the results of Fig. 4, but on a different domain, i.e., code. We again see
1902 a temporally disentangled, smoothly running trajectory for latent codes extracted using the predictive
1903 component of Temporal SAEs.

1904

1905 Temporal (Pred)

1906



1907

1908

1909

1910

1911

1912

1913

1914

1915

1916

1917

1918

1919

1920

1921

1922 Figure 38: **Dendrograms of Predictive and Novel Components from Temporal SAEs.** Reproduction
1923 of the results from Fig. 4b, but on a different domain shows similar results as with narrative-driven
1924 text (stories) for both the predictive and novel component.

1925

1926

1927

1928

1929

1930

1931

1932

1933

1934

1935

1936

1937

1938

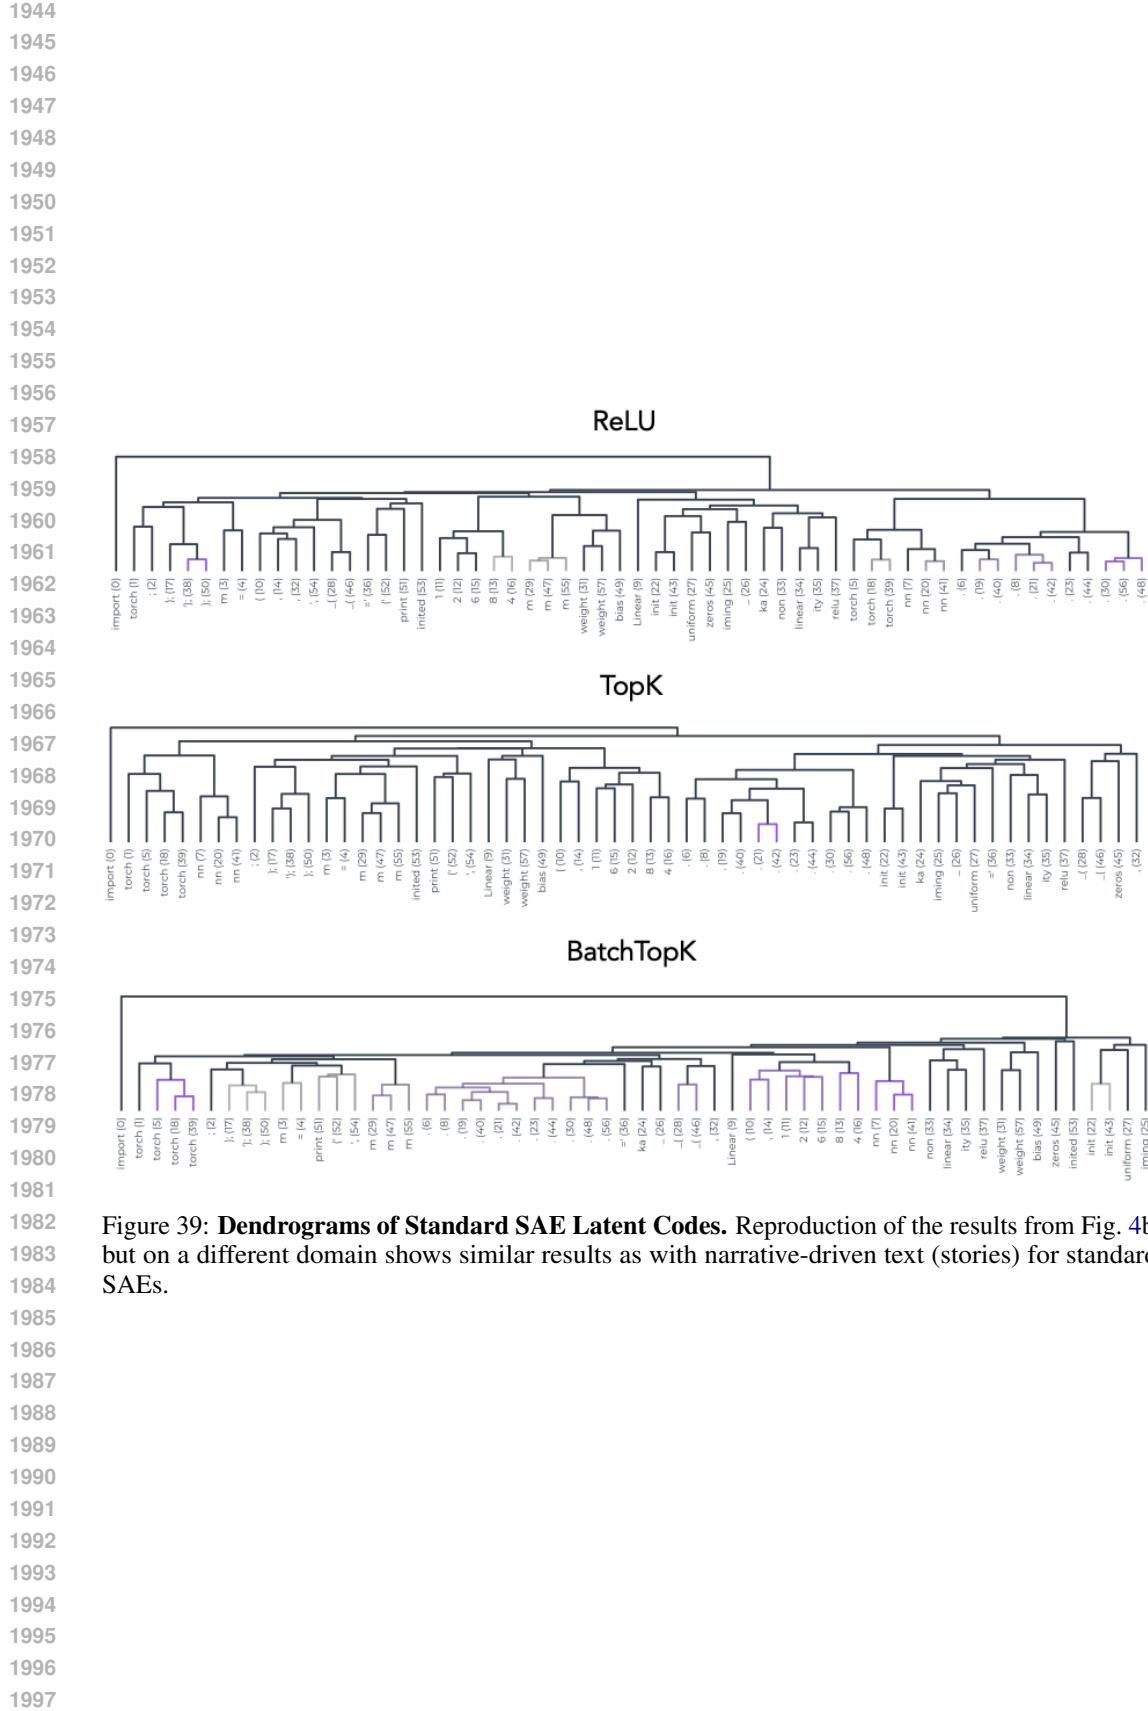
1939

1940

1941

1942

1943



1998
1999
2000
2001
2002
2003
2004
2005
2006
2007
2008
2009
2010
2011
2012
2013
2014
2015
2016
2017
2018
2019
2020

G.3 FURTHER RESULTS: SIMILARITY MAPS ON STORIES

Temporal (Pred)

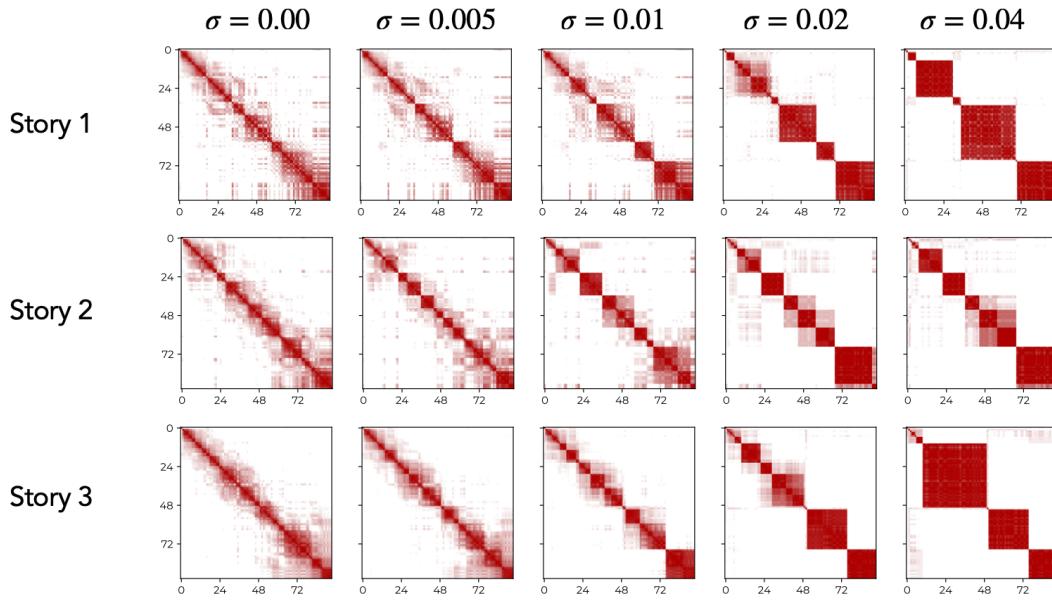
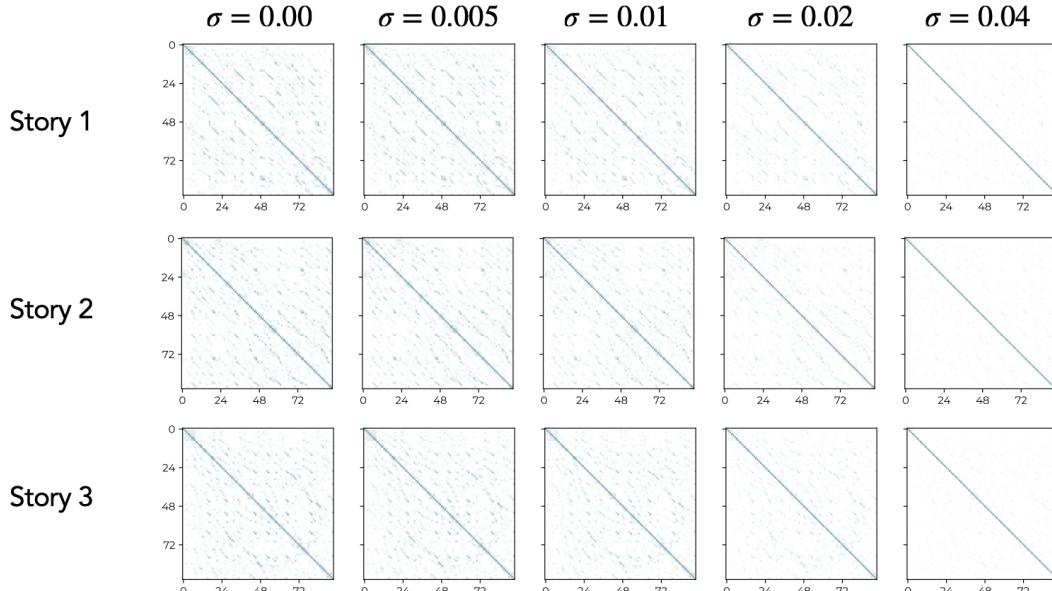


Figure 40: **Temporal (Pred) Similarity Maps Elicit Multi-Scale Structure with Noise.** Repeating the analysis shown in Fig. 6, we find the coarsening of temporal blocks is a robust result that continues to hold for different inputs.

