

SPARSE AUTOENCODERS REVEAL UNIVERSAL FEATURE SPACES ACROSS LARGE LANGUAGE MODELS

Anonymous authors

Paper under double-blind review

ABSTRACT

We investigate feature universality in large language models (LLMs), a research field that aims to understand how different models similarly represent concepts in the latent spaces of their intermediate layers. Demonstrating feature universality allows discoveries about latent representations to generalize across several models. However, comparing features across LLMs is challenging due to polysemanticity, in which individual neurons often correspond to multiple features rather than distinct ones. This makes it difficult to disentangle and match features across different models. To address this issue, we employ a method known as *dictionary learning* by using sparse autoencoders (SAEs) to transform LLM activations into more interpretable spaces spanned by neurons corresponding to individual features. After matching feature neurons across models via activation correlation, we apply representational space similarity metrics like Singular Value Canonical Correlation Analysis to analyze these SAE features across different LLMs. Our experiments reveal significant similarities in SAE feature spaces across various LLMs, providing new evidence for feature universality.

1 INTRODUCTION

Large language models (LLMs) have demonstrated remarkable capabilities across a wide range of natural language tasks (Bubeck et al., 2023; Naveed et al., 2024; Liu et al., 2023). However, understanding how these models represent and process information internally remains a significant challenge (Bereska and Gavves, 2024). A key question in this domain is whether different LLMs learn similar internal representations, or if each model develops its own unique way of encoding linguistic knowledge. This question of feature universality is crucial for several reasons: it impacts the generalizability of interpretability findings across models, it could accelerate the development of more efficient training techniques, and it may provide insights into safer and more controllable AI systems (Chughtai et al., 2023; Gurnee et al., 2024; Sharkey et al., 2024; Bricken et al., 2023).

Comparing features across LLMs is inherently difficult due to the phenomenon of polysemanticity, where individual neurons often correspond to multiple features rather than distinct ones Elhage et al. (2022). This superposition of features makes it challenging to disentangle and match representations across different models. To address these challenges, we propose a novel approach leveraging sparse autoencoders (SAEs) to transform LLM activations into more interpretable spaces. SAEs allow us to decompose the complex, superposed representations within LLMs into distinct features that are easier to analyze and compare across models (Makhzani and Frey, 2013; Cunningham et al., 2023; Bricken et al., 2023). We then apply specific representational space similarity metrics to these SAE features, enabling a rigorous quantitative analysis of feature universality across different LLMs (Klabunde et al., 2023). Our methodology involves several key steps: obtaining SAEs trained on the activations of multiple LLMs to extract interpretable features, developing and applying representational similarity metrics tailored for high-dimensional SAE feature spaces, and conducting extensive experiments comparing SAE features across models with varying architectures, sizes, and training regimes.

To verify the effectiveness of our approach, we present a comprehensive set of experiments and results. Our experimental approach involves solving two issues for representational space similarity measurement: the permutation issue, which involves aligning individual neurons, and the rotational issue, which involves comparing two spaces which have different basis that are rotationally-invariant. Additionally, we present an approach to filter out feature matches with low semantically-meaningful

similarity. We demonstrate high correlations between individual SAE features extracted from distinct models, providing evidence for feature universality. Furthermore, we show that semantically meaningful subspaces (e.g., features related to specific concepts, such as Emotions) exhibit even higher similarity across models. We also analyze how feature similarity varies across different layers of LLMs, revealing patterns in how representations evolve through the model depth.

This work opens up several exciting avenues for future research. Exploring whether universal features can be leveraged to improve transfer learning or model compression techniques may lead to more efficient training paradigms (Yosinski et al., 2014; Hinton et al., 2015). Additionally, examining the implications of feature universality for AI safety and control, particularly in identifying and mitigating potentially harmful representations, could contribute to the development of safer AI systems (Hendrycks et al., 2023). By shedding light on the internal representations of LLMs, our research contributes to the broader goal of developing more interpretable, efficient, and reliable AI systems (Barez et al., 2023). Understanding the commonalities and differences in how various models encode information is a crucial step toward building a more comprehensive theory of how large language models process and generate human language. **Our key contributions include:**

1. A novel framework for comparing LLM representations using SAEs and similarity metrics.
2. Empirical evidence supporting the existence of universal features across diverse LLMs.
3. A quantitative analysis of how feature similarity varies across layers and semantic domains.
4. A step-by-step approach to compare feature weight representations by paired features that differs from previous methods which compare feature activations.

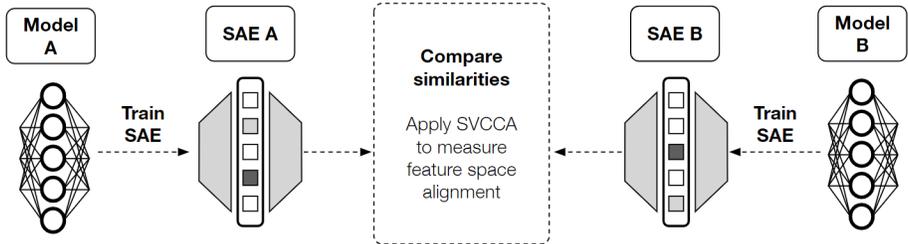


Figure 1: We train SAEs on LLMs, and then compare their SAE feature space similarity using metrics such as Singular Value Canonical Correlation Analysis (SVCCA).

2 BACKGROUND

Sparse Autoencoders. Sparse autoencoders (SAEs) are a type of autoencoder designed to learn efficient, sparse representations of input data by imposing a sparsity constraint on the hidden layer (Makhzani and Frey, 2013). The network takes in an input $\mathbf{x} \in \mathbb{R}^n$ and reconstructs it as an output $\hat{\mathbf{x}}$ using the equation $\hat{\mathbf{x}} = \mathbf{W}'\sigma(\mathbf{W}\mathbf{x} + \mathbf{b})$, where $\mathbf{W} \in \mathbb{R}^{h \times n}$ is the encoder weight matrix, \mathbf{b} is a bias term, σ is a nonlinear activation function, and \mathbf{W}' is the decoder matrix, which often uses the transpose of the encoder weights. SAE training aims to both encourage sparsity in the activations $\mathbf{h} = \sigma(\mathbf{W}\mathbf{x} + \mathbf{b})$ and to minimize the reconstruction loss $\|\mathbf{x} - \hat{\mathbf{x}}\|_2^2$.

The sparsity constraint encourages a large number of the neuron activations to remain inactive (close to zero) for any given input, while a few neurons, called *feature neurons*, are highly active. Recent work has replaced the regularization term by using the TopK activation function as the nonlinear activation function, which selects only the top K feature neurons with the highest activations (Gao et al., 2024), zeroing out the other activations. The active feature neurons tend to activate only for specific concepts in the data, promoting monosemanticity. Therefore, since there is a mapping from LLM neurons to SAE feature neurons that "translates" LLM activations to features, this method is a type of *dictionary learning* (Olshausen and Field, 1997).

SAE Feature Weight Spaces. By analyzing UMAPs of feature weights from an SAE trained on a middle residual stream layer of Claude 3 Sonnet, researchers discovered SAE feature spaces organized in neighborhoods of semantic concepts, allowing them to identify subspaces corresponding

to concepts such as biology and conflict (Templeton et al., 2024). Additionally, these concepts appear to be organized in hierarchical clusters, such as "disease" clusters that contain sub-clusters of specific diseases like flus. In Section §4.3, we measure the extent of similarity of these semantic feature spaces across LLMs.

Representational Space Similarity: Representational similarity measures typically compare neural network activations by assessing the similarity between activations from a consistent set of inputs (Klabunde et al., 2023). In this paper, we take a different approach by comparing the representations using the decoder weight matrices W' of SAEs, where the columns correspond to the feature neurons. We calculate a similarity score $m(W, W')$ for pairs of representations W in SAE A and W' from SAE B . We obtain two scores via the following two representational similarity metrics:

Singular Value Canonical Correlation Analysis (SVCCA). Canonical Correlation Analysis (CCA) seeks to identify pairs of canonical variables, \mathbf{u}_i and \mathbf{v}_i , from two sets of variables, $\mathbf{X} \in \mathbb{R}^{n \times d_1}$ and $\mathbf{Y} \in \mathbb{R}^{n \times d_2}$, that are maximally correlated with each other (Hotelling, 1936). Singular Value Canonical Correlation Analysis (SVCCA) (Raghu et al., 2017) enhances CCA by first applying Singular Value Decomposition (SVD) to both \mathbf{X} and \mathbf{Y} , which can be expressed as:

$$\mathbf{X} = \mathbf{U}_X \mathbf{S}_X \mathbf{V}_X^T, \quad \mathbf{Y} = \mathbf{U}_Y \mathbf{S}_Y \mathbf{V}_Y^T$$

where \mathbf{U}_X and \mathbf{U}_Y are the matrices containing the left singular vectors (informative directions), and \mathbf{S}_X and \mathbf{S}_Y are diagonal matrices containing the singular values. By projecting the original data onto these informative directions, noise is reduced. SVCCA then applies CCA on the transformed datasets, \mathbf{U}_X and \mathbf{U}_Y , resulting in correlation scores between the most informative components. These scores are averaged to obtain a similarity score.

Representational Similarity Analysis (RSA). Representational Similarity Analysis (RSA) (Kriegeskorte et al., 2008) works by first computing, for each space, a Representational Dissimilarity Matrix (RDM) $\mathbf{D} \in \mathbb{R}^{n \times n}$, where each element of the matrix represents the dissimilarity (or similarity) between every pair of data points within a space. The RDM essentially summarizes the pairwise distances between all possible data pairs in the feature space. A common distance metric used is the Euclidean distance. After an RDM is computed for each space, a correlation measurement like *Spearman's rank correlation coefficient* is applied to the pair of RDMs to obtain a similarity score. RSA has been applied to measure relational similarity of stimuli across brain regions and of activations across neural network layers (Klabunde et al., 2023).

SVCCA vs RSA. SVCCA focuses on the comparison of entire vector spaces by aligning the principal components (singular vectors) of two representation spaces and then computing the canonical correlations between them. In other words, it identifies directions in the space that carry the most variance and checks how well these align between two models or layers. On the other hand, RSA is primarily concerned with the structure of relationships, or *relational structure*, between different data points (like stimuli) within a representation space. It allows us to measure how well relations such as king-queen transfer across models. Thus, SVCCA focuses on comparing the aligned subspaces of two representation spaces, while RSA is more suitable for analyzing the structural similarity between representations at the level of pairwise relationships.

3 METHODOLOGY

Representational similarity is important for capturing the across-model generalizations of both: 1) Feature Spaces, which are spaces spanned by groups of features, and 2) Feature Relations. We compare SAEs trained on layer A_i from LLM A and layer B_j from LLM B . This is done for every layer pair. To compare latent spaces using our selected metrics, we have to solve both permutation and rotational alignment issues. For the permutation issue, we have to find a way to pair neuron weights. Typically, these metrics are performed by pairing activation vectors from two models according to common input instances. For example, input I_k is passed through models A and B , and then a vector of N_1 neuron activations from model A is compared to a paired vector of N_2 activations from model B . The common input instance creates individual mappings from one space to another, which is required for metrics like SVCCA and RSA (Klabunde et al., 2023).

In our paper, we do not pair activation vectors by input instances, but we pair the feature neuron weights themselves. This means we have to find individual pairings of "similar" neurons in SAE A to SAE B . However, we do not know which features map to which features; the order of these features are permuted in the weight matrices, and some features may not have a corresponding partner. Therefore, we have to pairwise match them via a correlation metric. For the rotational issue, models typically develop their own basis for latent space representations, so features may be similar relation-wise across spaces, but not rotation-wise due to differing basis. To solve this issue, we employ rotation-invariant metrics.

Assessing Scores with Baseline. We follow the approach of Kriegeskorte et al. (2008) to randomly pair features to obtain a baseline score. We compare the score of the features paired by correlation (which we refer to as the "paired features") with the average score of N runs of randomly paired features to obtain a p-value score. If the p-value is less than 0.05, the metric suggests that the similarity is statistically meaningful.

In summary, the steps to carry out our similarity comparison experiments are given below, and steps 1 to 3 are illustrated in Figure 8:

1. For the permutation issue: Get mean activation correlation of two decoder SAEs for similar LLMs. Pair each feature A with its max activation correlated feature from model B .
2. Rearrange the order of features in each matrix to pair them row by row.
3. For the rotational issue: Apply various rotation-invariant metrics to get a *Paired Score*.
4. Using the same metrics, obtain the similarity scores of N *random pairings* to estimate a null distribution. Obtain a p-value of where the paired score falls in the null distribution to determine statistical significance.

Semantic Subspace Matching Experiments. For these experiments, we first identify semantically similar subspaces, and then compare their similarity. This is summarized by the following steps:

1. We find a group of related features in each model by searching for features which activate highly on the top 5 samples' tokens that belong to a *concept category*. For example, the concept of "emotions" contains keywords such as "rage". Since the correlations are calculated by tokens, we use single-token keywords. If a keyword from a concept appears at least once, that feature is included. The keywords in each category are given in Appendix D.
2. We obtain mappings between these subsets of features using max activation correlation.
3. We characterize the relationships of features within these groups (eg. pairwise differences between features), and compare these relationships across models using similarity metrics.

Assessing Scores with Baseline. We use two types of tests which compare the paired score to a null distribution. Each test examines that the score is rare under a certain null distribution assumption:

1. Test 1: We compare the score of a semantic subspace mapping to the mean score of randomly shuffled pairings of the same subspaces. This test shows just comparing the subspace of features is not enough; the features must be paired.
2. Test 2: We compare the score of a semantic subspace mapping to a mean score of mappings between randomly selected feature subsets of the same sizes as the semantic subspaces. This uses a correlation matrix only between the two subsets, not between the entire space. This shows that the high similarity score does not hold for any two subspaces of the same size.

Metrics for the Permutation Issue: These metrics measure individual, local similarity.

Mean Activation Correlation. We take the mean activation correlation between neurons, following the approach of Bricken et al. (2023). We pass a dataset of samples through both models. Then, for each feature in model A , we find its activation correlation with every feature in model B , using tokens as the common instance pairing. For each feature in model A , we find its highest correlated feature in model B . We pair these features together when conducting our experiments. We refer to scores obtained by pairing via mean activation correlation as *Paired Scores*.

Filter by Top Tokens. We notice that there are many "non-concept mappings" with "non-concept features" that have top 5 token activations on end-of-text / padding tokens, spaces, new lines, and

punctuation; these non-concept mappings greatly bring our scores down as they are not accurately pairing features representing concepts, as we describe further in Appendix B. By filtering max activation correlation mappings to remove non-concept mapping features, we significantly raise the similarity scores. We also only keep mappings that share at least one keyword in their top 5 token activations. The list of non-concept keywords is given in the Appendix B.

Filter Many-to-1 Mappings. We find that some of these mappings are many-to-1, which means that more than one feature maps to the same feature. We found that removing many-to-1 mappings slightly increased the scores, and we discuss possible reasons why this may occur in Appendix A. However, the scores still show high similarity even with the inclusion of many-to-1 mappings, as seen in Figures 2. We show only scores of 1-to-1 mappings (filtering out the many-to-1 mappings) in the main text, and show scores of many-to-1 mappings in Appendix A. We use features from the SAEs of Pythia-70m and Gemma-1-2B as "many" features which map onto "one" feature in the SAEs of Pythia-160m and Gemma-2-2B.

Metrics for the Rotational Issue: Rotation-invariant metrics measure global similarity after feature pairing. We apply metrics which measure: 1) How well subspaces align using **SVCCA**, and 2) How similar feature relations like king-queen are using **RSA**. We set the inner similarity function using Pearson correlation the outer similarity function using Spearman correlation, and compute the RDM using Euclidean Distance.

4 EXPERIMENTS

4.1 EXPERIMENTAL SETUP

LLM Models. We compare LLMs that use the Transformer model architecture (Vaswani et al., 2017). We compare models that use the same tokenizer because the mean activation correlation pairing relies on comparing two activations using the same tokens. We compare the residual streams of Pythia-70m, which has 6 layers and 512 dimensions, vs Pythia-160m, which has 12 layers and 768 dimensions (Biderman et al., 2023). We compare the residual streams of Gemma-1-2B, which has 18 layers and 2048 dimensions, to Gemma-2-2B, which has 26 layers and 2304 dimensions Team et al. (2024a;b).¹

SAE Models. For Pythia, we use pretrained SAEs with 32768 feature neurons trained on residual stream layers from Eleuther (EleutherAI, 2023). For Gemma, we use pretrained SAEs with 16384 feature neurons trained on residual stream layers (Lieberum et al., 2024) We access pretrained Gemma-1-2B SAEs for layers 0, 6, 10, 12, and 17, and Gemma-2-2B SAEs for layers 0 to 26, through the SAELens library (Bloom and Chanin, 2024).