2021
2022
2023
2024
2025
2026
2027
2028
2029
2030
2031
2032
2033
2034
2035
2036
2037
2038
2039
2040
2041
2042
2043
2044
2045

Temporal (Novel)



2046
2047
2048
2049
2050
2051

Figure 41: **Temporal (Novel) Similarity Maps under Noise.** Repeating the analysis shown in Fig. 6, we find the novel component is only able to capture minimal local similarities, which when analyzed via dendograms, show clustering based on lexical information.

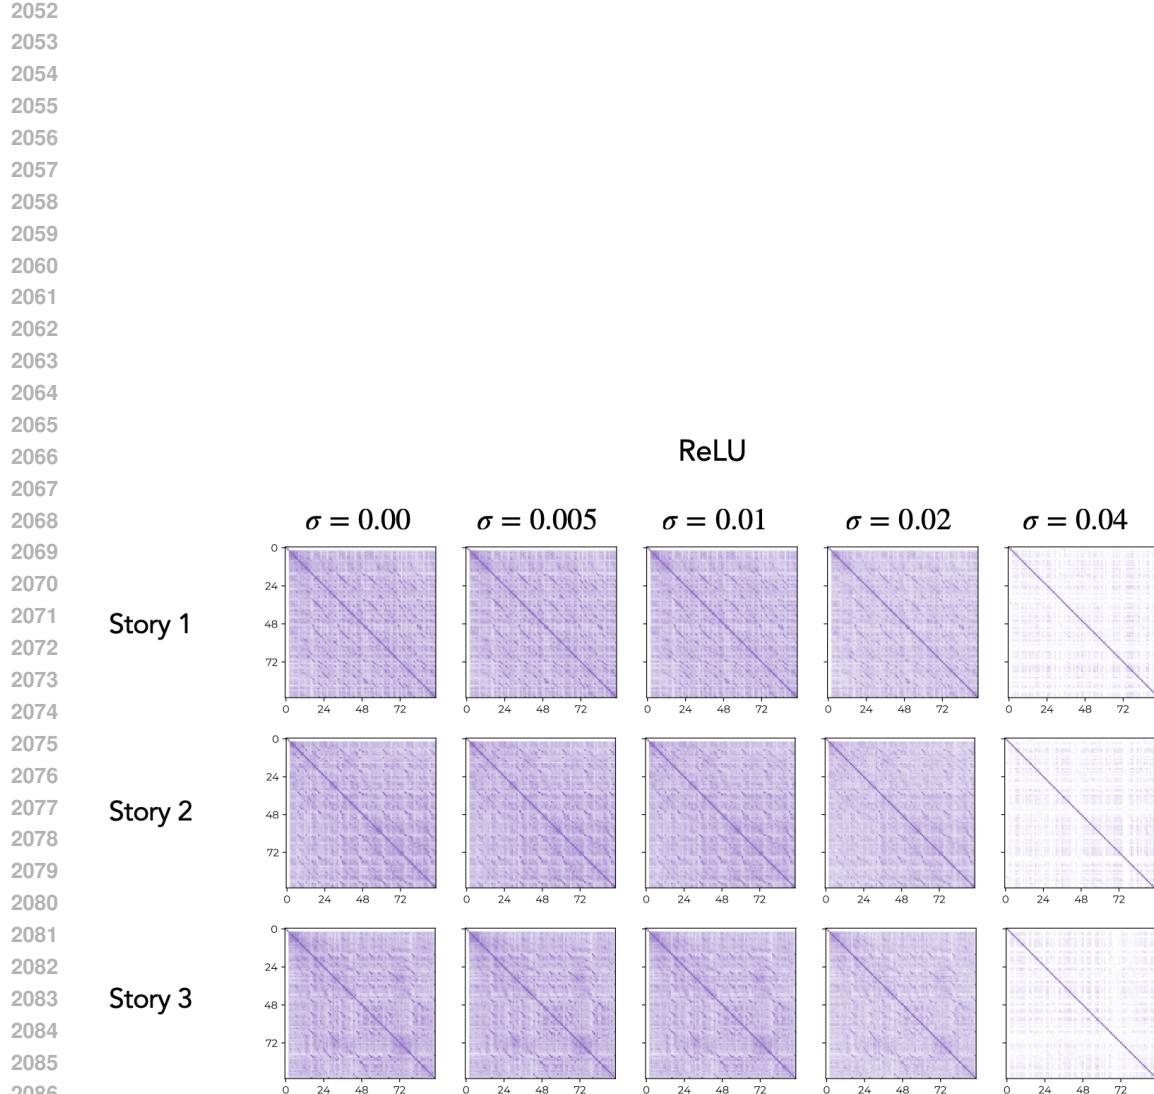
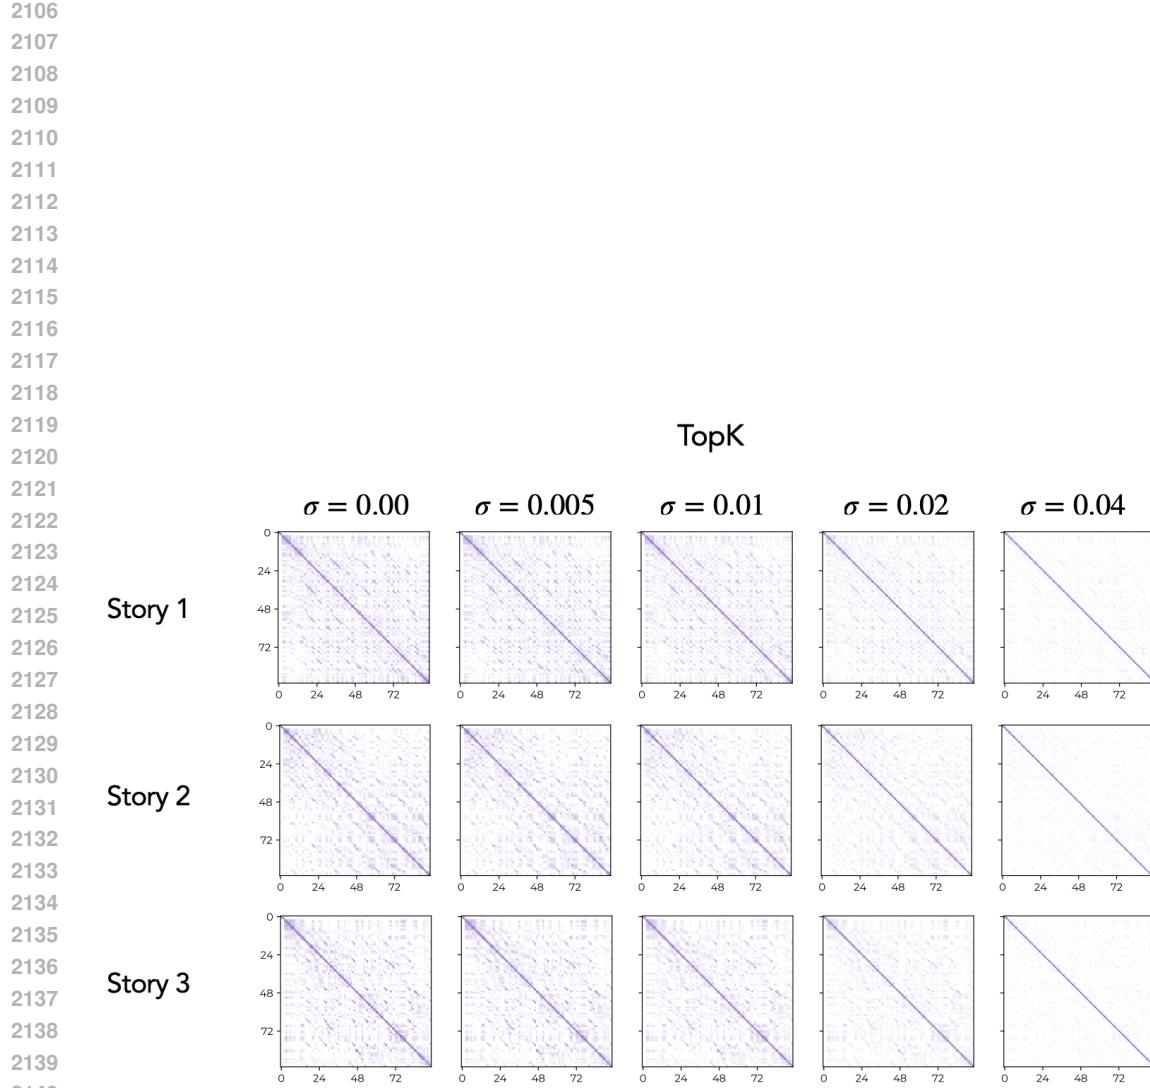
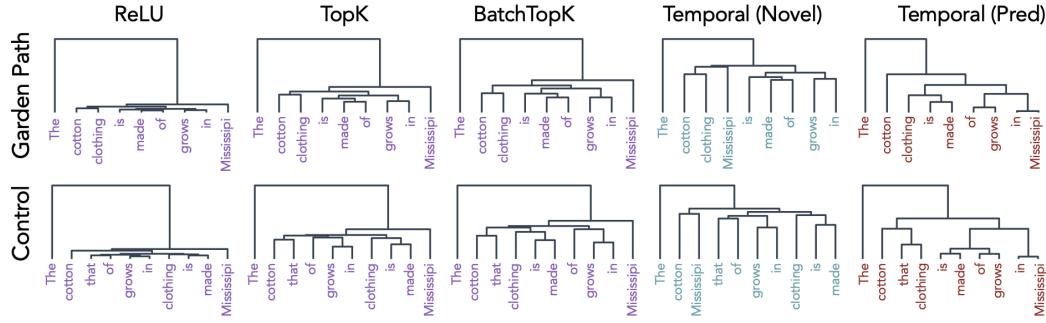


Figure 42: **ReLU Codes’ Similarity Maps under Noise.** Repeating the analysis shown in Fig. 6 on ReLU SAEs, we find the ReLU latent code has high similarity across the board, suggesting lack of meaningful temporal information. This similarity is entirely removed when noise scale increases too much, which, as shown in Fig. 6c, corresponds to the point that variance explained by ReLU SAEs drops to ~ 0 .



2141 **Figure 43: TopK Codes' Similarity Maps under Noise.** Repeating the analysis shown in Fig. 6,
 2142 we find TopK SAE's latent codes are only able to capture minimal local similarities, which when
 2143 analyzed via dendrograms, show clustering based on lexical information. Akin to ReLU SAEs, we
 2144 see this similarity map approximately turns into an identity matrix when the noise scale increases too
 2145 much, which, as shown in Fig. 6c, corresponds to the point that variance explained by TopK SAEs
 2146 drops to ~ 0 .