Datasets. We obtain SAE activations using OpenWebText, a dataset that scraps internet content (Gokaslan and Cohen, 2019). We use 100 samples with a max sequence length of 300 for Pythia for a total of 30k tokens, and we use 150 samples with a max sequence length of 150 for Gemma for a total of 22.5k tokens. Top activating dataset tokens were also obtained using OpenWebText samples.

Computing Resources. We run experiments on an A100 GPU.

4.2 SAE SPACE SIMILARITY

In summary, for almost all layer pairs after layer 0, we find the p-value of SAE experiments for SVCCA, which measures rotationally-invariant alignment, to be low; all of them are between 0.00 to 0.01 using 100 samples.² While RSA also passes many p-value tests, their values are much lower, suggesting a relatively weaker similarity in terms of pairwise differences. It is possible that many features are not represented in just one layer; hence, a layer from one model can have similarities to multiple layers in another model. We also note that the first layer of every model, layer 0, has a very high p-value when compared to every other layer, including when comparing layer 0 from another model. This may be because layer 0 contains few discernible, meaningful, and comparable features.

¹Gemma-2 introduces several architectural improvements over Gemma-1 (Team et al., 2024a;b).

²Only 100 samples were used as we found the variance to be low for the null distribution, and that the mean was similar for using 100 vs 1k samples.

As shown in Figure 9 for Pythia and Figure 10 for Gemma in Appendix E, we also find that mean activation correlation does not always correlate with the global similarity metric scores. For instance, a subset of feature pairings with a moderately high mean activation correlation (eg. 0.6) may yield a low SVCCA score (eg. 0.03).

Because LLMs, even with the same architecture, can learn many different features, we do not expect entire feature spaces to be highly similar; instead, we hypothesize there may be highly similar *subspaces*. In other words, there are many similar features universally found across SAEs, but not every feature learned by different SAEs is the same. Our experimental results support this hypothesis.

Pythia-70m vs Pythia-160m . In Figure 4, we compare Pythia-70m, which has 6 layers, to Pythia-160m, which has 12 layers. When compared to Figures 11 and 12 in Appendix E, for all layers pairs (excluding Layers 0), and for both SVCCA and RSA scores, almost all the p-values of the residual stream layer pairs are between 0% to 1%, indicating that the feature space similarities are statistically meaningful. In particular, Figure 4 shows that Layers 2 and 3, which are middle layer of Pythia-70m, are more similar by SVCCA and RSA scores to middle layers 4 to 7 of Pythia-160m compared to other layers. Figure 4 also shows that Layer 5, the last residual stream layer of Pythia-70m, is more similar to later layers of Pythia-160m than to other layers. The number of features after filtering non-concept features and many-to-1 features are given in Tables 5 and 6 in Appendix E.

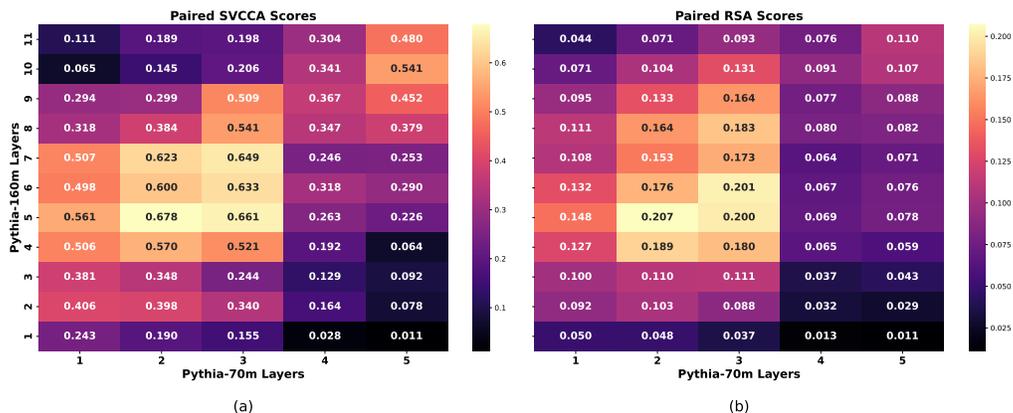


Figure 2: (a) SVCCA and (b) RSA 1-1 paired scores of SAEs for layers in Pythia-70m vs layers in Pythia-160m. We find that middle layers have the most similarity with one another (as shown by the high-similarity block spanned by layers 1 to 3 in Pythia-70m and Layers 4 to 7 in Pythia-160m). We exclude layers 0, as we observe they always have non-statistically significant similarity. The 1-1 scores are slightly higher for most of the Many-to-1 scores shown in Figure 4, and the SVCCA score at L2 vs L3 for 70m vs 160m is much higher.

Gemma-1-2B vs Gemma-2-2B. As shown by the paired SVCCA and RSA results for these models in Figure 6, we find that using SAEs, we can detect highly similar features in layers across Gemma-1-2B and Gemma-2-2B. Compared to mean random pairing scores and p-values in Figures 13 and 14 in Appendix E, for all layers pairs (excluding Layers 0), and for both SVCCA and RSA scores, almost all the p-values of the residual stream layer pairs are between 0% to 1%, indicating that the feature space similarities are statistically meaningful. The number of features after filtering non-concept features and many-to-1 features are given in Tables 7 and 8 in Appendix E.

Rather than just finding that middle layers are more similar to middle layers, we also find that they are the best at representing concepts. This is consistent with work by (Rimsky et al., 2024), which find the best steering vectors in middle layers for steering behaviors such as sycophancy. Other papers find that middle layers tend to provide better directions for eliminating refusal behavior (Arditi et al., 2024) and work well for representing goals and finding steering vectors in policy networks (Mini et al., 2023). In fact, we find that early layers do not represent universal features well, perhaps due to lacking content, while later layers do not represent them as well as middle layers, perhaps due to being too specific/complex.

324
325
326
327
328
329
330
331
332
333
334
335
336
337
338
339
340
341
342
343
344
345
346
347
348
349
350
351
352
353
354
355
356
357
358
359
360
361
362
363
364
365
366
367
368
369
370
371
372
373
374
375
376
377

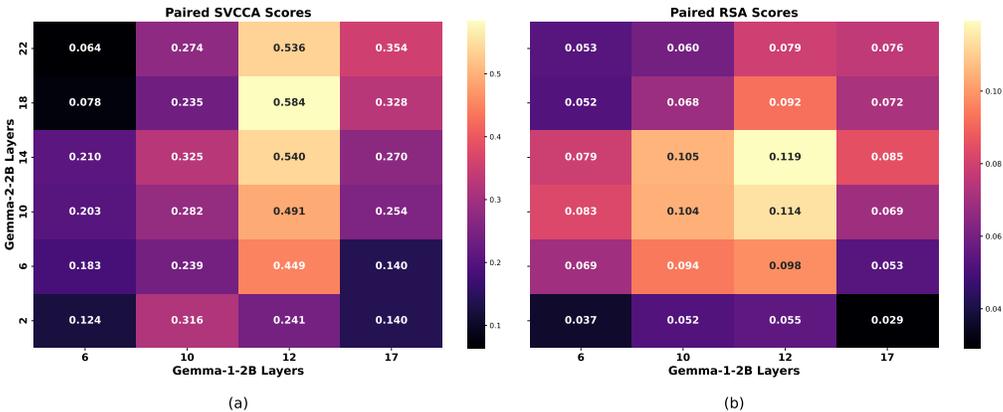


Figure 3: (a) SVCCA and (b) RSA 1-1 paired scores of SAEs for layers in Gemma-1-2B vs layers in Gemma-2-2B. Middle layers have the best performance. The later layer 17 in Gemma-1 is more similar to later layers in Gemma-2. Early layers like Layer 2 in Gemma-2 have very low similarity. We exclude layers 0, as we observe they always have non-statistically significant similarity. The 1-1 scores are slightly higher for most of the Many-to-1 scores shown in Figure 6.

Table 1: SVCCA scores and random mean results for 1-1 semantic subspaces of L3 of Pythia-70m vs L5 of Pythia-160m. P-values are taken for 1000 samples in the null distribution.

| Concept Subspace | Number of Features | Paired Mean | Random Shuffling Mean | p-value |
|------------------|--------------------|-------------|-----------------------|---------|
| Time | 228 | 0.59 | 0.05 | 0.00 |
| Calendar | 126 | 0.65 | 0.07 | 0.00 |
| Nature | 46 | 0.50 | 0.12 | 0.00 |
| Countries | 32 | 0.72 | 0.14 | 0.00 |
| People/Roles | 31 | 0.50 | 0.15 | 0.00 |
| Emotions | 24 | 0.83 | 0.15 | 0.00 |

4.3 SIMILARITY OF SEMANTICALLY MATCHED SAE FEATURE SPACES

We find that for all layers and for all concept categories, Test 2 described in §3 is passed. Thus, we only report specific results for Test 1 in Tables 1 and 2. Overall, in both Pythia and Gemma models and for many concept categories, we find that semantic subspaces are more similar to one another than non-semantic subspaces.

Pythia-70m vs Pythia-160m. We compare every layer of Pythia-70m to every layer of Pythia-160m for several concept categories. While many layer pairs have similar semantic subspaces. Middle layers appear to have the semantic subspaces with the highest similarity. Table 1 demonstrates one example of this by comparing the SVCCA score for layer 3 to layer 5 of Pythia-160m, which shows high similarity for semantic subspaces across models, as the concept categories all pass Test 1, having p-values below 0.05. We show the scores for other layer pairs and categories in Figures 15 and 17 in Appendix E, which include RSA scores in Table 9.

Gemma-1-2B vs Gemma-2-2B. In Table 2, we compare L12 of Gemma-1-2B vs L14 of Gemma-2-2B. As shown in Figure 6, this layer pair has a very high similarity for most concept spaces; as such, they likely have high similar semantically-meaningful feature subspaces. Notably, not all concept group are not highly similar; for instance, unlike in Pythia, the Country concept group does not pass the Test 1 as it has a p-value above 0.05. We show the scores for other layer pairs and categories in Figures 19 and 21 in Appendix E.

Table 2: SVCCA scores and random mean results for 1-1 semantic subspaces of L12 of Gemma-1-2B vs L14 of Gemma-2-2B. P-values are taken for 1000 samples in the null distribution.

| Concept Subspace | Number of Features | Paired SVCCA | Random Shuffling Mean | p-value |
|------------------|--------------------|--------------|-----------------------|---------|
| Time | 228 | 0.46 | 0.06 | 0.00 |
| Calendar | 105 | 0.54 | 0.07 | 0.00 |
| Nature | 51 | 0.30 | 0.11 | 0.01 |
| Countries | 21 | 0.39 | 0.17 | 0.07 |
| People/Roles | 36 | 0.66 | 0.15 | 0.00 |
| Emotions | 35 | 0.60 | 0.12 | 0.00 |

5 RELATED WORK

Superposition and Sparse Autoencoders. Superposition is a phenomenon in which a model, in response to the issue of having to represent more features than it has parameters, learns feature representations distributed across many parameters (Elhage et al., 2022). This causes its parameters, or neurons, to be polysemantic, which means that each neuron is involved in representing more than one feature. This leads to issues for interpretability, which aims to disentangle and cleanly identify features in models in order to manipulate them, such as by steering a model away from deception (Templeton et al., 2024). To address the issue of polysemanticity, Sparse Autoencoders (SAEs) have been applied to disentangle an LLM’s polysemantic neuron activations into monosemantic "feature neurons", which are encouraged to represent an isolated concept (Makhzani and Frey, 2013; Cunningham et al., 2023; Gao et al., 2024; Rajamanoharan et al., 2024a;b). Features interactions in terms of circuits have also been studied (Marks et al., 2024).

Feature Universality. To the best of our knowledge, only Bricken et al. (2023) has done a quantitative study on individual SAE feature similarity for two 1-layer toy models, finding that individual SAE feature correlations are stronger than individual LLM neuron correlations; however, this study did not analyze the global properties of feature spaces. Gurnee et al. (2024) found evidence of universal neurons across language models by applying pairwise correlation metrics, and taxonomized families of neurons based on patterns of downstream, functional effects. O’Neill et al. (2024) discover "feature families", which represent related hierarchical concepts, in SAEs across two different datasets. Recent work has shown that as vision and language models are trained with more parameters and with better methods, their representational spaces converge to more similar representations (Huh et al., 2024).

Representational Space Similarity. Previous work has studied neuron activation spaces by utilizing metrics to compare the geometric similarities of representational spaces (Raghu et al., 2017; Wang et al., 2018; Kornblith et al., 2019; Klabunde et al., 2024; 2023; Kriegeskorte et al., 2008). It was found that even models with different architectures may share similar representation spaces, hinting at feature universality. However, these techniques have yet to be applied to the feature spaces of sparse autoencoders trained on LLMs. Moreover, these techniques compare the similarity of paired input activations. Our work differs as it compares the similarity of paired *feature weights*. Essentially, previous works do not show that one can match LLMs by features that both LLMs find in common.

Feature Manifolds. In real world data, features may lie on a manifold, where nearby features respond to similar data. (Olah and Batson, 2023; Bricken et al., 2023). Previous work has discovered the existence of features corresponding to months and days of the week arranged in a circular structure in SAE feature spaces across models (Engels et al., 2024). Studying how feature arrangements generalize across models may shed light on how features lie on a manifold.

Mechanistic Interpretability. Recent work has made notable progress in neuron interpretation (Foote et al., 2023; Garde et al., 2023) and interpreting the interactions between Attention Heads and MLPs (Neo et al., 2024). Other work in mechanistic interpretability has traditionally focused on 'circuits'-style analysis (Elhage et al., 2021), as well as the use of steering vectors to direct models at inference using identified representations (Zou et al., 2023; Turner et al., 2023; Bereska and Gavves, 2024). These approaches have been combined with SAEs to steer models in the more interpretable SAE feature space (Nanda and Conmy, 2024). Steering vectors have been found across models;

432 for instance, Arditi et al. (2024) found "Refusal Vectors" that steered a model to refuse harmful
433 instructions across 13 different models, including Gemma and Llama-2.

434 **AI Safety.** Advanced AI systems may not be well aligned with the values of their designers (Ngo
435 et al., 2022), due to phenomena such as goal misgeneralization (Shah et al., 2022; Langosco et al.,
436 2022), deceptive alignment (Hubinger et al., 2024; 2021), sycophancy (Sharma et al., 2023), and
437 reward tampering (Denison et al., 2024). SAEs may be able to find AI safety-related features across
438 models; using SAEs to reliably find universal capabilities, especially in combination with control
439 methods such as steering, may help provide solutions to the above phenomena.

441 6 CONCLUSION

442
443 In this study, we address the underexplored domain of feature universality in sparse autoencoders
444 (SAEs) trained on LLMs with multiple layers and across multiple model pairs. To achieve these
445 insights, we developed a novel methodology to pair features with high mean activation correlations,
446 and then assess their global similarity of the resulting subspaces using SVCCA and RSA. Our findings
447 reveal a high degree of similarity in SAE feature spaces across various models. Furthermore, our
448 research reveals that subspaces of features associated with specific semantic concepts, such as calendar
449 or people tokens, demonstrate remarkably high similarity across different models. This suggests that
450 certain semantic feature subspaces are universally encoded across varied LLM architectures. Our
451 work enhances understanding of how LLMs represent information at a fundamental level, opening
452 new avenues for further research into model interpretability and the transferability of learned features.

453 LIMITATIONS

454
455 We perform analysis on available pre-trained SAEs for similar LLMs that use the same tokenizer, and
456 preferably on SAEs with the same number of features to match them. However, there are a limited
457 number of pre-trained SAEs on large models. Our study is also limited by the inherent challenges
458 in interpreting and comparing high-dimensional feature spaces. The methods we used to measure
459 similarity, while robust, may not capture all nuances of feature interactions and representations within
460 the SAEs. Additionally, our analysis is constrained by the computational resources available, which
461 may have limited the scale and depth of our comparisons. The generalizability of our findings to SAEs
462 trained on more diverse datasets or specialized domain-specific LLMs remains an open question.
463 Furthermore, the static nature of our analysis does not account for potential temporal dynamics in
464 feature representation that might occur during the training process of both LLMs and SAEs.