2147
 2148
 2149
 2150
 2151
 2152
 2153
 2154
 2155
 2156
 2157
 2158
 2159

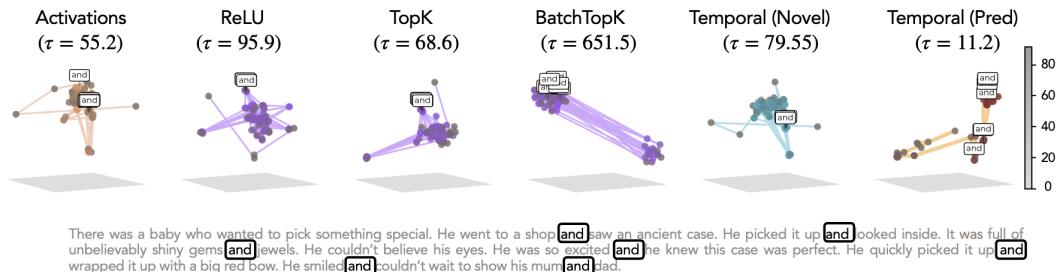
2214 G.4 FURTHER RESULTS: DENDROGRAMS AND EVALUATIONS ON GARDEN PATH SENTENCES
22152216 G.4.1 DENDROGRAMS
22172218 Figure 45: **Example Sentence 1.** Repeating the results of Fig. 7a on a different garden path sentence.
2219
2220
2221
2222
2223
2224
2225
2226
2227

2268 H FURTHER RESULTS: TEMPORAL SAEs ON LLAMA-3.1-8B

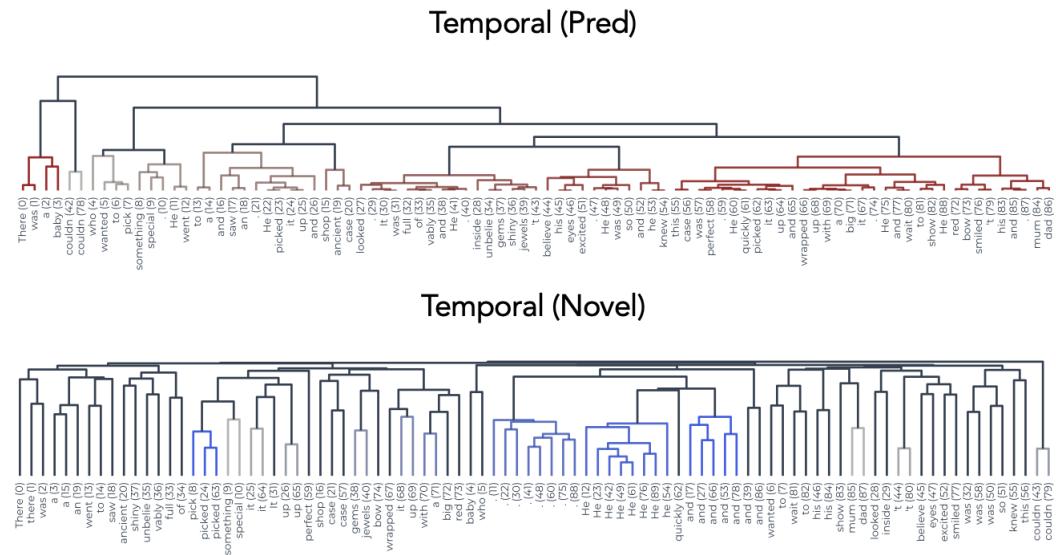
2270
 2271 We replicate a subset of the experiments evaluating Temporal SAEs on Llama-3.1-8B. The training
 2272 protocol remains the same as that of Gemma-2-2B model: train on 1B token activations with similar
 2273 normalization schema, but activations are harvested now from Layer 15 (i.e., at $\sim 50\%$ of model
 2274 depth). We note that the trained SAEs are not finished training, and hence the following results are
 2275 solely meant to be an impression of whether the qualitative trends observed with Gemma models
 2276 generalize to another model class.

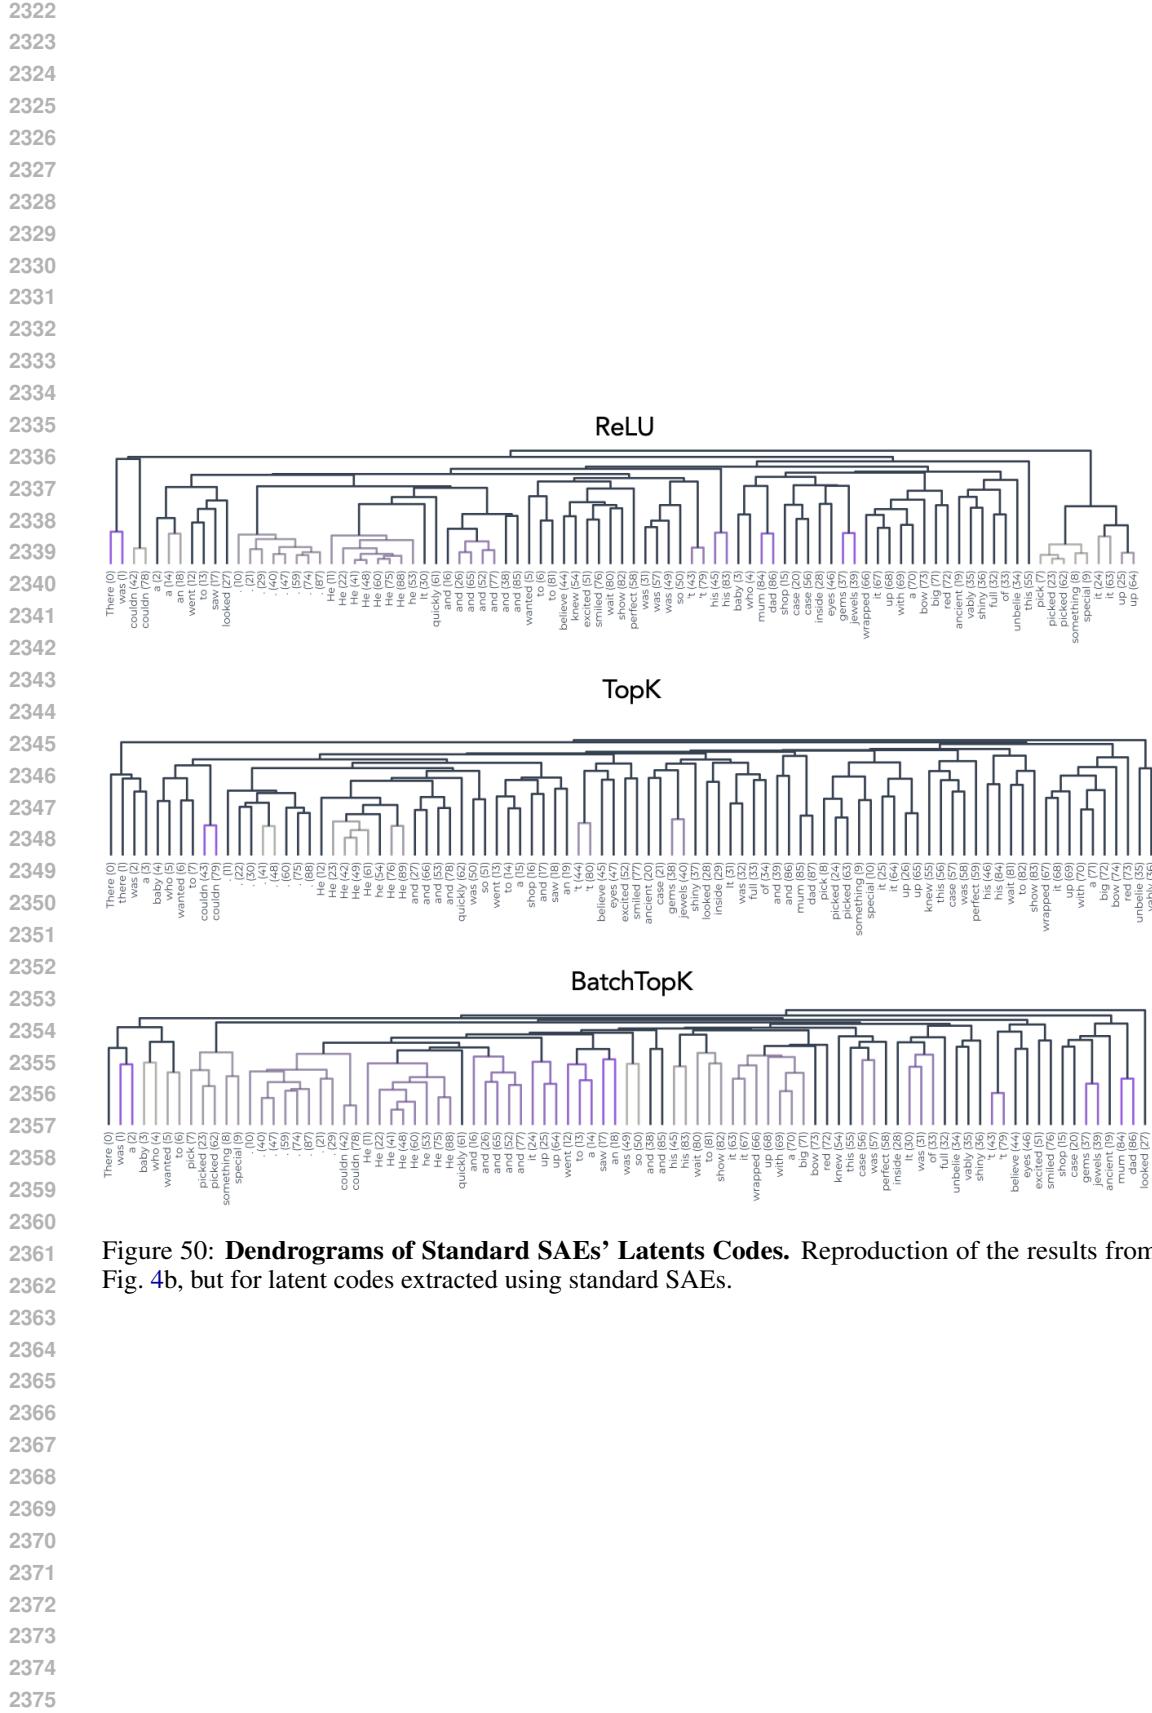
2277 H.1 GEOMETRY, DENDROGRAMS, AND SPECTRA

2279 H.1.1 STORY 1



2290 Figure 48: **Geometry.** Repeating results of Fig. 4, we again find smooth trajectories in UMAP
 2291 projections for Temporal SAEs.



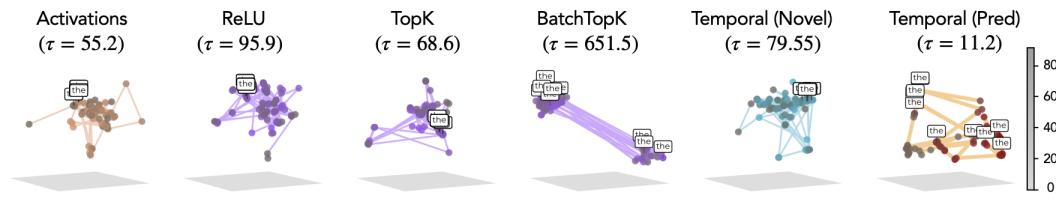


2376

2377

H.1.2 STORY 2

2378
2379
2380
2381
2382
2383
2384

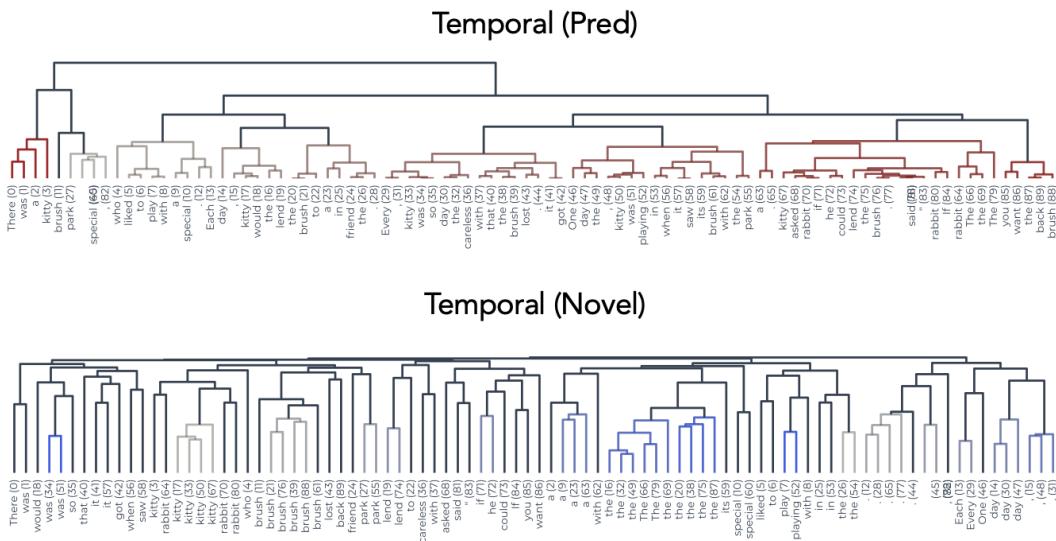


2385 There was a kitty who liked to play with a special brush. Each day, the kitty would lend the brush to a friend in the park. Every day, the kitty was so careless with the brush that it got lost. One day, the kitty was playing in the park when it saw its special brush with a rabbit. The kitty asked the rabbit if he could lend the brush. "No," the rabbit said, "If you want the brush back

2386

Figure 51: **Geometry.** Repeating results of Fig. 4, we again find smooth trajectories in UMAP projections for Temporal SAEs.

2390
2391
2392
2393
2394
2395
2396
2397



2407
2408
2409

Figure 52: **Dendrograms of Predictive and Novel Components from Temporal SAEs.** Reproduction of the results from Fig. 4b, but including the novel component’s dendrogram as well.

2410
2411
2412
2413
2414
2415
2416
2417
2418
2419
2420
2421
2422
2423
2424
2425
2426
2427
2428
2429

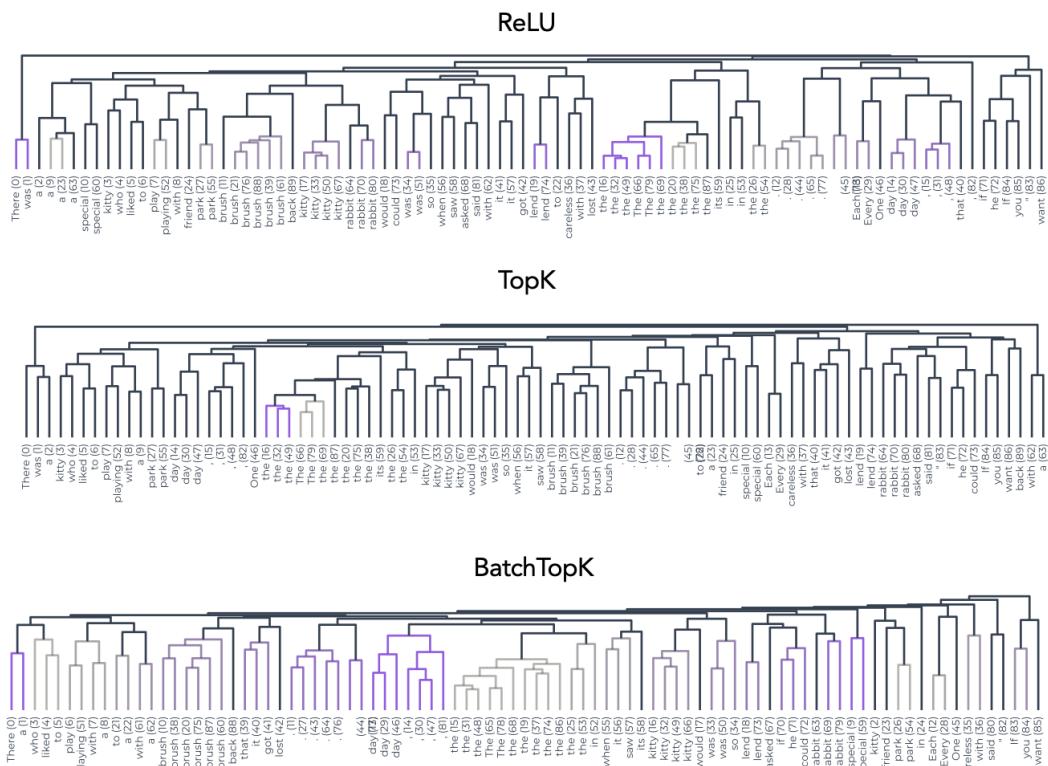


Figure 53: **Dendrograms of Standard SAEs' Latents Codes.** Reproduction of the results from Fig. 4b, but for latent codes extracted using standard SAEs.

2484
2485

H.1.3 STORY 3

2486

2487

2488

2489

2490

2491

2492

2493

2494

2495

2496

2497

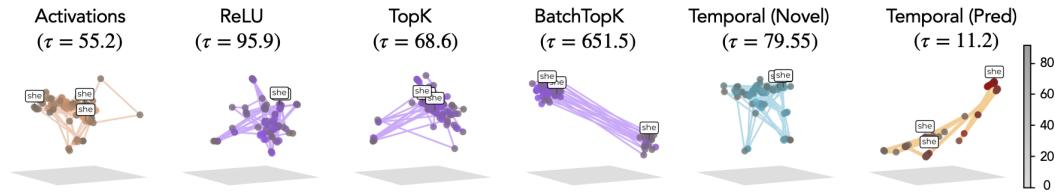


Figure 54: **Geometry.** Repeating results of Fig. 4, we again find smooth trajectories in UMAP projections for Temporal SAEs.

2498

2499

2500

2501

2502

2503

2504

2505

2506

2507

2508

2509

2510

2511

2512

2513

2514

2515

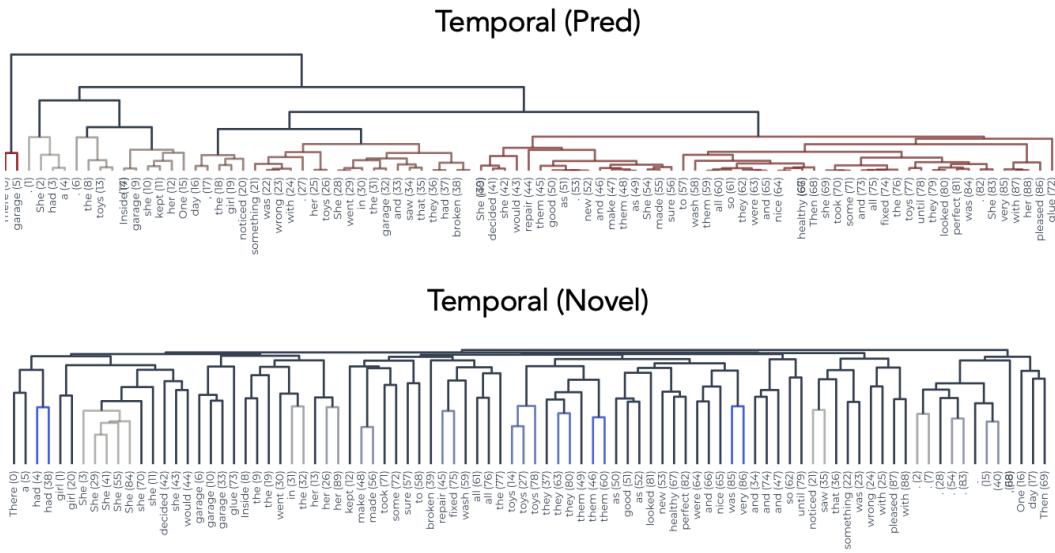


Figure 55: **Dendrograms of Predictive and Novel Components from Temporal SAEs.** Reproduction of the results from Fig. 4b, but including the novel component's dendrogram as well.