465 FUTURE WORK

466
467 We performed similar analysis on the Pythia and Gemma LLMs that the SAEs in this paper were
468 trained on, and found that they have some similarity, but much lower than the SAEs (around 0.1-0.2
469 for all layer pairs) with no patterns showing that the middle layers were the most similar; however,
470 the LLMs did not have corresponding weights for the "residual stream", which is an addition of
471 activations from the MLP and attention head layers, so we compared the MLPs of LLMs. Additionally,
472 unlike the SAE dimensions, the LLM layers they were trained on had different dimensions (eg. 512
473 for Pythia-70m vs 768 Pythia-160m), making feature matching less trivial. Thus, future work could
474 study the comparison of LLM similarity by feature weights to SAEs. Additional experiments could
475 include using ground-truth feature experiments for LLMs that partially share ground-truth features to
476 see if SAEs are modeling the same features, and performing functional similarity experiments.

477 REPRODUCIBILITY

478
479 We include the code and instructions for reproducing our main results in the Supplementary Materials.

480 REFERENCES

481
482
483 Andy Arditi, Oscar Obeso, Aquib Syed, Daniel Paleka, Nina Rimsky, Wes Gurnee, and Neel Nanda.
484 Refusal in language models is mediated by a single direction, 2024.

- 486 Fazl Barez, Hosien Hasanbieg, and Alesandro Abbate. System iii: Learning with domain knowledge
487 for safety constraints, 2023.
488
- 489 Leonard Bereska and Efstratios Gavves. Mechanistic interpretability for ai safety – a review, 2024.
490 URL <https://arxiv.org/abs/2404.14082>.
- 491 Stella Biderman, Hailey Schoelkopf, Quentin Anthony, Herbie Bradley, Kyle O’Brien, Eric Hallahan,
492 Mohammad Aflah Khan, Shivanshu Purohit, USVSN Sai Prashanth, Edward Raff, Aviya Skowron,
493 Lintang Sutawika, and Oskar van der Wal. Pythia: A suite for analyzing large language models
494 across training and scaling, 2023. URL <https://arxiv.org/abs/2304.01373>.
- 495 Joseph Bloom and David Chanin. Saelens. <https://github.com/jbloomAus/SAELens>,
496 2024.
497
- 498 Trenton Bricken, Adly Templeton, Joshua Batson, Brian Chen, Adam Jermyn, Tom Conerly, Nick
499 Turner, Cem Anil, Carson Denison, Amanda Askell, Robert Lasenby, Yifan Wu, Shauna Kravec,
500 Nicholas Schiefer, Tim Maxwell, Nicholas Joseph, Zac Hatfield-Dodds, Alex Tamkin, Karina
501 Nguyen, Brayden McLean, Josiah E Burke, Tristan Hume, Shan Carter, Tom Henighan, and
502 Christopher Olah. Towards monosemanticity: Decomposing language models with dictionary
503 learning. *Transformer Circuits Thread*, 2023. <https://transformer-circuits.pub/2023/monosemantic-features/index.html>.
504
- 505 Sébastien Bubeck, Varun Chandrasekaran, Ronen Eldan, Johannes Gehrke, Eric Horvitz, Ece Kamar,
506 Peter Lee, Yin Tat Lee, Yuanzhi Li, Scott Lundberg, Harsha Nori, Hamid Palangi, Marco Tulio
507 Ribeiro, and Yi Zhang. Sparks of artificial general intelligence: Early experiments with gpt-4,
508 2023. URL <https://arxiv.org/abs/2303.12712>.
- 509 Bilal Chughtai, Lawrence Chan, and Neel Nanda. A toy model of universality: Reverse engineer-
510 ing how networks learn group operations, 2023. URL <https://arxiv.org/abs/2302.03025>.
511
512
- 513 Hoagy Cunningham, Aidan Ewart, Logan Riggs, Robert Huben, and Lee Sharkey. Sparse autoen-
514 coders find highly interpretable features in language models, 2023. URL <https://arxiv.org/abs/2309.08600>.
515
- 516 Carson Denison, Monte MacDiarmid, Fazl Barez, David Duvenaud, Shauna Kravec, Samuel Marks,
517 Nicholas Schiefer, Ryan Soklaski, Alex Tamkin, Jared Kaplan, Buck Shlegeris, Samuel R. Bowman,
518 Ethan Perez, and Evan Hubinger. Sycophancy to Subterfuge: Investigating Reward-Tampering
519 in Large Language Models, June 2024. URL <http://arxiv.org/abs/2406.10162>.
520 arXiv:2406.10162 [cs].
521
- 522 EleutherAI. Sparse autoencoders (sae) repository. <https://github.com/EleutherAI/sae>,
523 2023.
- 524 Nelson Elhage, Neel Nanda, Catherine Olsson, Tom Henighan, Nicholas Joseph, Ben Mann, Amanda
525 Askell, Yuntao Bai, Anna Chen, Tom Conerly, Nova DasSarma, Dawn Drain, Deep Ganguli,
526 Zac Hatfield-Dodds, Danny Hernandez, Andy Jones, Jackson Kernion, Liane Lovitt, Kamal
527 Ndousse, Dario Amodei, Tom Brown, Jack Clark, Jared Kaplan, Sam McCandlish, and Chris
528 Olah. A mathematical framework for transformer circuits. *Transformer Circuits Thread*, 2021.
529 <https://transformer-circuits.pub/2021/framework/index.html>.
- 530 Nelson Elhage, Tristan Hume, Catherine Olsson, Nicholas Schiefer, Tom Henighan, Shauna
531 Kravec, Zac Hatfield-Dodds, Robert Lasenby, Dawn Drain, Carol Chen, Roger Grosse,
532 Sam McCandlish, Jared Kaplan, Dario Amodei, Martin Wattenberg, and Christopher Olah.
533 Toy models of superposition. *Transformer Circuits Thread*, 2022. https://transformer-circuits.pub/2022/toy_model/index.html.
534
- 535 Joshua Engels, Isaac Liao, Eric J. Michaud, Wes Gurnee, and Max Tegmark. Not all language model
536 features are linear, 2024. URL <https://arxiv.org/abs/2405.14860>.
537
- 538 Alex Foote, Neel Nanda, Esben Kran, Ioannis Konstas, Shay Cohen, and Fazl Barez. Neuron to
539 graph: Interpreting language model neurons at scale, 2023. URL <https://arxiv.org/abs/2305.19911>.

- 540 Leo Gao, Tom Dupré la Tour, Henk Tillman, Gabriel Goh, Rajan Troll, Alec Radford, Ilya Sutskever,
541 Jan Leike, and Jeffrey Wu. Scaling and evaluating sparse autoencoders, 2024. URL <https://arxiv.org/abs/2406.04093>.
542
543
- 544 Albert Garde, Esben Kran, and Fazl Barez. Deepdecipher: Accessing and investigating neuron
545 activation in large language models, 2023. URL <https://arxiv.org/abs/2310.01870>.
546
- 547 Aaron Gokaslan and Vanya Cohen. Openwebtext corpus. [http://Skyllion007.github.io/
548 OpenWebTextCorpus](http://Skyllion007.github.io/OpenWebTextCorpus), 2019. Accessed: [Insert date you accessed the resource].
- 549 Wes Gurnee, Theo Horsley, Zifan Carl Guo, Tara Rezaei Kheirkhah, Qinyi Sun, Will Hathaway,
550 Neel Nanda, and Dimitris Bertsimas. Universal neurons in gpt2 language models, 2024. URL
551 <https://arxiv.org/abs/2401.12181>.
552
- 553 Dan Hendrycks, Mantas Mazeika, and Thomas Woodside. An overview of catastrophic ai risks, 2023.
554 URL <https://arxiv.org/abs/2306.12001>.
555
- 556 Geoffrey E. Hinton, Oriol Vinyals, and Jeffrey Dean. Distilling the knowledge in a neural network.
557 *ArXiv*, abs/1503.02531, 2015. URL [https://api.semanticscholar.org/CorpusID:
558 7200347](https://api.semanticscholar.org/CorpusID:7200347).
- 559 Harold Hotelling. Relations between two sets of variates. *Biometrika*, 28(3/4):321–377, 1936.
560
- 561 Evan Hubinger, Chris van Merwijk, Vladimir Mikulik, Joar Skalse, and Scott Garrabrant. Risks
562 from Learned Optimization in Advanced Machine Learning Systems, December 2021. URL
563 <http://arxiv.org/abs/1906.01820>. arXiv:1906.01820 [cs].
564
- 565 Evan Hubinger, Carson Denison, Jesse Mu, Mike Lambert, Meg Tong, Monte MacDiarmid, Tamera
566 Lanham, Daniel M. Ziegler, Tim Maxwell, Newton Cheng, Adam Jermyn, Amanda Askeel, Ansh
567 Radhakrishnan, Cem Anil, David Duvenaud, Deep Ganguli, Fazl Barez, Jack Clark, Kamal
568 Ndousse, Kshitij Sachan, Michael Sellitto, Mrinank Sharma, Nova DasSarma, Roger Grosse,
569 Shauna Kravec, Yuntao Bai, Zachary Witten, Marina Favaro, Jan Brauner, Holden Karnofsky,
570 Paul Christiano, Samuel R. Bowman, Logan Graham, Jared Kaplan, Sören Mindermann, Ryan
571 Greenblatt, Buck Shlegeris, Nicholas Schiefer, and Ethan Perez. Sleeper Agents: Training
572 Deceptive LLMs that Persist Through Safety Training, January 2024. URL [http://arxiv.
573 org/abs/2401.05566](http://arxiv.org/abs/2401.05566). arXiv:2401.05566 [cs].
- 574 Minyoung Huh, Brian Cheung, Tongzhou Wang, and Phillip Isola. The platonic representation
575 hypothesis, 2024.
576
- 577 Max Klabunde, Tobias Schumacher, Markus Strohmaier, and Florian Lemmerich. Similarity of
578 neural network models: A survey of functional and representational measures, 2023. URL
579 <https://arxiv.org/abs/2305.06329>.
- 580 Max Klabunde, Tassilo Wald, Tobias Schumacher, Klaus Maier-Hein, Markus Strohmaier, and Florian
581 Lemmerich. Resi: A comprehensive benchmark for representational similarity measures, 2024.
582 URL <https://arxiv.org/abs/2408.00531>.
583
- 584 Simon Kornblith, Mohammad Norouzi, Honglak Lee, and Geoffrey Hinton. Similarity of neural
585 network representations revisited, 2019. URL <https://arxiv.org/abs/1905.00414>.
586
- 587 Nikolaus Kriegeskorte, Marieke Mur, and Peter Bandettini. Representational similarity analysis –
588 connecting the branches of systems neuroscience. *Frontiers in Systems Neuroscience*, 2:4, 2008.
589 doi: 10.3389/neuro.06.004.2008. URL [https://www.frontiersin.org/articles/
590 10.3389/neuro.06.004.2008/full](https://www.frontiersin.org/articles/10.3389/neuro.06.004.2008/full).
- 591 Lauro Langosco Di Langosco, Jack Koch, Lee D. Sharkey, Jacob Pfau, and David Krueger. Goal
592 Misgeneralization in Deep Reinforcement Learning. In *Proceedings of the 39th International
593 Conference on Machine Learning*, pages 12004–12019. PMLR, June 2022. URL [https://
proceedings.mlr.press/v162/langosco22a.html](https://proceedings.mlr.press/v162/langosco22a.html). ISSN: 2640-3498.

- 594 Tom Lieberum, Senthoran Rajamanoharan, Arthur Conmy, Lewis Smith, Nicolas Sonnerat, Vikrant
595 Varma, János Kramár, Anca Dragan, Rohin Shah, and Neel Nanda. Gemma scope: Open sparse
596 autoencoders everywhere all at once on gemma 2, 2024. URL [https://arxiv.org/abs/
597 2408.05147](https://arxiv.org/abs/2408.05147).
- 598 Yiheng Liu, Tianle Han, Siyuan Ma, Jiayue Zhang, Yuanyuan Yang, Jiaming Tian, Hao He, Antong
599 Li, Mengshen He, Zhengliang Liu, Zihao Wu, Lin Zhao, Dajiang Zhu, Xiang Li, Ning Qiang,
600 Dingang Shen, Tianming Liu, and Bao Ge. Summary of chatgpt-related research and perspective
601 towards the future of large language models. *Meta-Radiology*, 1(2):100017, September 2023. ISSN
602 2950-1628. doi: 10.1016/j.metrad.2023.100017. URL [http://dx.doi.org/10.1016/j.
604 metrad.2023.100017](http://dx.doi.org/10.1016/j.
603 metrad.2023.100017).
- 605 Alireza Makhzani and Brendan Frey. k-sparse autoencoders. *arXiv preprint arXiv:1312.5663*, 2013.
- 606 Samuel Marks, Can Rager, Eric J. Michaud, Yonatan Belinkov, David Bau, and Aaron Mueller.
607 Sparse feature circuits: Discovering and editing interpretable causal graphs in language models,
608 2024.
- 609 Ulisse Mini, Peli Grietzer, Mrinank Sharma, Austin Meek, Monte MacDiarmid, and Alexander Matt
610 Turner. Understanding and Controlling a Maze-Solving Policy Network, October 2023. URL
611 <http://arxiv.org/abs/2310.08043>. arXiv:2310.08043 [cs].
- 612 Neel Nanda and Arthur Conmy. Progress update #1 from the gdm mech interp team: Full
613 update. [https://www.alignmentforum.org/posts/C5KAZQib3bzzpeyrg/
614 progress-update-1-from-the-gdm-mech-interp-team-full-update#
615 Activation_Steering_with_SAEs](https://www.alignmentforum.org/posts/C5KAZQib3bzzpeyrg/progress-update-1-from-the-gdm-mech-interp-team-full-update#Activation_Steering_with_SAEs), 2024. Accessed: 2024-06-19.
- 616 Humza Naveed, Asad Ullah Khan, Shi Qiu, Muhammad Saqib, Saeed Anwar, Muhammad Usman,
617 Naveed Akhtar, Nick Barnes, and Ajmal Mian. A comprehensive overview of large language
618 models, 2024. URL <https://arxiv.org/abs/2307.06435>.
- 619 Clement Neo, Shay B. Cohen, and Fazl Barez. Interpreting context look-ups in transformers: Investi-
620 gating attention-mlp interactions, 2024. URL <https://arxiv.org/abs/2402.15055>.
- 621 Richard Ngo, Lawrence Chan, and Sören Mindermann. The Alignment Problem from a Deep Learning
622 Perspective, 2022. URL <http://arxiv.org/abs/2209.00626>. arXiv:2209.00626 [cs].
- 623 Chris Olah and Josh Batson. Feature manifold toy model. Transformer Circuits May Up-
624 date, May 2023. URL [https://transformer-circuits.pub/2023/may-update/
625 index.html#feature-manifolds](https://transformer-circuits.pub/2023/may-update/index.html#feature-manifolds).
- 626 Bruno A. Olshausen and David J. Field. Sparse coding with an overcomplete basis set: A strategy
627 employed by v1? *Vision Research*, 37(23):3311–3325, 1997. ISSN 0042-6989. doi: [https:
628 //doi.org/10.1016/S0042-6989\(97\)00169-7](https://doi.org/10.1016/S0042-6989(97)00169-7). URL [https://www.sciencedirect.com/
629 science/article/pii/S0042698997001697](https://www.sciencedirect.com/science/article/pii/S0042698997001697).
- 630 Charles O’Neill, Christine Ye, Kartheik Iyer, and John F. Wu. Disentangling dense embeddings with
631 sparse autoencoders, 2024. URL <https://arxiv.org/abs/2408.00657>.
- 632 Maithra Raghu, Justin Gilmer, Jason Yosinski, and Jascha Narain Sohl-Dickstein. Svcca: Singular
633 vector canonical correlation analysis for deep learning dynamics and interpretability. In *Neu-
634 ral Information Processing Systems*, 2017. URL [https://api.semanticscholar.org/
635 CorpusID:23890457](https://api.semanticscholar.org/CorpusID:23890457).
- 636 Senthoran Rajamanoharan, Arthur Conmy, Lewis Smith, Tom Lieberum, Vikrant Varma, János
637 Kramár, Rohin Shah, and Neel Nanda. Improving dictionary learning with gated sparse autoen-
638 coders, 2024a. URL <https://arxiv.org/abs/2404.16014>.
- 639 Senthoran Rajamanoharan, Tom Lieberum, Nicolas Sonnerat, Arthur Conmy, Vikrant Varma, János
640 Kramár, and Neel Nanda. Jumping ahead: Improving reconstruction fidelity with jumprelu sparse
641 autoencoders, 2024b. URL <https://arxiv.org/abs/2407.14435>.
- 642 Nina Rimsky, Nick Gabrieli, Julian Schulz, Meg Tong, Evan Hubinger, and Alexander Matt Turner.
643 Steering llama 2 via contrastive activation addition, 2024.