2516

2517

2518

2519

2520

2521

2522

2523

2524

2525

2526

2527

2528

2529

2530

2531

2532

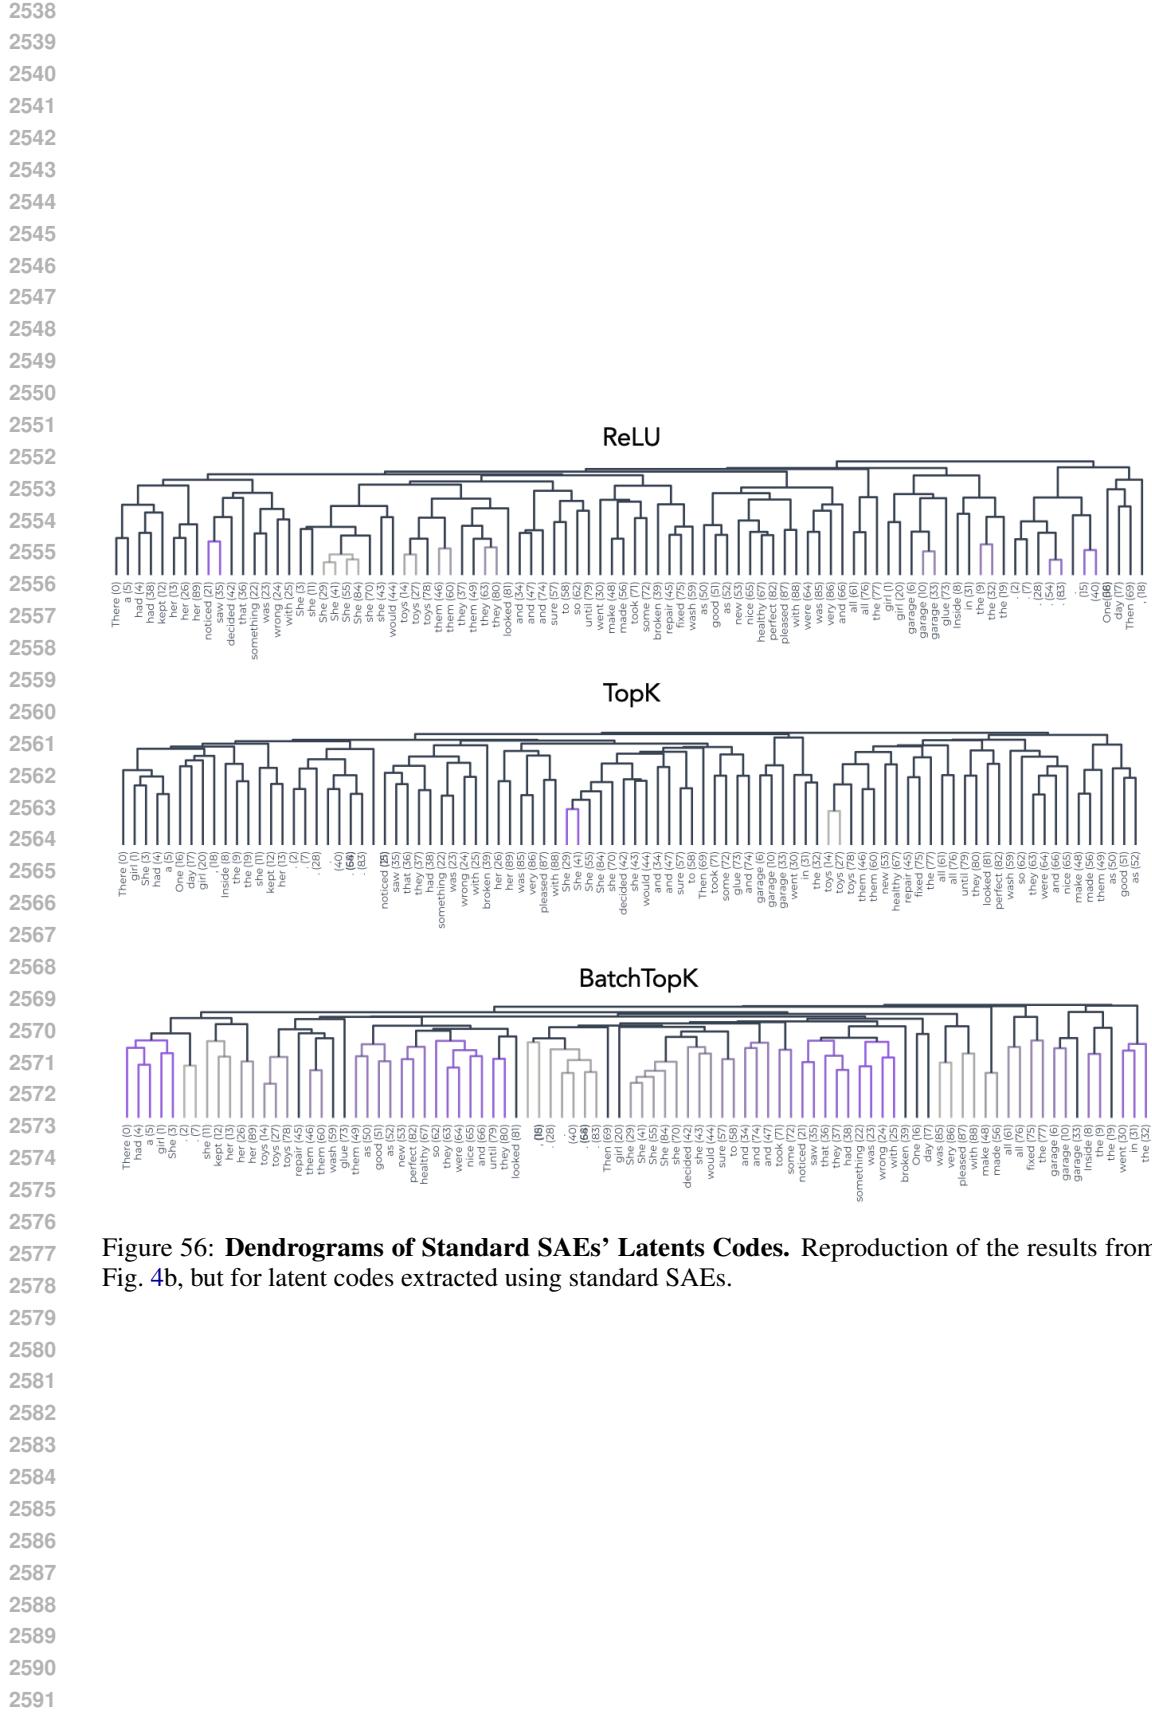
2533

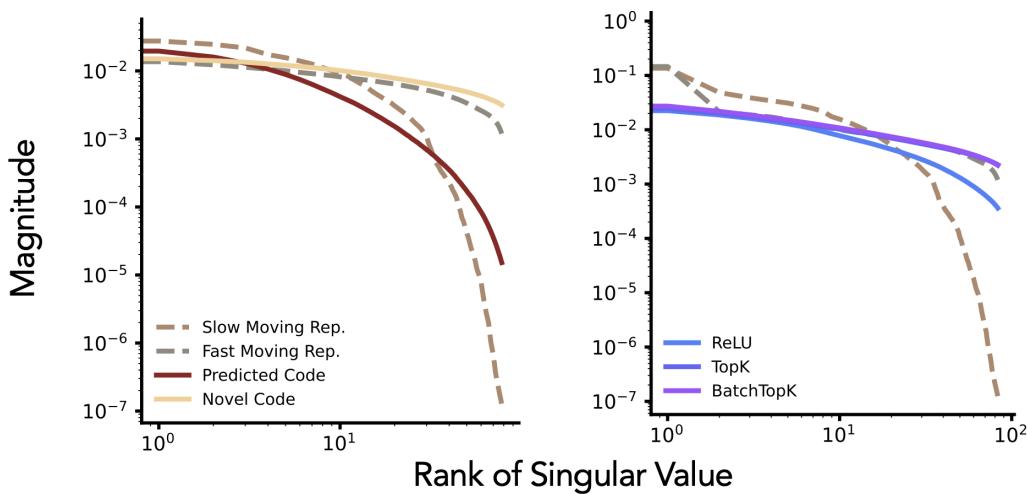
2534

2535

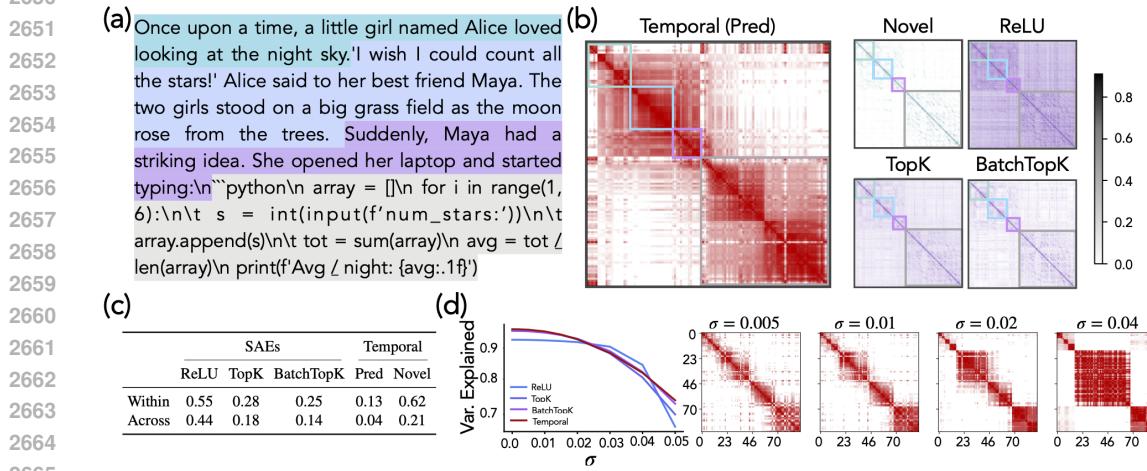
2536

2537



2592 H.1.4 COMPARING EIGENSPECTRUM WITH SLOW VS. FAST FEATURES
2593
2594
2595
2596
2597
2598
2599
2600
2601
2602
2603
2604
2605
2606
2607

2609
2610
2611
2612 **Figure 57: Kernel spectrum for latent codes and model representations.** Kernels defined using
2613 novel code from Temporal SAEs and standard SAEs both align well with the fast-changing part
2614 of model representations; meanwhile, only the predictive code shows strong similarity to the slow
2615 changing part.
2616
2617
2618
2619
2620
2621
2622
2623
2624
2625
2626
2627
2628
2629
2630
2631
2632
2633
2634
2635
2636
2637
2638
2639
2640
2641
2642
2643
2644
2645

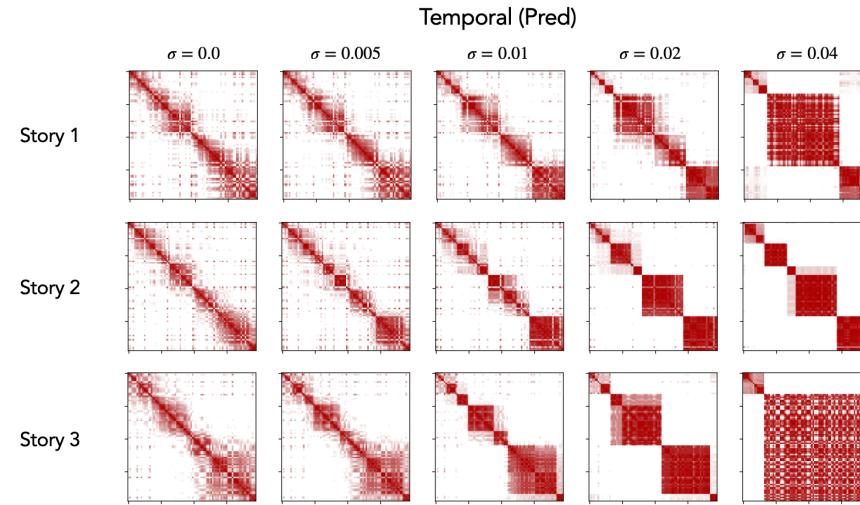
2646 H.2 EVENT BOUNDARIES AND NOISE STABILITY
26472648 H.2.1 EVENT BOUNDARIES
2649

2666 Figure 58: **Analysis of Event Boundaries.** Reproducing results from Fig. 6 on Llama-3.1-8B, we
2667 see similar behavior as the Gemma analysis done previously.

2668

2669

2670



2688 Figure 59: **Temporal (Pred) Similarity Maps Elicit Multi-Scale Structure with Noise.** Repeating
2689 the analysis shown in Fig. 6, we find the coarsening of temporal blocks is a robust result that continues
2690 to hold for different inputs.

2691

2692

2693

2694

2695

2696

2697

2698

2699

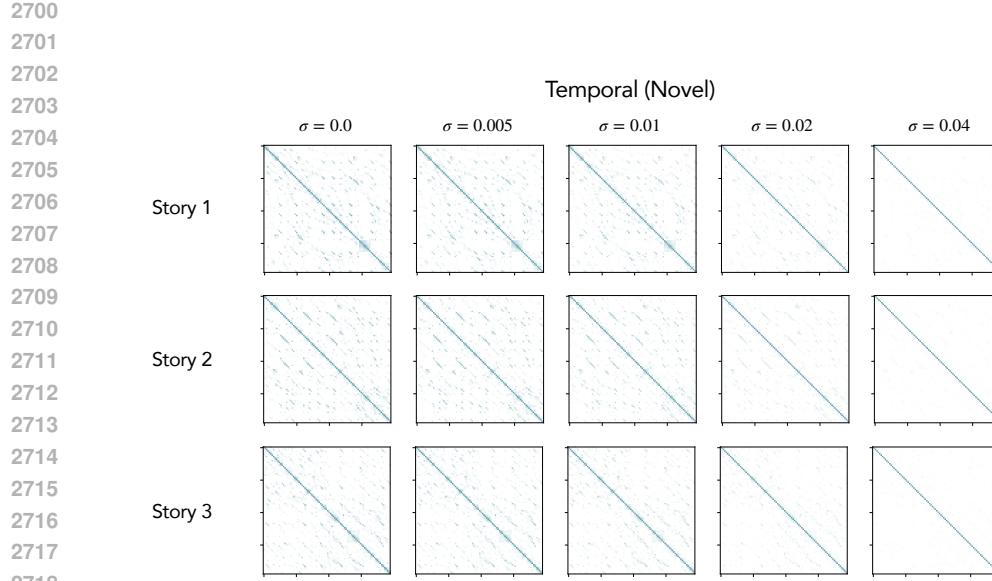
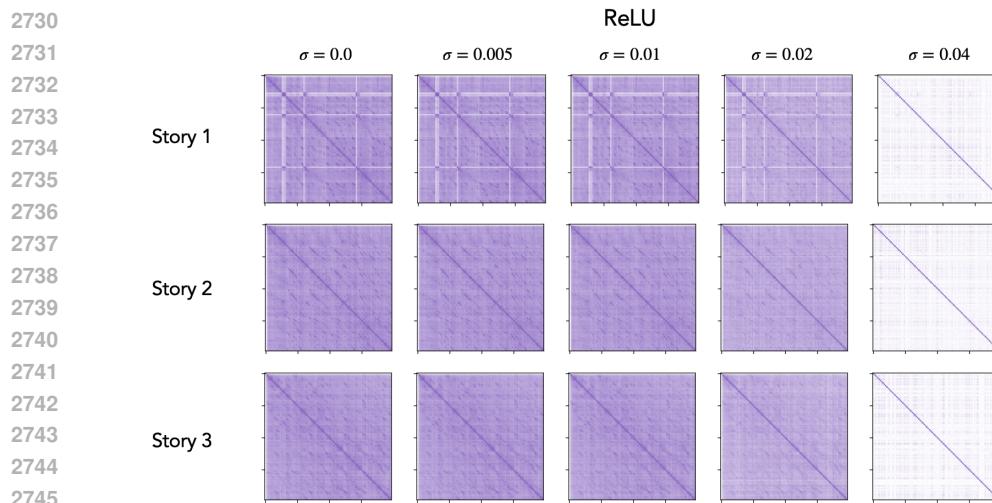


Figure 60: **Temporal (Novel) Similarity Maps under Noise.** Repeating the analysis shown in Fig. 6, we find the novel component is only able to capture minimal local similarities, which when analyzed via dendograms, show clustering based on lexical information.



2749
 2750
 2751
 2752
 2753

Figure 61: **ReLU Codes' Similarity Maps under Noise.** Repeating the analysis shown in Fig. 6 on ReLU SAEs, we find the ReLU latent code has high similarity across the board, suggesting lack of meaningful temporal information.

2754

2755

2756

2757

2758

2759

2760

2761

2762

2763

2764

2765

2766

2767

2768

2769

2770

2771

2772

2773

2774

2775

2776

2777

2778

2779

2780

2781

2782

2783

2784

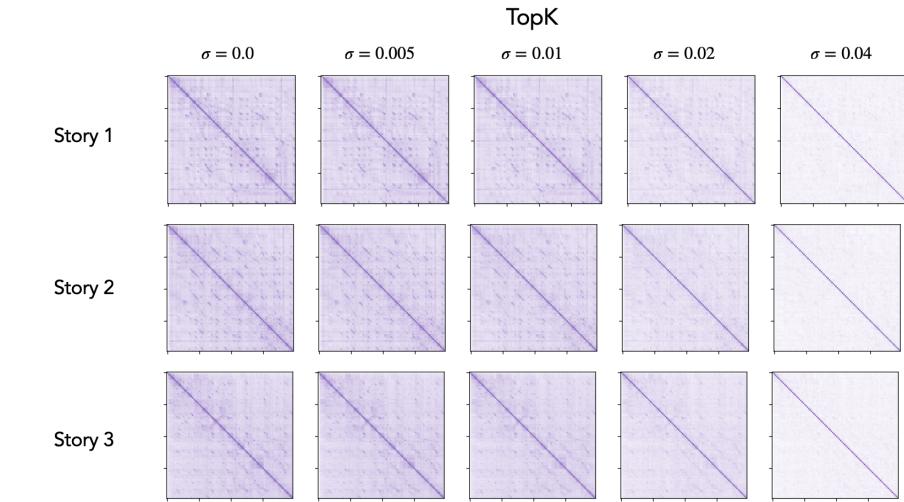


Figure 62: **TopK Codes’ Similarity Maps under Noise.** Repeating the analysis shown in Fig. 6, we find TopK SAE’s latent codes are only able to capture minimal local similarities, which when analyzed via dendograms, show clustering based on lexical information.

2785

2786

2787

2788

2789

2790

2791

2792

2793

2794

2795

2796

2797

2798

2799

2800

2801

2802

2803

2804

2805

2806

2807

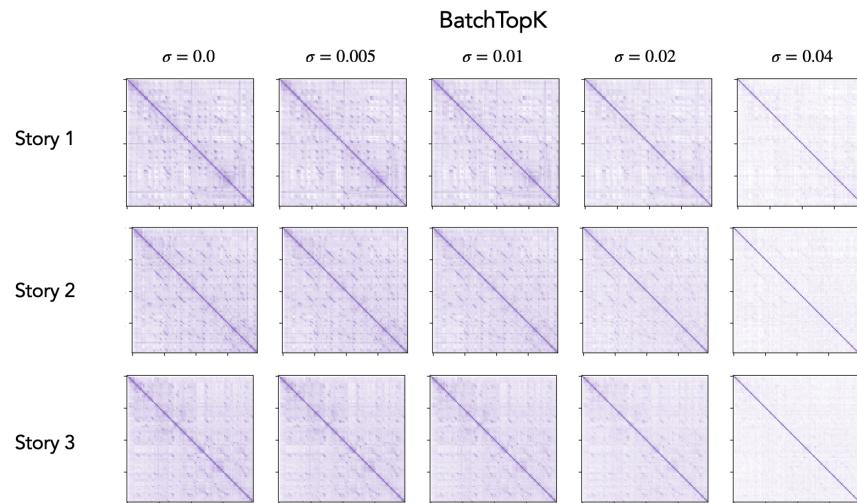


Figure 63: **BatchTopK Codes’ Similarity Maps under Noise.** Repeating the analysis shown in Fig. 6, we find BatchTopK SAE’s latent codes are only able to capture minimal local similarities, which when analyzed via dendograms, show clustering based on lexical information.

2808 **I FURTHER THEORY RESULTS**
 2809

2810 **I.1 PRIORS ON THE SPARSE CODE FOR VARIOUS SAEs**
 2811

2812 We restate and prove the proposition on independence priors of SAEs over time (Proposition 4.1)
 2813 below.

2814 **Proposition I.1** (Independence priors over time). *Consider the SAE maximum a posteriori (MAP)*
 2815 *objective for ReLU, JumpReLU, TopK and BatchTopK SAEs. The sparsity constraints for these SAEs*
 2816 *are additive over time, resulting in:*

$$\begin{aligned} & \arg \min_{\mathbf{D}, \mathbf{z}} \frac{1}{T} \sum_{i=1}^T \|\mathbf{x}_i - \mathbf{D}\mathbf{z}_i\|_2^2 + \lambda \mathcal{R}(\mathbf{z}_i), \\ & \text{s.t. } \mathbf{z}_k = f_{\text{SAE}}(\mathbf{x}_k) \forall k, \tilde{g}(\mathbf{z}_1, \dots, \mathbf{z}_T) = \sum_i \tilde{g}(\mathbf{z}_i) = 0. \end{aligned} \quad (8)$$

2820 *This MAP objective has an independent and identically distributed (i.i.d.) prior over time i.e.,*
 2821

$$P(\mathbf{z}_1, \dots, \mathbf{z}_T) \propto \prod_{t=1}^T \exp \left(-\lambda \mathcal{R}(\mathbf{z}_i) - \tilde{\lambda} \tilde{g}(\mathbf{z}_i) \right) = \prod_i P(\mathbf{z}_i),$$

2824 *Proof.* The sparsity constraints and sparsity-promoting regularizers for the SAEs under study are
 2825 specified in the table below.

SAE	Regularizer $\mathcal{R}(\mathbf{z}_i)$	Sparsity Constraint $\tilde{g}(\mathbf{z}_1, \dots, \mathbf{z}_T) = 0$
ReLU	$\ \mathbf{z}_i\ _1$	0
JumpReLU	$\ \mathbf{z}_i\ _0$	0
TopK	0	$\sum_{i=1}^T (\ \mathbf{z}_i\ _0 - K)^2$
BatchTopK	0	$\frac{1}{T} \sum_{i=1}^T \ \mathbf{z}_i\ _0 - K$

2826 Table 5: Sparsity constraints and regularizers for SAEs
 2827

2828 Note that TopK imposes a pointwise hard sparsity constraint, which has been restated using sum-of-
 2829 squares above for convenience. While BatchTopK imposes the fixed mean sparsity for each mini
 2830 batch, we take the batch to capture the entire timeseries in the above formulation. The above table
 2831 shows us that the sparsity constraint is additive over time in all cases:

$$\tilde{g}(\mathbf{z}_1, \dots, \mathbf{z}_T) = \sum_{i=1}^T \tilde{g}(\mathbf{z}_i) = 0, \quad (9)$$

$$\text{where } \tilde{g}(\mathbf{z}_i) = \begin{cases} (\|\mathbf{z}_i\|_0 - K)^2 & \text{TopK} \\ \frac{1}{T} (\|\mathbf{z}_i\|_0 - K) & \text{BatchTopK} \\ 0 & \text{ReLU, JumpReLU} \end{cases} \quad (10)$$

2832 Recall that SAEs solve the following constrained optimization problem (restated from Eq. 1).
 2833

$$\begin{aligned} & \arg \min_{\mathbf{D}, \mathbf{z}} \frac{1}{N} \sum_{i=1}^T \|\mathbf{x}_i - \mathbf{D}\mathbf{z}_i\|_2^2 + \lambda \mathcal{R}(\mathbf{z}_i), \\ & \text{s.t. } \mathbf{z}_k = f_{\text{SAE}}(\mathbf{x}_k) \forall k, \tilde{g}(\mathbf{z}_1, \dots, \mathbf{z}_T) = 0. \end{aligned} \quad (11)$$

2834 We rewrite the above problem using Lagrange multipliers on the sparsity constraints, and further
 2835 simplify using the above result on constraints being additive over time, as:

$$\begin{aligned} & \arg \min_{\mathbf{D}, \mathbf{z}} \frac{1}{N} \sum_{i=1}^T \|\mathbf{x}_i - \mathbf{D}\mathbf{z}_i\|_2^2 + \lambda \mathcal{R}(\mathbf{z}_i) + \tilde{\lambda} \tilde{g}(\|\mathbf{z}_i\|_0), \\ & \text{s.t. } \mathbf{z}_k = f_{\text{SAE}}(\mathbf{x}_k) \forall k. \end{aligned} \quad (12)$$

2836 Note that we don't use Lagrange multipliers on the SAE architecture constraint $\mathbf{z} = f_{\text{SAE}}(\mathbf{x})$ since
 2837 we only care about \mathbf{z} -specific constraints (which don't include the data \mathbf{x}) for the prior.