- 648 Rohin Shah, Vikrant Varma, Ramana Kumar, Mary Phuong, Victoria Krakovna, Jonathan Uesato,
649 and Zac Kenton. Goal Misgeneralization: Why Correct Specifications Aren't Enough For Correct
650 Goals, November 2022. URL <http://arxiv.org/abs/2210.01790>. arXiv:2210.01790
651 [cs].
- 652 Lee Sharkey, Lucius Bushnaq, Dan Braun, Stefan Hex, and Nicholas
653 Goldowsky-Dill. A list of 45 mech interp project ideas from apollo re-
654 search. [https://www.lesswrong.com/posts/KfkpgXdgRheSRWdy8/
655 a-list-of-45-mech-interp-project-ideas-from-apollo-research,](https://www.lesswrong.com/posts/KfkpgXdgRheSRWdy8/a-list-of-45-mech-interp-project-ideas-from-apollo-research)
656 2024. Accessed: 2024-07-25.
- 657
658 Mrinank Sharma, Meg Tong, Tomasz Korbak, David Duvenaud, Amanda Askill, Samuel R. Bowman,
659 Newton Cheng, Esin Durmus, Zac Hatfield-Dodds, Scott R. Johnston, Shauna Kravec, Timothy
660 Maxwell, Sam McCandlish, Kamal Ndousse, Oliver Rausch, Nicholas Schiefer, Da Yan, Miranda
661 Zhang, and Ethan Perez. Towards Understanding Sycophancy in Language Models, October 2023.
662 URL <http://arxiv.org/abs/2310.13548>. arXiv:2310.13548 [cs, stat].
- 663 Gemma Team, Thomas Mesnard, Cassidy Hardin, Robert Dadashi, Surya Bhupatiraju, Shreya Pathak,
664 Laurent Sifre, Morgane Rivière, Mihir Sanjay Kale, Juliette Love, Pouya Tafti, Léonard Hussenot,
665 Pier Giuseppe Sessa, Aakanksha Chowdhery, Adam Roberts, Aditya Barua, Alex Botev, Alex
666 Castro-Ros, Ambrose Slone, Amélie Héliou, Andrea Tacchetti, Anna Bulanova, Antonia Paterson,
667 Beth Tsai, Bobak Shahriari, Charline Le Lan, Christopher A. Choquette-Choo, Clément Crepy,
668 Daniel Cer, Daphne Ippolito, David Reid, Elena Buchatskaya, Eric Ni, Eric Noland, Geng Yan,
669 George Tucker, George-Christian Muraru, Grigory Rozhdestvenskiy, Henryk Michalewski, Ian
670 Tenney, Ivan Grishchenko, Jacob Austin, James Keeling, Jane Labanowski, Jean-Baptiste Lespiau,
671 Jeff Stanway, Jenny Brennan, Jeremy Chen, Johan Ferret, Justin Chiu, Justin Mao-Jones, Katherine
672 Lee, Kathy Yu, Katie Millican, Lars Lowe Sjoesund, Lisa Lee, Lucas Dixon, Machel Reid, Maciej
673 Miłkuła, Mateo Wirth, Michael Sharman, Nikolai Chinaev, Nithum Thain, Olivier Bachem, Oscar
674 Chang, Oscar Wahltinez, Paige Bailey, Paul Michel, Petko Yotov, Rahma Chaabouni, Ramona
675 Comanescu, Reena Jana, Rohan Anil, Ross McIlroy, Ruibo Liu, Ryan Mullins, Samuel L Smith,
676 Sebastian Borgeaud, Sertan Girgin, Sholto Douglas, Shree Pandya, Siamak Shakeri, Soham De,
677 Ted Klimenko, Tom Hennigan, Vlad Feinberg, Wojciech Stokowiec, Yu hui Chen, Zafarali Ahmed,
678 Zhitao Gong, Tris Warkentin, Ludovic Peran, Minh Giang, Clément Farabet, Oriol Vinyals, Jeff
679 Dean, Koray Kavukcuoglu, Demis Hassabis, Zoubin Ghahramani, Douglas Eck, Joelle Barral,
680 Fernando Pereira, Eli Collins, Armand Joulin, Noah Fiedel, Evan Senter, Alek Andreev, and
681 Kathleen Kenealy. Gemma: Open models based on gemini research and technology, 2024a. URL
682 <https://arxiv.org/abs/2403.08295>.
- 683 Gemma Team, Morgane Riviere, Shreya Pathak, Pier Giuseppe Sessa, Cassidy Hardin, Surya
684 Bhupatiraju, Léonard Hussenot, Thomas Mesnard, Bobak Shahriari, Alexandre Ramé, Johan
685 Ferret, Peter Liu, Pouya Tafti, Abe Friesen, Michelle Casbon, Sabela Ramos, Ravin Kumar,
686 Charline Le Lan, Sammy Jerome, Anton Tsitsulin, Nino Vieillard, Piotr Stanczyk, Sertan Girgin,
687 Nikola Momchev, Matt Hoffman, Shantanu Thakoor, Jean-Bastien Grill, Behnam Neyshabur,
688 Olivier Bachem, Alanna Walton, Aliaksei Severyn, Alicia Parrish, Aliya Ahmad, Allen Hutchison,
689 Alvin Abdagic, Amanda Carl, Amy Shen, Andy Brock, Andy Coenen, Anthony Laforge, Antonia
690 Paterson, Ben Bastian, Bilal Piot, Bo Wu, Brandon Royal, Charlie Chen, Chintu Kumar, Chris
691 Perry, Chris Welty, Christopher A. Choquette-Choo, Danila Sinopalnikov, David Weinberger,
692 Dimple Vijaykumar, Dominika Rogozińska, Dustin Herbison, Elisa Bandy, Emma Wang, Eric
693 Noland, Erica Moreira, Evan Senter, Evgenii Eltyshev, Francesco Visin, Gabriel Rasskin, Gary
694 Wei, Glenn Cameron, Gus Martins, Hadi Hashemi, Hanna Klimczak-Plucińska, Harleen Batra,
695 Harsh Dhand, Ivan Nardini, Jacinda Mein, Jack Zhou, James Svensson, Jeff Stanway, Jetha
696 Chan, Jin Peng Zhou, Joana Carrasqueira, Joana Iljazi, Jocelyn Becker, Joe Fernandez, Joost
697 van Amersfoort, Josh Gordon, Josh Lipschultz, Josh Newlan, Ju yeong Ji, Kareem Mohamed,
698 Kartikeya Badola, Kat Black, Katie Millican, Keelin McDonell, Kelvin Nguyen, Kiranbir Sodhia,
699 Kish Greene, Lars Lowe Sjoesund, Lauren Usui, Laurent Sifre, Lena Heuermann, Leticia Lago,
700 Lilly McNealus, Livio Baldini Soares, Logan Kilpatrick, Lucas Dixon, Luciano Martins, Machel
701 Reid, Manvinder Singh, Mark Iverson, Martin Görner, Mat Velloso, Mateo Wirth, Matt Davidow,
Matt Miller, Matthew Rahtz, Matthew Watson, Meg Risdal, Mehran Kazemi, Michael Moynihan,
Ming Zhang, Minsuk Kahng, Minwoo Park, Mofi Rahman, Mohit Khatwani, Natalie Dao, Nenshad
Bardoliwalla, Nesh Devanathan, Neta Dumai, Nilay Chauhan, Oscar Wahltinez, Pankil Botarda,

702 Parker Barnes, Paul Barham, Paul Michel, Pengchong Jin, Petko Georgiev, Phil Culliton, Pradeep
 703 Kuppala, Ramona Comanescu, Ramona Merhej, Reena Jana, Reza Ardeshtir Rokni, Rishabh
 704 Agarwal, Ryan Mullins, Samaneh Saadat, Sara Mc Carthy, Sarah Perrin, Sébastien M. R. Arnold,
 705 Sebastian Krause, Shengyang Dai, Shruti Garg, Shruti Sheth, Sue Ronstrom, Susan Chan, Timothy
 706 Jordan, Ting Yu, Tom Eccles, Tom Hennigan, Tomas Kocisky, Tulsee Doshi, Vihan Jain, Vikas
 707 Yadav, Vilobh Meshram, Vishal Dharmadhikari, Warren Barkley, Wei Wei, Wenming Ye, Woohyun
 708 Han, Woosuk Kwon, Xiang Xu, Zhe Shen, Zhitao Gong, Zichuan Wei, Victor Cotruta, Phoebe
 709 Kirk, Anand Rao, Minh Giang, Ludovic Peran, Tris Warkentin, Eli Collins, Joelle Barral, Zoubin
 710 Ghahramani, Raia Hadsell, D. Sculley, Jeanine Banks, Anca Dragan, Slav Petrov, Oriol Vinyals,
 711 Jeff Dean, Demis Hassabis, Koray Kavukcuoglu, Clement Farabet, Elena Buchatskaya, Sebastian
 712 Borgeaud, Noah Fiedel, Armand Joulin, Kathleen Kenealy, Robert Dadashi, and Alek Andreev.
 713 Gemma 2: Improving open language models at a practical size, 2024b. URL <https://arxiv.org/abs/2408.00118>.

714
 715 Adly Templeton, Tom Conerly, Jonathan Marcus, Jack Lindsey, Trenton Bricken, Brian Chen,
 716 Adam Pearce, Craig Citro, Emmanuel Ameisen, Andy Jones, Hoagy Cunningham, Nicholas L
 717 Turner, Callum McDougall, Monte MacDiarmid, C. Daniel Freeman, Theodore R. Sumers,
 718 Edward Rees, Joshua Batson, Adam Jermyn, Shan Carter, Chris Olah, and Tom Henighan.
 719 Scaling monosemanticity: Extracting interpretable features from claude 3 sonnet. *Trans-*
 720 *former Circuits Thread*, 2024. URL <https://transformer-circuits.pub/2024/scaling-monosemanticity/index.html>.

721
 722 Alexander Matt Turner, Lisa Thiergart, David Udell, Gavin Leech, Ulisse Mini, and Monte MacDi-
 723 armid. Activation addition: Steering language models without optimization, 2023.

724
 725 Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz
 726 Kaiser, and Illia Polosukhin. Attention is all you need. *Advances in neural information processing*
 727 *systems*, 30, 2017.

728
 729 Liwei Wang, Lunjia Hu, Jiayuan Gu, Yue Wu, Zhiqiang Hu, Kun He, and John Hopcroft. Towards
 730 understanding learning representations: To what extent do different neural networks learn the
 731 same representation. In *Proceedings of the 32nd International Conference on Neural Information*
 732 *Processing Systems*, NIPS’18, pages 9607–9616, Red Hook, NY, USA, 2018. Curran Associates
 733 Inc. ISBN 9781510860964.

734
 735 Jason Yosinski, Jeff Clune, Yoshua Bengio, and Hod Lipson. How transferable are features in deep
 736 neural networks? In *Proceedings of the 27th International Conference on Neural Information*
 737 *Processing Systems - Volume 2*, NIPS’14, page 3320–3328, Cambridge, MA, USA, 2014. MIT
 738 Press.

739
 740 Andy Zou, Long Phan, Sarah Chen, James Campbell, Phillip Guo, Richard Ren, Alexander Pan,
 741 Xuwang Yin, Mantas Mazeika, Ann-Kathrin Dombrowski, Shashwat Goel, Nathaniel Li, Michael J.
 742 Byun, Zifan Wang, Alex Mallen, Steven Basart, Sanmi Koyejo, Dawn Song, Matt Fredrikson,
 743 J. Zico Kolter, and Dan Hendrycks. Representation engineering: A top-down approach to ai
 744 transparency, 2023.

745 A NOISE OF MANY-TO-1 MAPPINGS

746
 747 Overall, we do not aim for mappings that uniquely pair every feature of both spaces as we hypothesize
 748 that it is unlikely that all the features are the same in each SAE; rather, we are looking for large
 749 subspaces where the features match by activations and we can map one subspace to another.

750
 751 We hypothesize that these many-to-1 features may contribute a lot of noise to feature space alignment
 752 similarity, as information from a larger space (eg. a ball) is collapsed into a dimension with much
 753 less information (eg. a point). Thus, when we only include 1-1 features in our similarity scores, we
 754 receive scores with much less noise.

755
 756 As shown in the comparisons of Figure 4 to Figure 5, we find the 1-1 feature mappings give slightly
 757 higher scores for many layer pairs, like for the SVCCA scores at L2 vs L3 for Pythia-70m vs Pythia-
 758 160m, though they give slightly lower scores for a few layer pairs, like for the SVCCA scores at L5

756
757
758
759
760
761
762
763
764
765
766
767
768
769
770
771
772
773
774
775
776
777
778
779
780
781
782
783
784
785
786
787
788
789
790
791
792
793
794
795
796
797
798
799
800
801
802
803
804
805
806
807
808
809

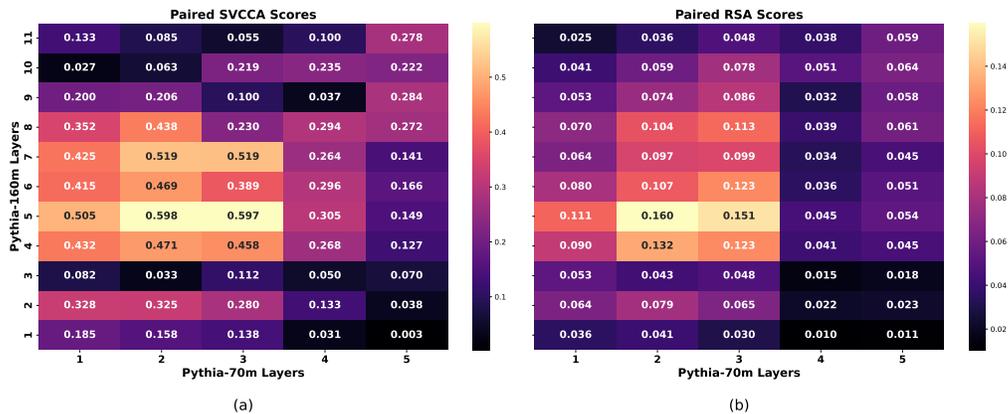


Figure 4: (a) SVCCA and (b) RSA **Many-to-1** paired scores of SAEs for layers in Pythia-70m vs layers in Pythia-160m. We note there appears to be an "anomaly" at L2 vs L3 with a low SVCCA score; we find that taking 1-1 features greatly increases this score from 0.03 to 0.35, as shown and explained in Figure 2.

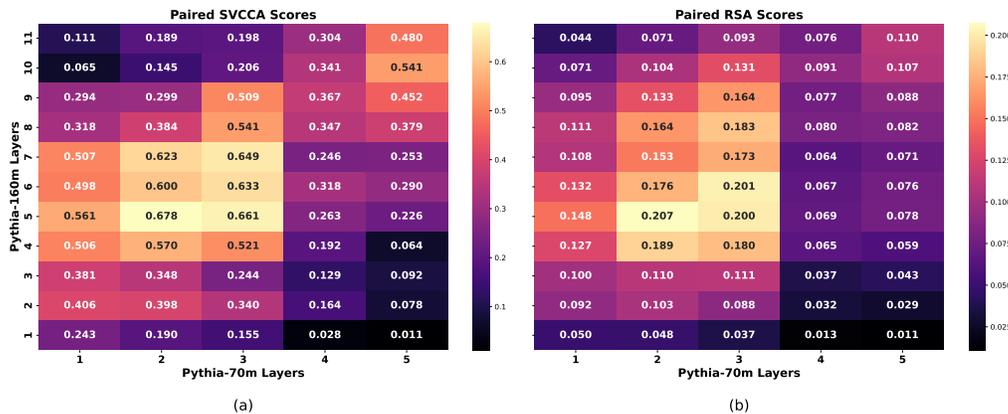


Figure 5: (a) SVCCA and (b) RSA **1-1** paired scores of SAEs for layers in Pythia-70m vs layers in Pythia-160m. Compared to Figure 4, some of the scores are slightly higher, and the SVCCA score at L2 vs L3 for 70m vs 160m is much higher. On the other hand, some scores are slightly lower, such as SVCCA at L5 vs L4 for 70m vs 160m. This is the same figure as Figure 2; it is shown here again for easier comparison to the Many-to-1 scores in Figure 4.

vs L4 for Pythia-70m vs Pythia-160m. Given that most layer pair scores are slightly higher, we use 1-1 feature mappings in our main text results.

Additionally, we notice that for semantic subspace experiments, 1-1 gave vastly better scores than many-to-1. We hypothesize that this is because since these feature subspaces are smaller, the noise from the many-to-1 mappings have a greater impact on the score.

We also looked into using more types of 1-1 feature matching, such as pairing many-to-1 features with their "next highest" or using efficient methods to first select the mapping pair with the highest correlation, taking those features off as candidates for future mappings, and continuing this for the next highest correlations. This also appeared to work well, though further investigation is needed. More analysis can also be done for mapping one feature from SAE A to many features in SAE B.

810
811
812
813
814
815
816
817
818
819
820
821
822
823
824
825
826
827
828
829
830
831
832
833
834
835
836
837
838
839
840
841
842
843
844
845
846
847
848
849
850
851
852
853
854
855
856
857
858
859
860
861
862
863

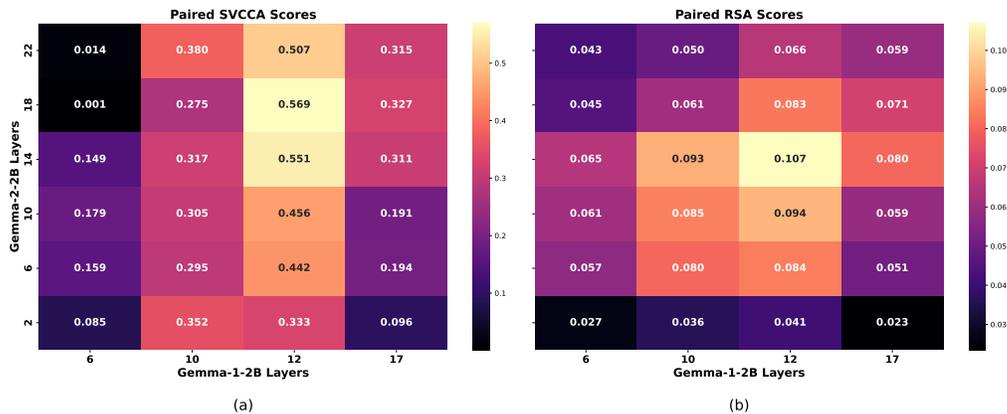


Figure 6: (a) SVCCA and (b) RSA **Many-to-1** paired scores of SAEs for layers in Gemma-1-2B vs layers in Gemma-2-2B. We obtain similar scores compared to the 1-1 paired scores in Figure 3.

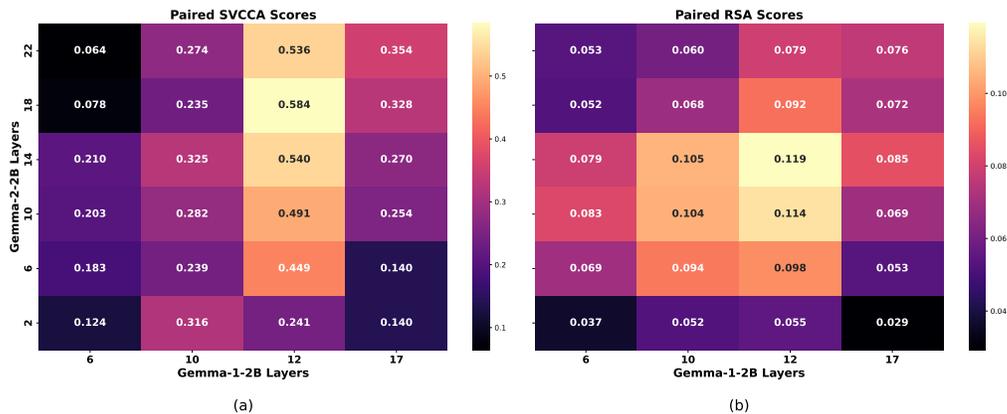


Figure 7: (a) SVCCA and (b) RSA **1-1** paired scores of SAEs for layers in Gemma-1-2B vs layers in Gemma-2-2B. Compared to Figure 6, some of the scores are slightly higher, while some are slightly lower. This is the same figure as Figure 3; it is shown here again for easier comparison to the Many-to-1 scores in Figure 4.

B NOISE OF NON-CONCEPT FEATURES

We define "non-concept features" as features which are not modeling specific concepts that can be mapped well across models. Their highest activations are on new lines, spaces, punctuation, and EOS tokens. As such, they may introduce noise when computing similarity scores, as their removal appears to greatly improve similarity scores. We hypothesize that one reason they introduce noise is that these could be tokens that aren't specific to concepts, but multiple concepts; thus, they are mixing concept spaces up when matching is done.

For instance, say feature X from model A fires highly on "!" in the sample "I am BOB!", while feature Y from model B fires highly on "!" in the sample, "that cat ate!"

A feature X that activates on "!" in "BOB!" may be firing for names, while a feature Y that activates on "!" in "ate!" may fire on food-related tokens. As such, features activating on non-concept features are not concept specific, and we could map a feature firing on names to a feature firing on food. We call these possibly "incoherent" mappings "non-concept mappings". By removing non-concept mappings, we show that SAEs reveal highly similar feature spaces.

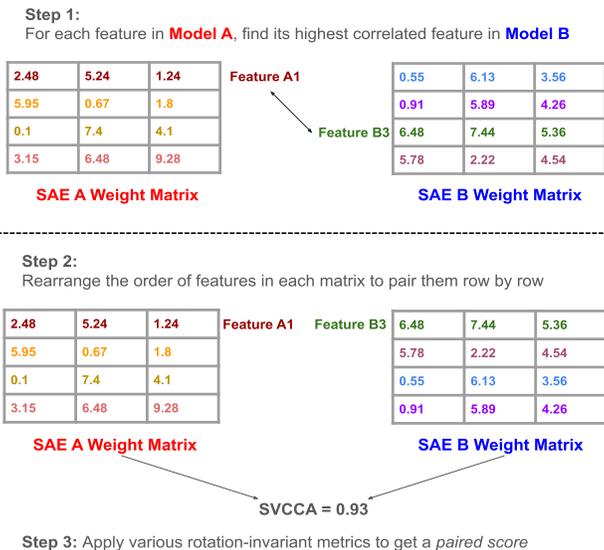
We filter "non-concept features" that fire, for at least one of their top 5 dataset examples, on a "non-concept" keyword. The list of "non-concept" keywords we use is given below:

- 864 • \n
- 865 • \n
- 866 • (empty string)
- 867 • (space)
- 868 • .
- 869 • ,
- 870 • !
- 871 • ?
- 872 • -
- 873 • <bos>
- 874 • <|endoftext|>

879 There may exist other non-concept features that we did not filter, such as colons, brackets and carats.
 880 We do not include arithmetic symbols, as those are domain specific to specific concepts. Additionally,
 881 we have yet to perform a detailed analysis on which non-concept features influence the similarity
 882 scores more than others; this may be done in future work.

884 C ILLUSTRATED STEPS OF COMPARISON METHOD

886 Figure 8 demonstrates steps 1 to 3 to carry out our similarity comparison experiments.



908 Figure 8: Steps to Main Results. We first find correlated pairs to solve neuron permutation, and then
 909 apply similarity metrics to solve latent space rotation issues.

912 D CONCEPT CATEGORIES LIST AND FURTHER ANALYSIS

914 The keywords used for each concept group for the semantic subspace experiments are given in Table
 915 4. We generate these keywords using GPT-4 Bubeck et al. (2023) using the prompt "give a list of N
 916 single token keywords as a python list that are related to the concept of C", such that N is a number
 917 (we use N=50 or 100) and C is a concept like Emotions. We then manually filter these lists to avoid
 using keywords which have meanings outside of the concept group (eg. "even" does not largely mean

"divisible by two" because it is often used to mean "even if..."). We also aim to select keywords which are single-tokens with no spaces, and are case-insensitive.³

When searching for keyword matches from each feature’s list of top 5 highest activating tokens, we use keyword matches that avoid features with dataset samples which use the keywords in compound words. For instance, when searching for features that activate on the token "king" from the People/Roles concept, we avoid using features that activate on "king" when "king" is a part of the word "seeking", as that is not related to the People/Roles concept.

We perform interpretability experiments to check which keywords of the features kept are activated on. We find that for many feature pairs, they activate on the same keyword and are monosemantic. Most keywords in each concept group are not kept. Additionally, for the layer pairs we checked on, after filtering there were only a few keywords that multiplied features fired on. This shows that the high similarity is not because the same keyword is over-represented in a space.

For instance, for Layer 3 in Pythia-70m vs Layer 5 in Pythia-160m, and for the concept category of Emotions, we find which keywords from the category appear in the top 5 activating tokens of each of the Pythia-70m features from the 24 feature mappings described in Table 1. We only count a feature if it appears once in the top 5 of a feature’s top activating tokens; if three out of five of a feature’s top activating tokens are "smile", then smile is only counted once. The results for Models A and B are given in Table 3. Not only are their counts similar, indicating similar features, but there is not a great over-representation of features. There are a total of 24 features, and a total of 25 keywords (as some features can activate on multiple keywords in their top 5). We note that even if a feature fires for a keyword does not mean that feature’s purpose is only to "recognize that keyword", as dataset examples as just one indicator of what a feature’s purpose is (which still can have multiple, hard-to-interpret purposes).

| Model A | | Model B | |
|---------|-------|---------|-------|
| Keyword | Count | Keyword | Count |
| free | 4 | free | 5 |
| pain | 4 | pain | 4 |
| smile | 3 | love | 3 |
| love | 2 | hate | 2 |
| hate | 2 | calm | 2 |
| calm | 2 | smile | 2 |
| sad | 2 | kind | 1 |
| kind | 1 | shy | 1 |
| shy | 1 | doubt | 1 |
| doubt | 1 | trust | 1 |
| trust | 1 | peace | 1 |
| peace | 1 | joy | 1 |
| joy | 1 | sad | 1 |

Table 3: Count of Model A vs Model B features with keywords from the Emotion category in the semantic subspace found in Table 1.

However, as there are many feature pairs, we did not perform a thorough analysis of this yet, and thus did not include this analysis in this paper. We also check this in LLMs, and find that the majority of LLM neurons are polysemantic. We do not just check the top 5 dataset samples for LLM neurons, but the top 90, as the SAEs have an expansion factor that is around 16x or 32x bigger than the LLM dimensions they take in as input.

Calendar is a subset of Time, removing keywords like “after” and “soon”, and keeping only “day” and “month” type of keywords pertaining to dates.

³However, not all the tokens in the Table 4 are single token.

972
973
974
975
976
977
978
979
980
981
982
983
984
985
986
987
988
989
990
991
992
993
994
995
996
997
998
999
1000
1001
1002
1003
1004
1005
1006
1007
1008
1009
1010
1011
1012
1013
1014
1015
1016
1017
1018
1019
1020
1021
1022
1023
1024
1025

Table 4: Keywords Organized by Category

| Category | Keywords |
|--------------|---|
| Time | day, night, week, month, year, hour, minute, second, now, soon, later, early, late, morning, evening, noon, midnight, dawn, dusk, past, present, future, before, after, yesterday, today, tomorrow, next, previous, instant, era, age, decade, century, millennium, moment, pause, wait, begin, start, end, finish, stop, continue, forever, constant, frequent, occasion, season, spring, summer, autumn, fall, winter, anniversary, deadline, schedule, calendar, clock, duration, interval, epoch, generation, period, cycle, timespan, shift, quarter, term, phase, lifetime, timeline, delay, prompt, timely, recurrent, daily, weekly, monthly, yearly, annual, biweekly, timeframe |
| Calendar | day, night, week, month, year, hour, minute, second, morning, evening, noon, midnight, dawn, dusk, yesterday, today, tomorrow, decade, century, millennium, season, spring, summer, autumn, fall, winter, calendar, clock, daily, weekly, monthly, yearly, annual, biweekly, timeframe |
| People/Roles | man, girl, boy, kid, dad, mom, son, sis, bro, chief, priest, king, queen, duke, lord, friend, clerk, coach, nurse, doc, maid, clown, guest, peer, punk, nerd, jock |
| Nature | tree, grass, stone, rock, cliff, hill, dirt, sand, mud, wind, storm, rain, cloud, sun, moon, leaf, branch, twig, root, bark, seed, tide, lake, pond, creek, sea, wood, field, shore, snow, ice, flame, fire, fog, dew, hail, sky, earth, glade, cave, peak, ridge, dust, air, mist, heat |
| Emotions | joy, glee, pride, grief, fear, hope, love, hate, pain, shame, bliss, rage, calm, shock, dread, guilt, peace, trust, scorn, doubt, hurt, wrath, laugh, cry, smile, frown, gasp, blush, sigh, grin, woe, spite, envy, glow, thrill, mirth, bored, cheer, charm, grace, shy, brave, proud, glad, mad, sad, tense, free, kind |
| MonthNames | January, February, March, April, May, June, July, August, September, October, November, December |
| Countries | USA, Canada, Brazil, Mexico, Germany, France, Italy, Spain, UK, Australia, China, Japan, India, Russia, Korea, Argentina, Egypt, Iran, Turkey |
| Biology | gene, DNA, RNA, virus, bacteria, fungus, brain, lung, blood, bone, skin, muscle, nerve, vein, organ, evolve, enzyme, protein, lipid, membrane, antibody, antigen, ligand, substrate, receptor, cell, chromosome, nucleus, cytoplasm |

Table 5: Number of Features in each Layer Pair Mapping after filtering Non-Concept Features for Pythia-70m (cols) vs Pythia-160m (rows) out of a total of 32768 SAE features in both models.

| Layer | 1 | 2 | 3 | 4 | 5 |
|-------|-------|-------|-------|-------|-------|
| 1 | 23600 | 21423 | 26578 | 19696 | 19549 |
| 2 | 16079 | 14201 | 17578 | 12831 | 12738 |
| 3 | 7482 | 7788 | 7841 | 7017 | 6625 |
| 4 | 12756 | 11195 | 13349 | 9019 | 9357 |
| 5 | 7987 | 6825 | 8170 | 5367 | 5624 |
| 6 | 10971 | 9578 | 10937 | 8099 | 8273 |
| 7 | 15074 | 12988 | 16326 | 12720 | 12841 |
| 8 | 14445 | 12580 | 15300 | 11779 | 11942 |
| 9 | 13320 | 11950 | 14338 | 11380 | 11436 |
| 10 | 9834 | 8573 | 9742 | 7858 | 8084 |
| 11 | 6936 | 6551 | 7128 | 5531 | 6037 |

Table 6: Number of **1-1** Features in each Layer Pair Mapping after filtering Non-Concept Features for Pythia-70m (cols) vs Pythia-160m (rows) out of a total of 32768 SAE features in both models.

| Layer | 1 | 2 | 3 | 4 | 5 |
|-------|------|------|------|------|------|
| 1 | 7553 | 5049 | 5853 | 3244 | 3115 |
| 2 | 6987 | 4935 | 5663 | 3217 | 3066 |
| 3 | 4025 | 3152 | 3225 | 2178 | 1880 |
| 4 | 6649 | 5058 | 6015 | 3202 | 3182 |
| 5 | 5051 | 4057 | 4796 | 2539 | 2596 |
| 6 | 5726 | 4528 | 5230 | 3226 | 3031 |
| 7 | 7017 | 5659 | 6762 | 4032 | 3846 |
| 8 | 6667 | 5179 | 6321 | 3887 | 3808 |
| 9 | 6205 | 4820 | 5773 | 3675 | 3682 |
| 10 | 4993 | 3859 | 4602 | 3058 | 3356 |
| 11 | 3609 | 2837 | 3377 | 2339 | 2691 |

E ADDITIONAL RESULTS

As shown in Figure 9 for Pythia and Figure 10 for Gemma, we also find that mean activation correlation does not always correlate with the global similarity metric scores. For instance, a subset of feature pairings with a moderately high mean activation correlation (eg. 0.6) may yield a low SVCCA score (eg. 0.03).

The number of features after filtering non-concept features and many-to-1 features are given in Tables 5 and 6. The number of features after filtering non-concept features and many-to-1 features are given in Tables 7 and 8.