Bayesian Interpretation. The objective function above (sans the SAE architecture constraint) can be thought of as minimizing the negative log posterior, which is proportional to log prior added to log likelihood:

$$-\log P(\mathbf{z}_1, \dots, \mathbf{z}_T \mid \mathbf{x}_1, \dots, \mathbf{x}_T) \propto \underbrace{\frac{1}{N} \sum_{i=1}^T \|\mathbf{x}_i - \mathbf{D}\mathbf{z}_i\|_2^2}_{-\log P(\mathbf{x}_1, \dots, \mathbf{x}_T \mid \mathbf{z}_1, \dots, \mathbf{z}_T)} + \underbrace{\frac{1}{N} \sum_{i=1}^T \lambda \mathcal{R}(\mathbf{z}_i) + \tilde{\lambda} \tilde{g}(\|\mathbf{z}_i\|_0)}_{-\log P(\mathbf{z}_1, \dots, \mathbf{z}_T)} \quad (13)$$

The prior over latents \mathbf{z} is:

$$\log P(\mathbf{z}_1, \dots, \mathbf{z}_T) = -\frac{1}{N} \sum_{i=1}^T \lambda \mathcal{R}(\mathbf{z}_i) + \tilde{\lambda} \tilde{g}(\|\mathbf{z}_i\|_0), \quad (14)$$

$$\implies P(\mathbf{z}_1, \dots, \mathbf{z}_T) = \prod_{i=1}^T \exp \left(-\lambda \mathcal{R}(\mathbf{z}_i) + \tilde{\lambda} \tilde{g}(\|\mathbf{z}_i\|_0) \right) = \prod_i P(\mathbf{z}_i). \quad (15)$$

Therefore, the prior is multiplicative over time, implying independence, and the distribution at each time t has the same form, implying that the prior is independent and identically distributed (i.i.d.). This completes the proof. \square

I.2 PRIORS OVER CONCEPTS AND GENERATIVE PRIORS OVER TIME

We can further think of the priors of each SAE over concepts as well as over time in a generative fashion. In some cases, this manifests as a hierarchical latent variable model $n \rightarrow S \rightarrow \mathbf{z}$, where $n = \|\mathbf{z}\|_0$ is the sparsity, $S = \text{supp}(\mathbf{z}) = \{k : z_t^k > 0\}$ is the support, and \mathbf{z} is the SAE latent code.

Proposition I.2 (SAE Priors on Sparse Code). *Let $S_t = \text{supp}(\mathbf{z}_t) = \{k : z_t^k > 0\}$ be the set of active latents in the sparse code \mathbf{z} at time t , and $n_t = |S_t|$ be the cardinality of S_t (the number of active latents). Each SAE imposes a prior distribution on the sparse code \mathbf{z} , arising from its sparsity penalty $\mathcal{R}(\mathbf{z})$ or implicit conditions imposed on the sparse code. These conditions are highlighted in Table 6.*

Table 6: Priors over concept interactions and dynamics for various SAEs

$f_{\text{SAE}}, \mathcal{R}(\mathbf{z})$	Across-Concept Prior (interaction)	Across-time Prior (dynamics)
ReLU, L_1 -norm	$z_t^1, \dots, z_t^M \stackrel{\text{i.i.d.}}{\sim} \text{Laplace}(0, \cdot)$	$\mathbf{z}_1, \dots, \mathbf{z}_T \stackrel{\text{i.i.d.}}{\sim} P_{\mathbf{z}}$
TopK	$z_t^{i_1}, \dots, z_t^{i_K} \mid S_t \stackrel{\text{i.i.d.}}{\sim} U(0, \cdot) \forall i \in S_t,$ $S_t \sim U([M]^K)$	$(\mathbf{z}_1, S_1), \dots, (\mathbf{z}_t, S_t) \stackrel{\text{i.i.d.}}{\sim} P_S P_{\mathbf{z} S},$ $S_1, \dots, S_t \stackrel{\text{i.i.d.}}{\sim} U([M]^K)$
JumpReLU, L_0 -norm	$z_t^{i_1}, \dots, z_t^{i_{n_t}} \mid S_t \stackrel{\text{i.i.d.}}{\sim} U(0, \cdot) \forall i \in S_t,$ $S_t \mid n_t \sim U([M]^{n_t})$	$(\mathbf{z}_1, S_1, n_1), \dots, (\mathbf{z}_t, S_t, n_t) \stackrel{\text{i.i.d.}}{\sim} P_n P_{S n} P(\mathbf{z} \mid S),$ $n_1, \dots, n_t \stackrel{\text{i.i.d.}}{\sim} P_n$
BatchTopK	$z_t^{i_1}, \dots, z_t^{i_{n_t}} \mid S_t \stackrel{\text{i.i.d.}}{\sim} U(0, \cdot) \forall i \in S_t,$ $S_t \mid n_t \sim U([M]^{n_t})$	$(\mathbf{z}_1, S_1, n_1), \dots, (\mathbf{z}_t, S_t, n_t) \stackrel{\text{i.i.d.}}{\sim} P_n P_{S n} P(\mathbf{z} \mid S),$ $n_1, \dots, n_t \stackrel{\text{i.i.d.}}{\sim} P_n, \mathbb{E}[n_t] = K$

I.2.1 RELU SAE

The vanilla ReLU SAE (Bricken et al. (2023), Cunningham et al. (2023)) is trained with the L_1 -norm penalty:

$$\mathcal{R}(\mathbf{z}) = \|\mathbf{z}\|_1. \quad (16)$$

The prior over \mathbf{z} for the above case is:

$$\log P(\mathbf{z}_1, \dots, \mathbf{z}_N) \propto - \sum_{i=1}^N \sum_{k=1}^M |z_i^k|, \quad (17)$$

$$\implies P(\mathbf{z}_1, \dots, \mathbf{z}_N) \propto \prod_{i=1}^N \left(\prod_{k=1}^M \exp -\nu |z_i^k| \right). \quad (18)$$

2916 This joint distribution implies that for each sample i , different indices k are sampled i.i.d. from the
 2917 same distribution:
 2918

$$2919 \quad z_i^1, \dots, z_i^M \stackrel{\text{i.i.d.}}{\sim} \text{Laplace}(0, 1/\nu), \quad (19)$$

2920 and different samples are all independently sampled from the same product Laplace distribution:
 2921

$$2922 \quad \mathbf{z}_1, \dots, \mathbf{z}_N \stackrel{\text{i.i.d.}}{\sim} \text{Laplace}^M(0, 1/\nu). \quad (20)$$

2923 This concludes the proof for priors of ReLU SAE trained with L_1 norm sparsity penalty. \square
 2924

2925 I.2.2 TOPK SAE

2927 The TopK SAE (Makhzani and Frey (2013), Gao et al. (2024)) directly controls the sparsity of the
 2928 representation \mathbf{z} by fixing it at $\|\mathbf{z}\|_0 = K$, instead of imposing an explicit sparsity penalty $\mathcal{R}(\mathbf{z})$ in
 2929 the loss function. The objective function for TopK SAE is:

$$2931 \quad \arg \min_{\mathbf{D}, \mathbf{z}} \sum_{i=1}^N \frac{1}{N} \|\mathbf{x}_i - \mathbf{D}\mathbf{z}_i\|_2^2, \quad (21)$$

$$2933 \quad \text{s.t. } \forall j, \mathbf{z}_j = f_{TopK}(\mathbf{x}_j), \quad \|\mathbf{z}_j\|_0 = K. \quad (22)$$

2935 Since the fixed sparsity is a hard constraint that depends on \mathbf{z} alone (and not the data \mathbf{x}), it can
 2936 further be simplified as a sum-of-squares constraint: $\sum_j (\|\mathbf{z}_j\|_0 - K)^2 = 0$. We can use Lagrange
 2937 multipliers to reformulate it as an effective prior:

$$2939 \quad \arg \min_{\mathbf{D}, \mathbf{z}, \{\lambda_i\}} \sum_{i=1}^N \frac{1}{N} \left(\|\mathbf{x}_i - \mathbf{D}\mathbf{z}_i\|_2^2 + \lambda (\|\mathbf{z}_i\|_0 - K)^2 \right), \quad (23)$$

$$2943 \quad \text{s.t. } \forall j, \mathbf{z}_j = f_{TopK}(\mathbf{x}_j). \quad (24)$$

2944 The prior over \mathbf{z} for the above (effective) regularizer is:

$$2945 \quad \log P(\mathbf{z}_1, \dots, \mathbf{z}_N) \propto - \sum_{i=1}^N \lambda (\|\mathbf{z}_i\|_0 - K)^2 \quad (25)$$

$$2949 \quad \implies P(\mathbf{z}_1, \dots, \mathbf{z}_N) \propto \prod_{i=1}^N \exp(-\lambda (\|\mathbf{z}_i\|_0 - K)^2) \quad (26)$$

2952 Note that the above prior is finite for finite values of λ , but the overall objective optimizes over λ ,
 2953 resulting in a *hard* prior peaked at $\|\mathbf{z}_i\|_0 = K$ for each sample i .

2954 The factorization over samples i implies mutual independence of $\mathbf{z}_1, \dots, \mathbf{z}_n$: $P(\mathbf{z}_1, \dots, \mathbf{z}_n) =$
 2955 $\prod_{i=1}^N P(\mathbf{z}_i)$.

2957 As defined in Theorem I.2 (and restated here for convenience), let $S_i = \text{supp}(\mathbf{z}_i) = \{k : z_i^k >$
 2958 $0\}$, $n_i = |S_i| = \|\mathbf{z}_i\|_0$ denote the active indices and their number (sparsity) respectively.