Table 7: Number of Features in each Layer Pair Mapping after filtering Non-Concept Features for Gemma-1-2B (cols) vs Gemma-2-2B (rows) out of a total of 16384 SAE features in both models.

| Layer | 6 | 10 | 12 | 17 |
|-------|------|------|------|------|
| 2 | 8926 | 8685 | 8647 | 4816 |
| 6 | 4252 | 4183 | 4131 | 2474 |
| 10 | 6458 | 6320 | 6277 | 3658 |
| 14 | 4213 | 4208 | 4214 | 2672 |
| 18 | 3672 | 3679 | 3771 | 2515 |
| 22 | 4130 | 4100 | 4168 | 3338 |

Table 8: Number of **1-1** Features in each Layer Pair Mapping after filtering Non-Concept Features for Gemma-1-2B (cols) vs Gemma-2-2B (rows) out of a total of 16384 SAE features in both models.

| Layer | 6 | 10 | 12 | 17 |
|-------|------|------|------|------|
| 2 | 3427 | 3874 | 3829 | 1448 |
| 6 | 3056 | 3336 | 3261 | 1266 |
| 10 | 4110 | 4650 | 4641 | 1616 |
| 14 | 2967 | 3524 | 3553 | 1414 |
| 18 | 2636 | 3058 | 3239 | 1543 |
| 22 | 2646 | 3037 | 3228 | 1883 |

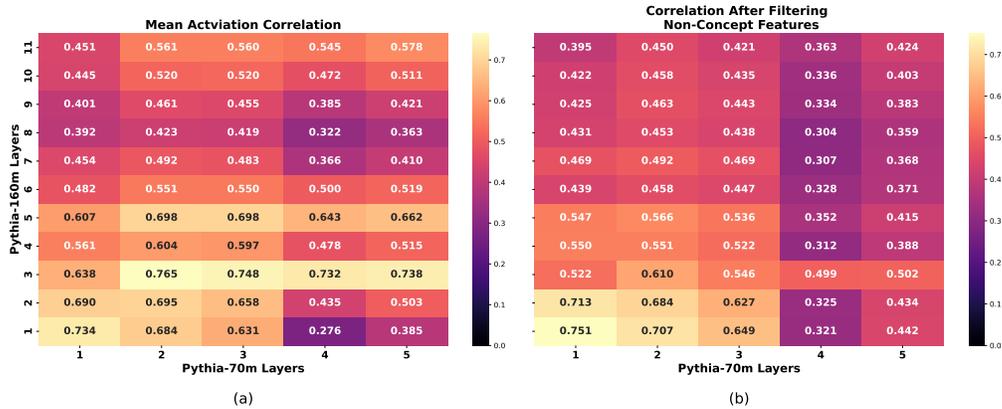


Figure 9: Mean Activation Correlation before (a) and after (b) filtering non-concept features for Pythia-70m vs Pythia-160m. We note these patterns are different from those of the SVCCA and RSA scores in Figure 4, indicating that these three metrics each reveal different patterns not shown by other metrics.

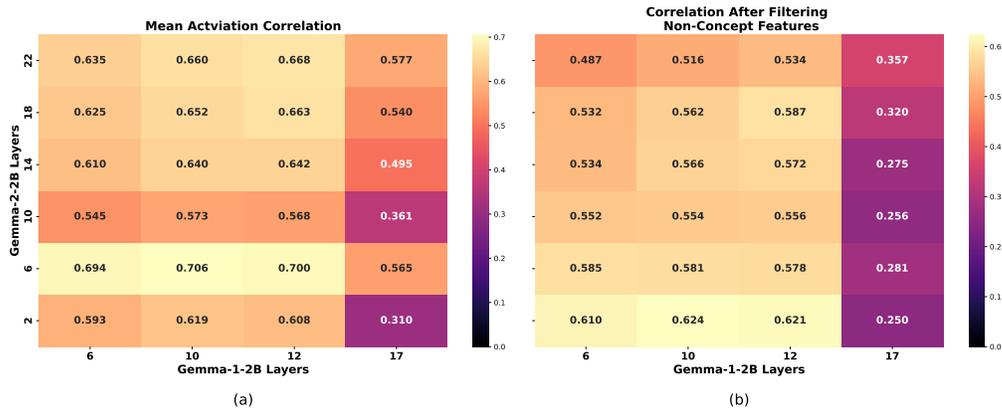


Figure 10: Mean Activation Correlation before (a) and after (b) filtering non-concept features for Gemma-1-2B vs Gemma-2-2B. We note these patterns are different from those of the SVCCA and RSA scores in Figure 6, indicating that these three metrics each reveal different patterns not shown by other metrics.

1134
1135
1136
1137
1138
1139
1140
1141
1142
1143
1144
1145
1146
1147
1148
1149
1150
1151
1152
1153
1154
1155
1156
1157
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

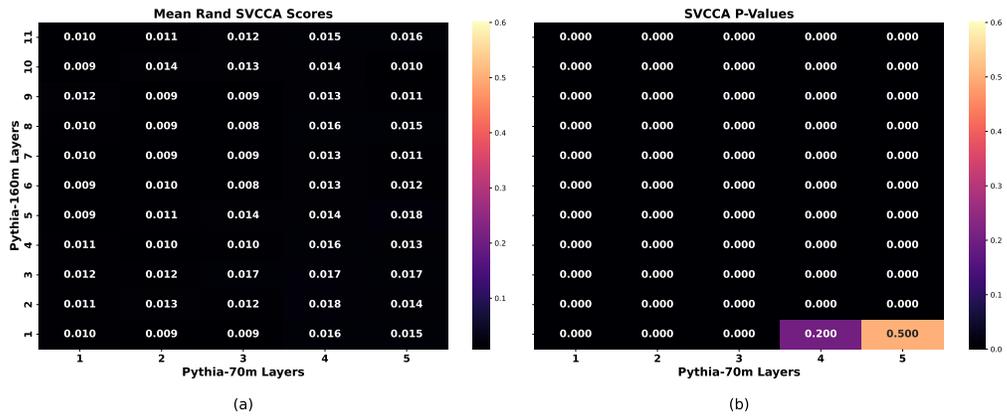


Figure 11: Mean Randomly Paired SVCCA scores and P-values of SAEs for layers in Pythia-70m vs Pythia-160m. Compared to Paired Scores in Figure 4, these are all very low.

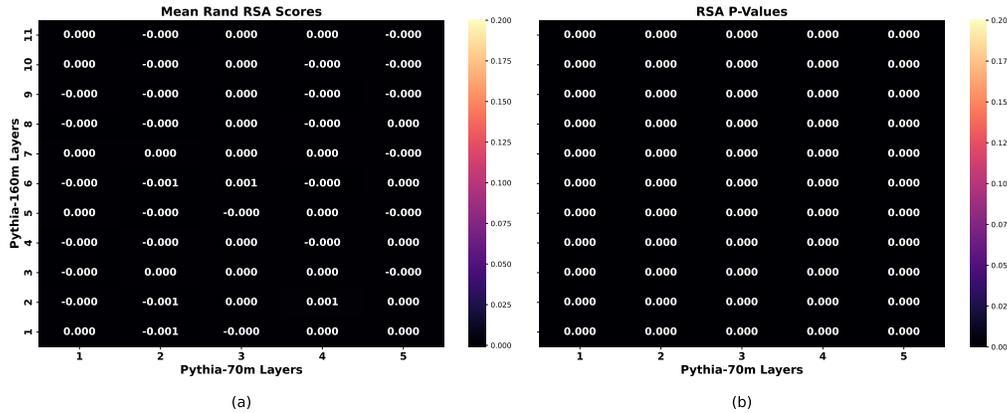


Figure 12: Mean Randomly Paired RSA scores and P-values of SAEs for layers in Pythia-70m vs Pythia-160m. Compared to Paired Scores in Figure 4, these are all very low.

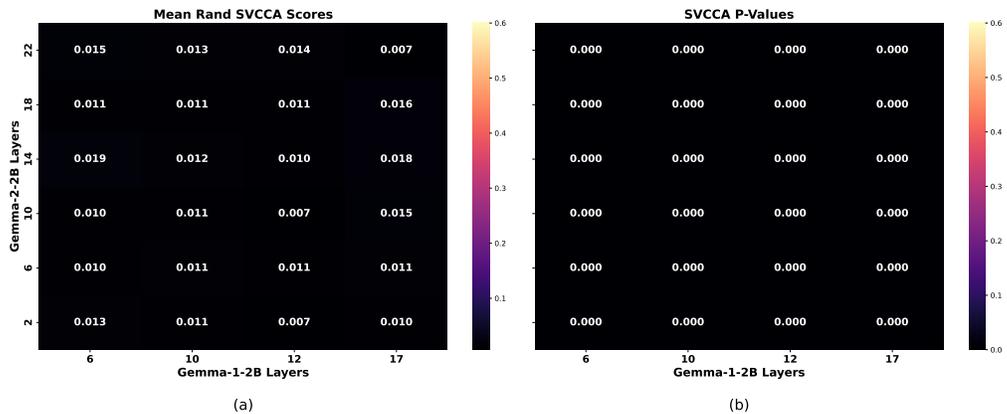


Figure 13: Mean Randomly Paired SVCCA scores and P-values of SAEs for layers in Gemma-1-2B vs layers in Gemma-2-2B. Compared to Paired Scores in Figure 6, these are all very low.

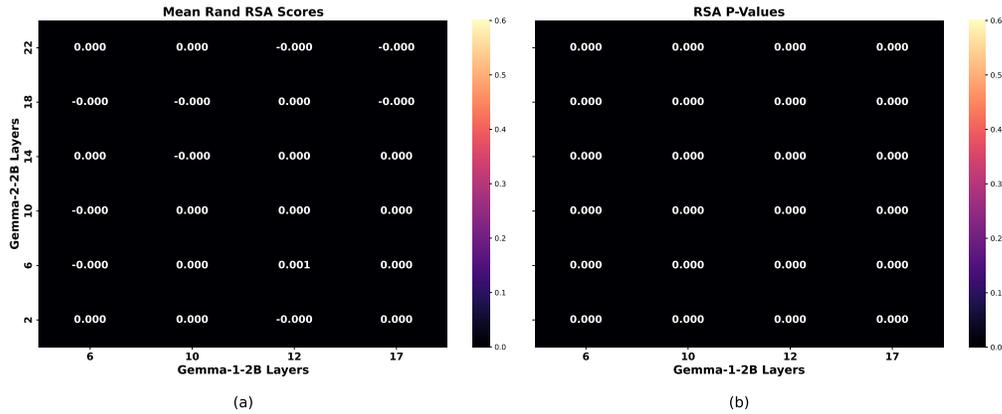


Figure 14: Mean Randomly Paired RSA scores and P-values of SAEs for layers in Gemma-1-2B vs layers in Gemma-2-2B. Compared to Paired Scores in Figure 6, these are all very low.

Table 9: RSA scores and random mean results for 1-1 semantic subspaces of L3 of Pythia-70m vs L5 of Pythia-160m. P-values are taken for 1000 samples in the null distribution.

| Concept Subspace | Number of 1-1 Features | Paired Mean | Random Shuffling Mean | p-value |
|------------------|------------------------|-------------|-----------------------|---------|
| Time | 228 | 0.10 | 0.00 | 0.00 |
| Calendar | 126 | 0.09 | 0.00 | 0.00 |
| Nature | 46 | 0.22 | 0.00 | 0.00 |
| MonthNames | 32 | 0.76 | 0.00 | 0.00 |
| Countries | 32 | 0.10 | 0.00 | 0.03 |
| People/Roles | 31 | 0.18 | 0.00 | 0.01 |
| Emotions | 24 | 0.46 | 0.00 | 0.00 |

1242
 1243
 1244
 1245
 1246
 1247
 1248
 1249
 1250
 1251
 1252
 1253
 1254
 1255
 1256
 1257
 1258
 1259
 1260
 1261
 1262
 1263
 1264
 1265
 1266
 1267
 1268
 1269
 1270
 1271
 1272
 1273
 1274
 1275
 1276
 1277
 1278
 1279
 1280
 1281
 1282
 1283
 1284
 1285
 1286
 1287
 1288
 1289
 1290
 1291
 1292
 1293
 1294
 1295

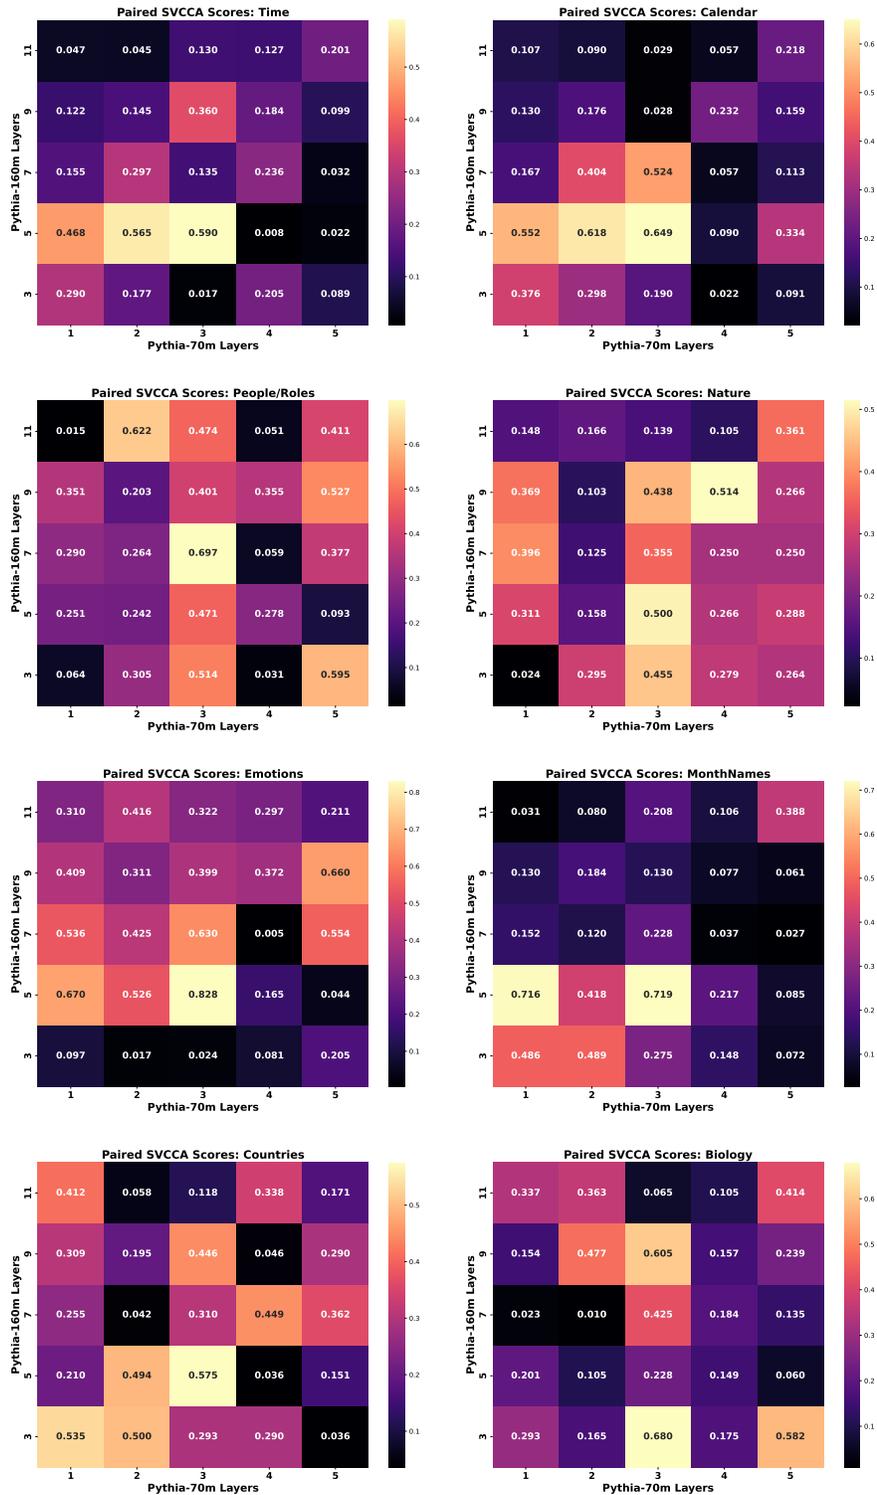


Figure 15: 1-1 Paired SVCCA scores of SAEs for layers in Pythia-70m vs layers in Pythia-160m for Concept Categories. Middle layers appear to be the most similar with one another.

1296
 1297
 1298
 1299
 1300
 1301
 1302
 1303
 1304
 1305
 1306
 1307
 1308
 1309
 1310
 1311
 1312
 1313
 1314
 1315
 1316
 1317
 1318
 1319
 1320
 1321
 1322
 1323
 1324
 1325
 1326
 1327
 1328
 1329
 1330
 1331
 1332
 1333
 1334
 1335
 1336
 1337
 1338
 1339
 1340
 1341
 1342
 1343
 1344
 1345
 1346
 1347
 1348
 1349

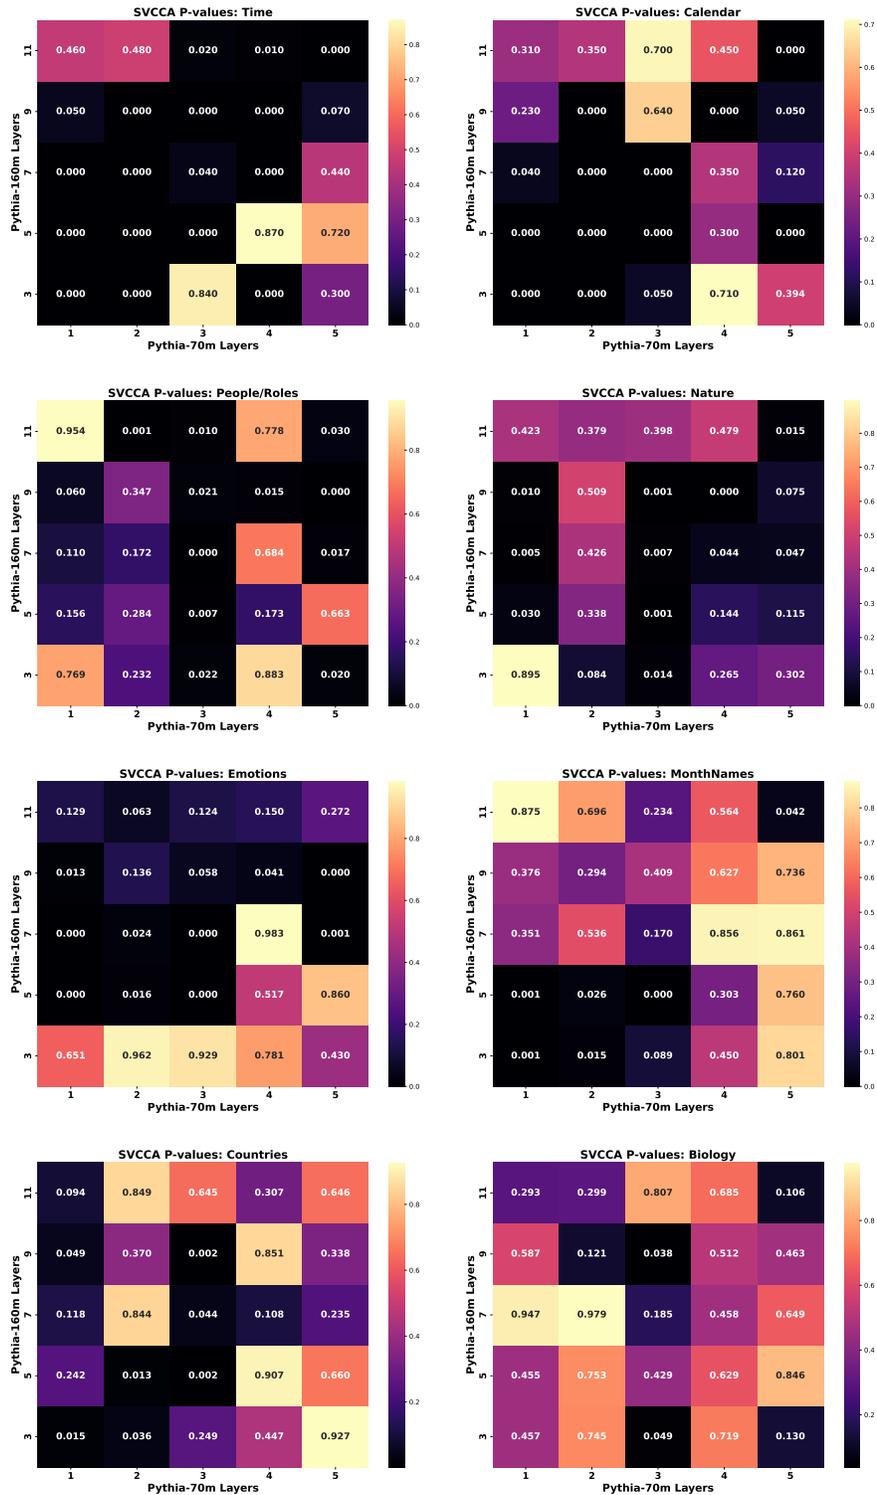


Figure 16: 1-1 SVCCA p-values of SAEs for layers in Pythia-70m vs layers in Pythia-160m for Concept Categories. A lower p-value indicates more statistically meaningful similarity.

1350
 1351
 1352
 1353
 1354
 1355
 1356
 1357
 1358
 1359
 1360
 1361
 1362
 1363
 1364
 1365
 1366
 1367
 1368
 1369
 1370
 1371
 1372
 1373
 1374
 1375
 1376
 1377
 1378
 1379
 1380
 1381
 1382
 1383
 1384
 1385
 1386
 1387
 1388
 1389
 1390
 1391
 1392
 1393
 1394
 1395
 1396
 1397
 1398
 1399
 1400
 1401
 1402
 1403

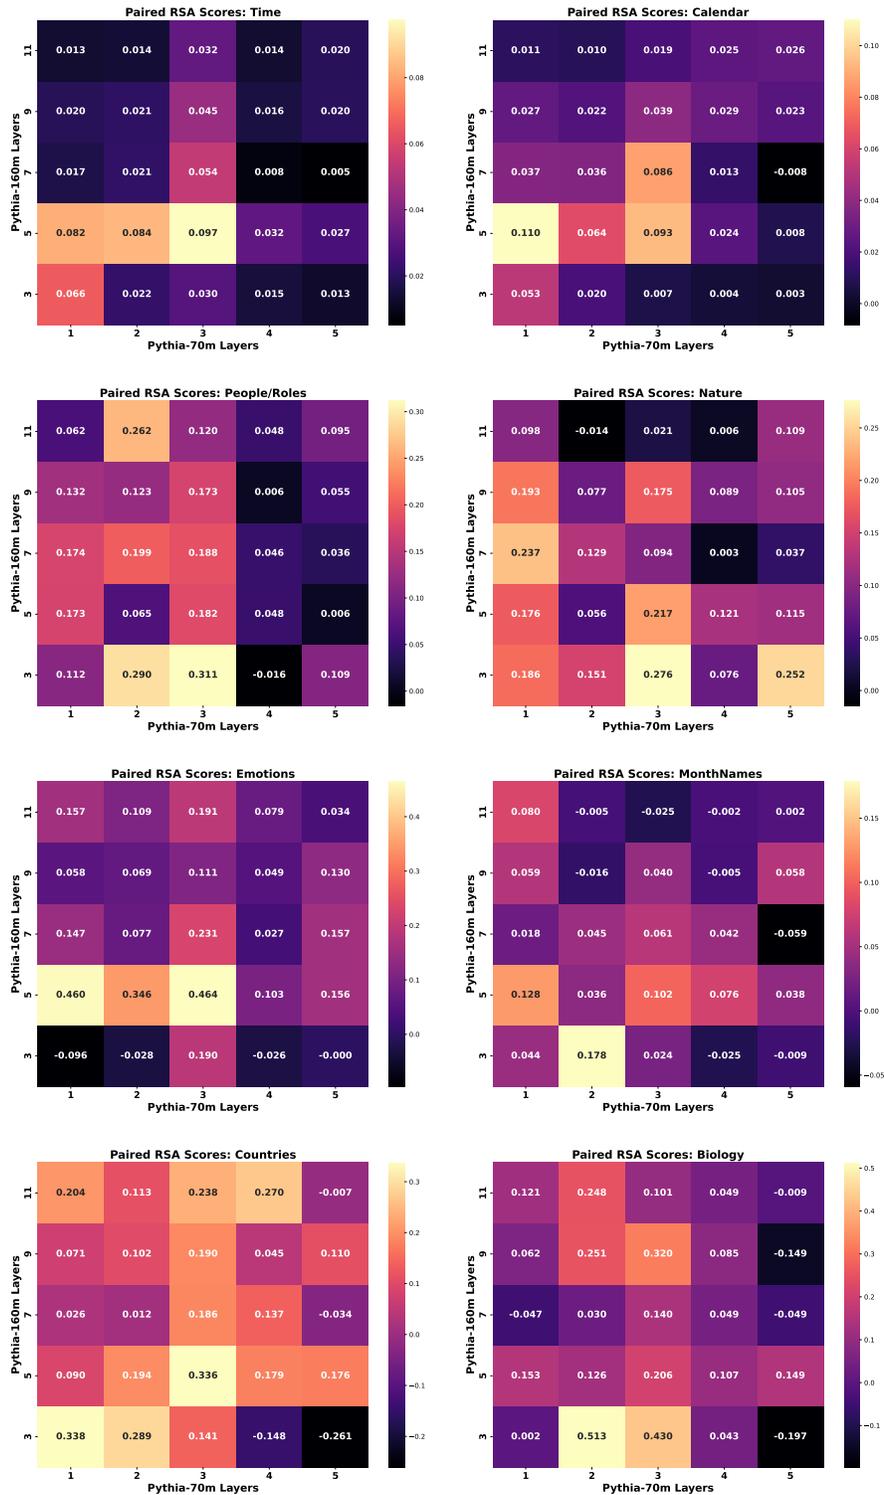


Figure 17: 1-1 Paired RSA scores of SAEs for layers in Pythia-70m vs layers in Pythia-160m for Concept Categories. Middle layers appear to be the most similar with one another.

1404
1405
1406
1407
1408
1409
1410
1411
1412
1413
1414
1415
1416
1417
1418
1419
1420
1421
1422
1423
1424
1425
1426
1427
1428
1429
1430
1431
1432
1433
1434
1435
1436
1437
1438
1439
1440
1441
1442
1443
1444
1445
1446
1447
1448
1449
1450
1451
1452
1453
1454
1455
1456
1457

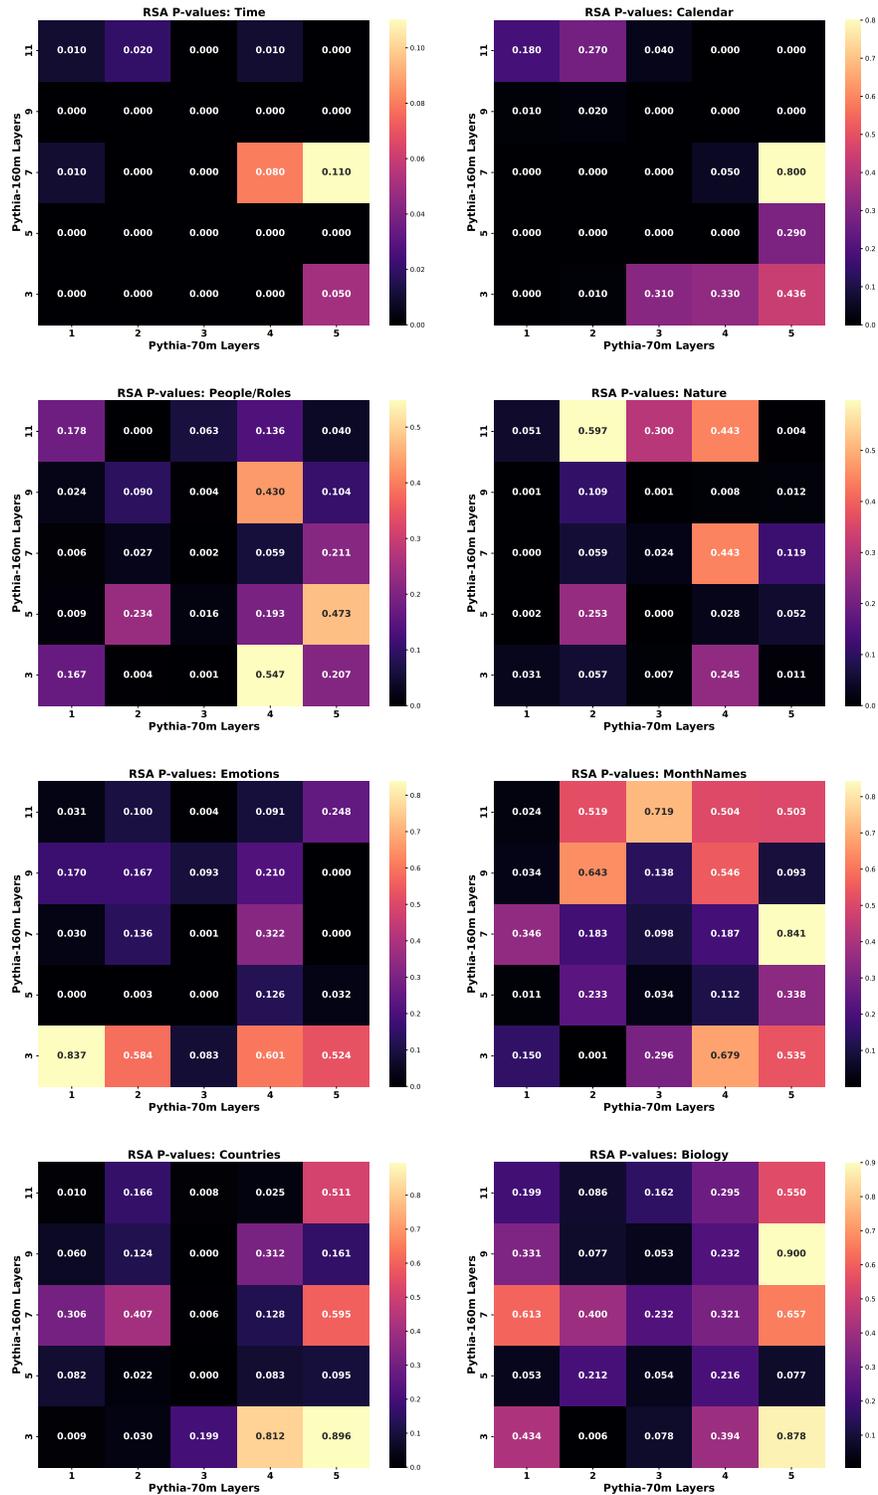


Figure 18: 1-1 RSA p-values of SAEs for layers in Pythia-70m vs layers in Pythia-160m for Concept Categories. A lower p-value indicates more statistically meaningful similarity.

1458
 1459
 1460
 1461
 1462
 1463
 1464
 1465
 1466
 1467
 1468
 1469
 1470
 1471
 1472
 1473
 1474
 1475
 1476
 1477
 1478
 1479
 1480
 1481
 1482
 1483
 1484
 1485
 1486
 1487
 1488
 1489
 1490
 1491
 1492
 1493
 1494
 1495
 1496
 1497
 1498
 1499
 1500
 1501
 1502
 1503
 1504
 1505
 1506
 1507
 1508
 1509
 1510
 1511

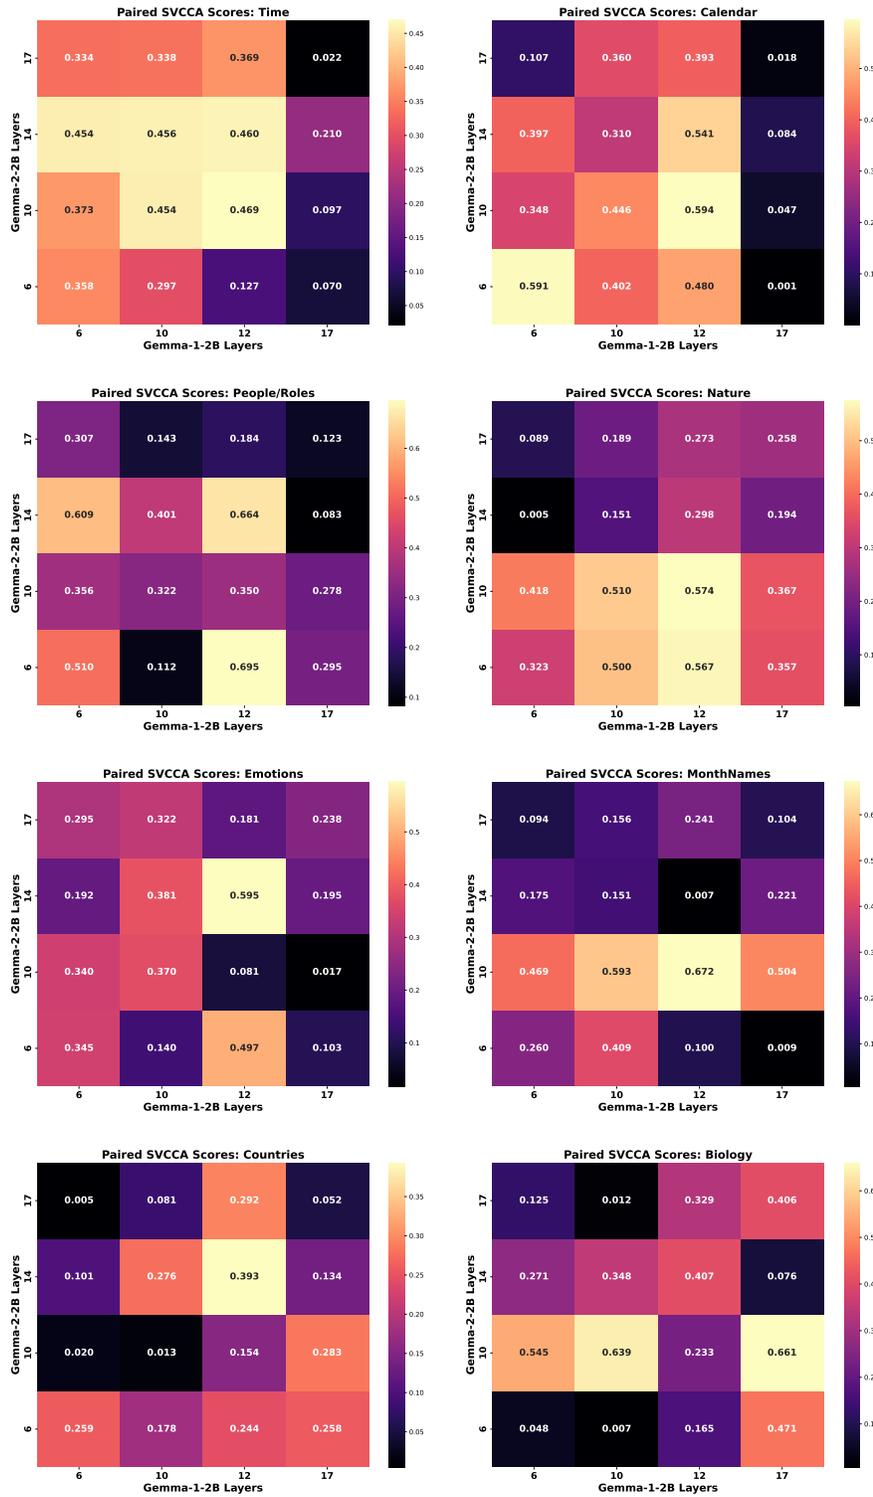


Figure 19: 1-1 Paired SVCCA scores of SAEs for layers in Gemma-1-2b vs layers in Gemma-2-2b for Concept Categories. Middle layers appear to be the most similar with one another.