2959 For individual samples \mathbf{z}_i , if we condition on the set of active indices S_i , the sparsity gets fixed since
 2960 $\|\mathbf{z}_i\|_0 = |S_i| = n_i$, and the distribution becomes constant:

$$2962 \quad P(\mathbf{z}_i | S_i) = C \quad (27)$$

$$2963 \quad \implies z_i^\mu | S_i \sim \begin{cases} U(0, \kappa) & \mu \in S_i \\ \delta_0 & \mu \notin S_i \end{cases}, \text{ and} \quad (28)$$

$$2966 \quad z_i^{\mu_1}, \dots, z_i^{\mu_{|S_i|}} | S_i \stackrel{\text{i.i.d.}}{\sim} U(0, \kappa) \text{ for } \mu_i \in S_i \quad (29)$$

2967 where C, κ are appropriate constants.
 2968

2969 Since $\{\mathbf{z}_i\}_i$ s are mutually independent, any measurable function of each is also independent. The
 indices of nonzero entries of \mathbf{z}_j , i.e., S_j is a measurable function since it is a map $S : \mathbb{R}_+^M \rightarrow 2^M$

which is discrete valued, and pre images of each value—a set of nonzero indices—are measurable since they equal the cartesian products of the measurable sets $\{z = 0\}, \{z > 0\}$ over all indices. Hence, S_1, \dots, S_n are also independent.

Since $S_i = g(\mathbf{z}_i)$ and the distribution of \mathbf{z}_i depends only on $n_i = \|\mathbf{z}_i\|_0$ (Eq. 25), the distribution of S_i will also depend only on n_i , becoming uniform when conditioned on n_i . In TopK SAE, $n_i = K$ is a constant. Therefore, each $S_i \sim U([M]^K)$, and together with independence argued above,

$$S_1, \dots, S_N \stackrel{\text{i.i.d.}}{\sim} U([M]^K) \quad (30)$$

This completes the proof for the priors of TopK SAE. \square

I.2.3 BATCHTOPK SAE

BatchTopK SAE (Bussmann et al. (2024)) is a modification of the TopK SAE. Instead of fixing sparsity like TopK, BatchTopK allows variable sparsity per input while fixing the mean sparsity over a batch at K . The objective function for BatchTopK SAE can equivalently be written as:

$$\arg \min_{\mathbf{D}, \mathbf{z}} \sum_{i=1}^N \frac{1}{N} \|\mathbf{x}_i - \mathbf{D}\mathbf{z}_i\|_2^2, \quad (31)$$

$$\text{s.t. } \forall j, \mathbf{z}_j = f_{TopK}(\mathbf{x}_j), \frac{1}{N} \sum_{j=1}^N \|\mathbf{z}_j\|_0 = K. \quad (32)$$

While BatchTopK imposes a mean sparsity per batch, for simplicity, we use the batch size to match the size of the entire dataset (WLOG). Smaller batch sizes can easily be incorporated by adding separate constraints, each over the entire batch (only leads to a change in constants—lagrange multipliers—in the analysis).

Following similar analysis as for TopK SAE (App. I.2.2), we can derive an equivalent prior over \mathbf{z} for BatchTopK SAE:

$$P(\mathbf{z}_1, \dots, \mathbf{z}_N) \propto \prod_{i=1}^N \exp(-\lambda |\|\mathbf{z}_i\|_0 - K|) \quad (33)$$

The sparse codes for different samples $\{\mathbf{z}_i\}_i$ are thus sampled i.i.d. from a distribution that only depends on the sparsity penalty. While this prior looks very similar to the prior of TopK SAE, the difference is that in TopK, the fixed sparsity constraint is imposed per sample, leading to a different Lagrange multiplier λ_i per sample to optimize over, while in BatchTopK, we have a common multiplier λ over all examples in a batch (with multiple batches, we will have one multiplier per batch), which is then optimized over to ensure that average sparsity per batch constraint is met.

Similar to the analysis for the TopK SAE, we get the following prior over different latents per sample:

$$z_i^\mu \mid S_i \sim \begin{cases} U(0, \kappa) & \mu \in S_i \\ \delta_0 & \mu \notin S_i \end{cases}, \text{ and} \quad (34)$$

$$z_i^{\mu_1}, \dots, z_i^{\mu_{|S_i|}} \mid S_i \stackrel{\text{i.i.d.}}{\sim} U(0, \kappa) \text{ for } \mu \in S_i \quad (35)$$

The active indices S_i are sampled uniformly conditioned on the number of active indices n_i :

$$S_i \mid n_i \sim U([M]^{n_i}) \quad (36)$$

The number of active latents n_i are themselves sampled i.i.d. (since $n_i = \tilde{g}(\mathbf{z}_i)$ and $\{\mathbf{z}_i\}_i$ are i.i.d.) from a distribution whose mean is fixed:

$$n_1, \dots, n_N \stackrel{\text{i.i.d.}}{\sim} P, \text{ s.t. } \mathbb{E}[n_i] = K \quad (37)$$

This completes the derivation for the BatchTopK prior. \square

3024 I.2.4 JUMPRELU SAE
30253026 JumpReLU SAE (Rajamanoharan et al. (2024)) is trained with the L_0 (pseudo-)norm regularizer.
3027 This leads to the following optimization problem:

3028
$$\arg \min_{\mathbf{D}, \mathbf{z}} \sum_{i=1}^N \frac{1}{N} (\|\mathbf{x}_i - \mathbf{D}\mathbf{z}_i\|_2^2 + \lambda \|\mathbf{z}_i\|_0)$$

3029
3030
3031 s.t. $\forall k, \mathbf{z}_k = f_{JumpReLU}(\mathbf{x}_k)$

3032 This objective is equivalent to the following prior over \mathbf{z} :
3033

3034
$$P(\mathbf{z}_1, \dots, \mathbf{z}_N) \propto \prod_{i=1}^N \exp \left(-\eta \|\mathbf{z}_i\|_0 \right) \quad (38)$$

3035
3036

3037 Noting the similarity with the TopK/ BatchTopK cases, we use the same analysis to derive the
3038 following conditions:

3039
$$z_i^\mu \mid S_i \sim \begin{cases} U(0, \kappa) & \mu \in S_i \\ \delta_0 & \mu \notin S_i \end{cases}, \text{ and} \quad (39)$$

3040
3041

3042
$$z_i^{\mu_1}, \dots, z_i^{\mu_{|S_i|}} \mid S_i \stackrel{\text{i.i.d.}}{\sim} U(0, \kappa) \text{ for } \mu \in S_i \quad (40)$$

3043
3044

$$S_i \mid n_i \sim U([M]^{n_i}) \quad (41)$$

3044 The number of active latents n_i are again i.i.d., but there is no constraint on the mean of the distribution
3045 (unlike BatchTopK which constrained the mean of n_i to equal K):

3046
$$n_1, \dots, n_N \stackrel{\text{i.i.d.}}{\sim} P, \quad (42)$$

3047

3048 which completes the analysis for JumpReLU SAE. \square 3049 J STATIONARITY MEASURES
30503052 LLM activations are empirically non-stationary across the sequence. We quantify the non-stationary
3053 nature by measuring autocorrelations and the U-statistic.3054 J.1 AUTOCORRELATION
30553057 We compute autocorrelation by selecting evenly spaced tokens across the sequence and measuring the
3058 cosine similarity between each token and tokens at various lags in the past. Specifically, for tokens at
3059 position t , we compute similarities to tokens at $t - w$ where lag w ranges from 5 to 20. This creates a
3060 heatmap where rows represent lag offsets and columns represent token positions.3061 For a stationary process, we expect the autocorrelation pattern to remain consistent across time—that
3062 is, the relationship between a token and its historical context should be similar regardless of position
3063 in the sequence. This would manifest as similar autocorrelation patterns repeating horizontally across
3064 token positions. In contrast, for a non-stationary process where representations evolve over time, we
3065 expect the autocorrelation patterns to vary systematically across positions, with columns showing
3066 different temporal dependency structures as the sequence progresses.3067 J.2 U-STATISTIC
30683069 We measure the effective dimensionality of LLM representations using a U-statistic based on pairwise
3070 cosine similarities.

3072
$$\text{U-stat}(t) = \frac{M^2 - M}{\|\mathbf{G}_t\|_F^2 - M} \quad (43)$$

3073
3074

3075 where $\|\mathbf{G}_t\|_F^2$ is the squared Frobenius norm of the Gram matrix. This quantity estimates the
3076 effective rank $1/\text{tr}(\mathbf{C}_t^2)$, where $\mathbf{C}_t = \mathbb{E}[\hat{\mathbf{x}}_t \hat{\mathbf{x}}_t^T]$ is the second moment matrix and $\hat{\mathbf{x}}_t$ is the normalized
3077 activation vector at time t . Under stationarity, U-stat remains constant. When representations evolve
over time, U-stat increases systematically as more orthogonal directions become active.

3078
3079

J.3 SURROGATE

3080 For the U-statistic and Autocorrelation metrics, we compare LLM activations to surrogate distributions
3081 that preserve certain statistical properties while removing temporal structure. We operate on
3082 representation vectors $\mathbf{X} \in \mathbb{R}^{B \times T \times d}$, where B denotes batch size, T denotes sequence length, and d
3083 denotes the model dimension.3084 **U-statistic surrogate (Fig. 2 a, e):** For each sequence $i \in \{1, \dots, B\}$, we construct the surrogate
3085 $\tilde{\mathbf{X}}_i$ by applying a random permutation $\pi_i : \{1, \dots, T\} \rightarrow \{1, \dots, T\}$ to the temporal positions:
3086

3087
$$\tilde{\mathbf{X}}_{i,t,:} = \mathbf{X}_{i,\pi_i(t),:} \quad \forall t \in \{1, \dots, T\}$$

3088 This preserves the marginal distribution of activations within each sequence while destroying temporal
3089 dependencies.
30903091 **Autocorrelation surrogate (Fig. 2 c, g):** Given the similarity matrix $\mathbf{S} \in \mathbb{R}^{T \times T}$ where $S_{ij} =$
3092 $\text{sim}(\mathbf{X}_{\cdot, i, \cdot}, \mathbf{X}_{\cdot, j, \cdot})$, we construct the surrogate similarity matrix $\tilde{\mathbf{S}}$ by replacing each diagonal with its
3093 mean:
3094

3095
$$\tilde{S}_{ij} = \bar{S}_k \quad \text{where } k = |i - j|, \quad \bar{S}_k = \frac{1}{T - k} \sum_{t=1}^{T-k} S_{t,t+k}$$

3096 This preserves the average correlation structure at each lag while removing position-specific temporal
3097 patterns.
30983099
3100
3101
3102
3103
3104
3105
3106
3107
3108
3109
3110
3111
3112
3113
3114
3115
3116
3117
3118
3119
3120
3121
3122
3123
3124
3125
3126
3127
3128
3129
3130
3131

3132 **K USE OF LARGE LANGUAGE MODELS**
31333134 LLMs were used in this work for the following:
31353136 • **Polish writing:** Although major parts of the writing were done by the authors themselves, LLMs
3137 (ChatGPT) were used to critique and iteratively improve the writing.
3138 • **Research ideation:** In the initial stages of the project, conversations with LLMs (ChatGPT, Gemini)
3139 aided in refining the overall storyline of the project, as well as to get feedback on theory sections.
3140 In all cases, LLM outputs were only used by the authors to refine their ideas.
3141
3142
3143
3144
3145
3146
3147
3148
3149
3150
3151
3152
3153
3154
3155
3156
3157
3158
3159
3160
3161
3162
3163
3164
3165
3166
3167
3168
3169
3170
3171
3172
3173
3174
3175
3176
3177
3178
3179
3180
3181
3182
3183
3184
3185