1512
 1513
 1514
 1515
 1516
 1517
 1518
 1519
 1520
 1521
 1522
 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

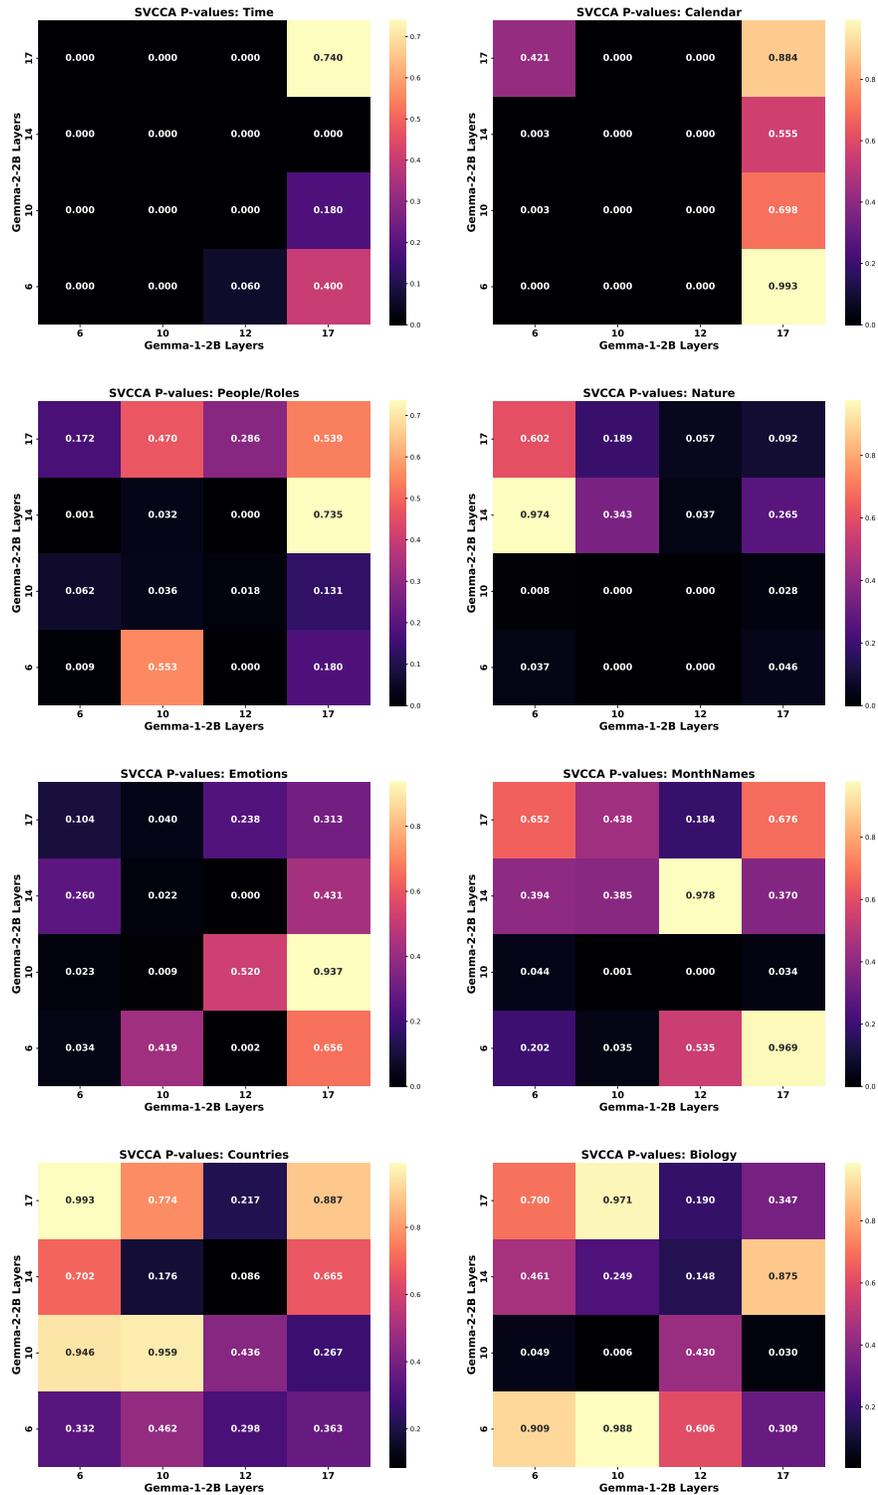


Figure 20: 1-1 SVCCA p-values of SAEs for layers in Gemma-1-2b vs layers in Gemma-2-2b for Concept Categories. A lower p-value indicates more statistically meaningful similarity.

1566
 1567
 1568
 1569
 1570
 1571
 1572
 1573
 1574
 1575
 1576
 1577
 1578
 1579
 1580
 1581
 1582
 1583
 1584
 1585
 1586
 1587
 1588
 1589
 1590
 1591
 1592
 1593
 1594
 1595
 1596
 1597
 1598
 1599
 1600
 1601
 1602
 1603
 1604
 1605
 1606
 1607
 1608
 1609
 1610
 1611
 1612
 1613
 1614
 1615
 1616
 1617
 1618
 1619

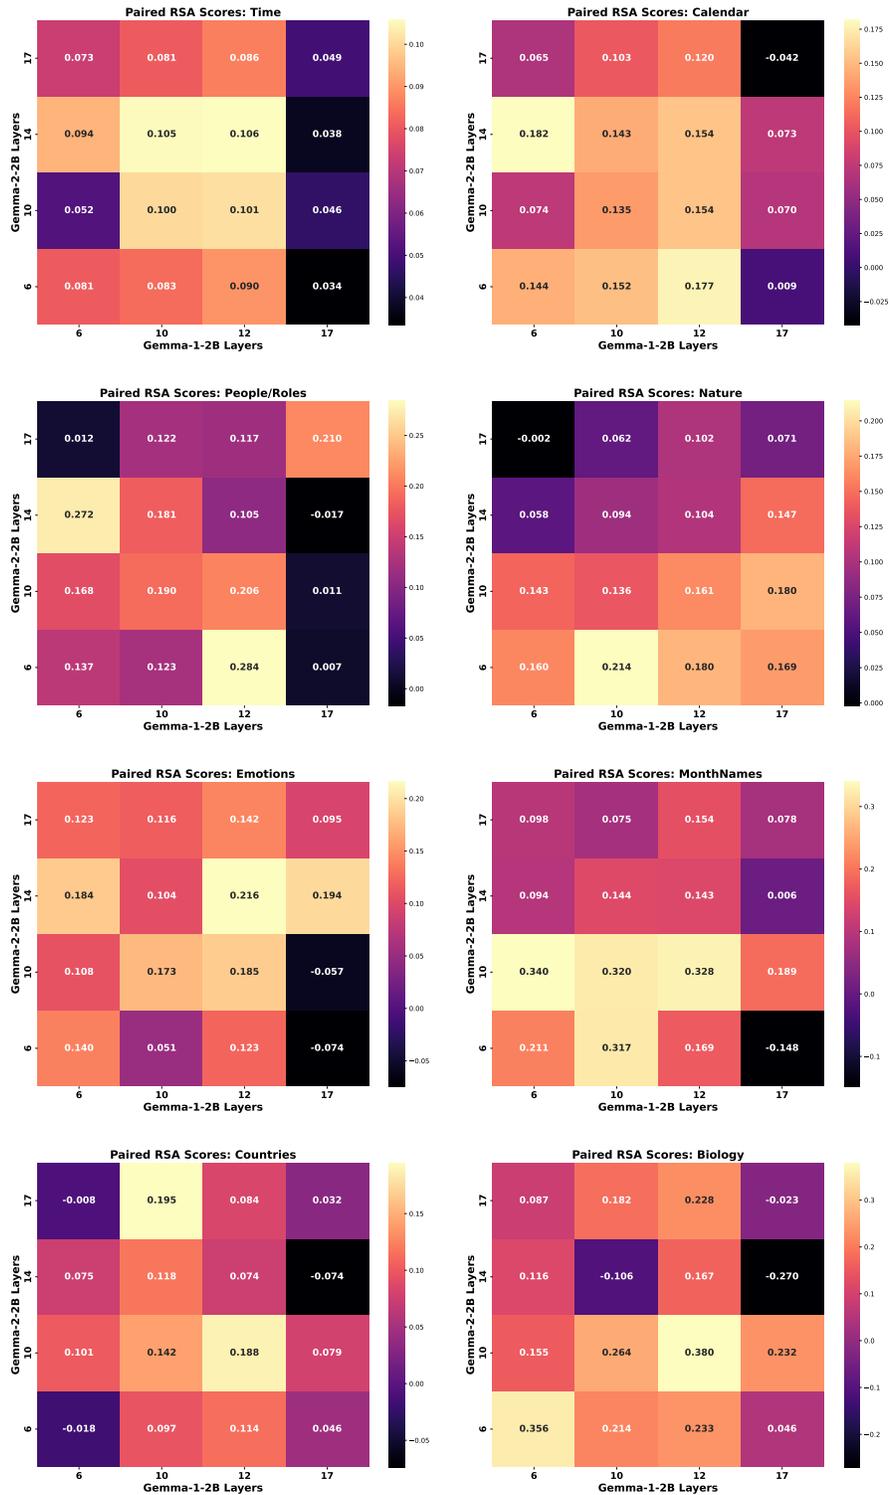


Figure 21: 1-1 Paired RSA scores of SAEs for layers in Gemma-1-2b vs layers in Gemma-2-2b for Concept Categories. Middle layers appear to be the most similar with one another.

1620
1621
1622
1623
1624
1625
1626
1627
1628
1629
1630
1631
1632
1633
1634
1635
1636
1637
1638
1639
1640
1641
1642
1643
1644
1645
1646
1647
1648
1649
1650
1651
1652
1653
1654
1655
1656
1657
1658
1659
1660
1661
1662
1663
1664
1665
1666
1667
1668
1669
1670
1671
1672
1673

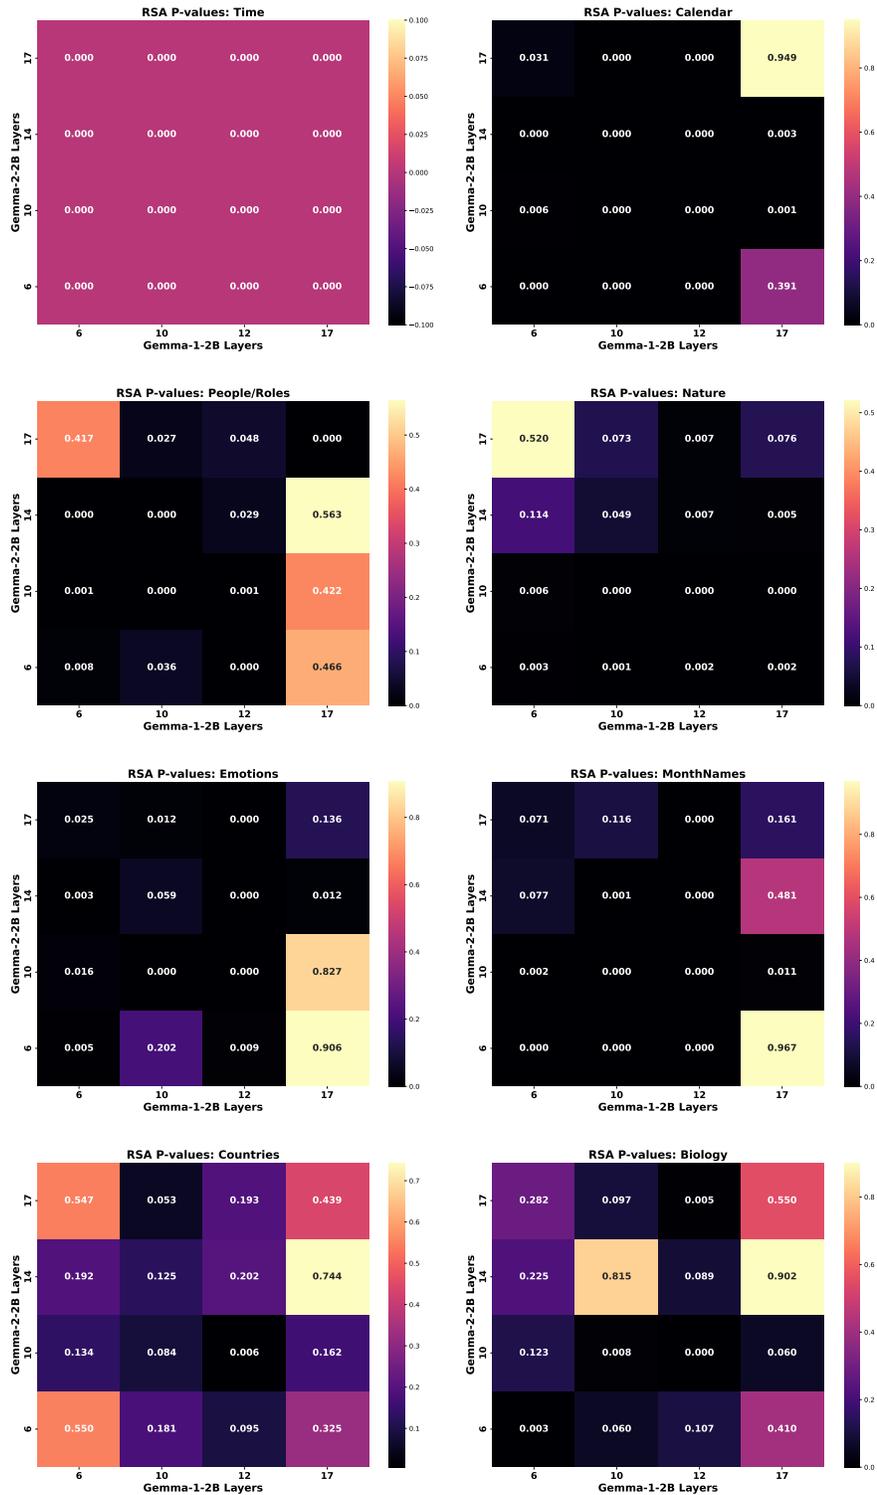


Figure 22: 1-1 RSA p-values of SAEs for layers in Gemma-1-2b vs layers in Gemma-2-2b for Concept Categories. A lower p-value indicates more statistically meaningful similarity.