

PROBING NEURAL TOPOLOGY OF LARGE LANGUAGE MODELS

Anonymous authors

Paper under double-blind review

ABSTRACT

Probing large language models (LLMs) has yielded valuable insights into their internal mechanisms by linking neural activations to interpretable semantics. However, the complex mechanisms that link neuron’s functional co-activation with the emergent model capabilities remains largely unknown, hindering a deeper understanding and safer development of LLMs. In this work, we introduce graph probing, a method for uncovering the functional connectivity of LLM neurons and relating it to language generation performance. By probing models across diverse LLM families and scales, we discover a universal predictability of next-token prediction performance using only neural topology, which persists even when retaining just 1% of neuron connections. **Strikingly, probing on topology outperforms probing on activation by up to 130.4% and 67.7% on perplexity and space/time semantic regression respectively, suggesting that neural topology contains orders of richer information of LLM performance than neural activation, which can be easily extracted with simple linear or MLP probes.** To explain the dependence between neural topology and language performance, we identify default networks and hub neurons in LLMs and provide causal evidence by interventional experiments on multiple benchmarks, showing that LLMs actually exploit these topological information. Further analyses suggest that neural topology can be effectively leveraged to improve the efficiency, reliability, and safety of LLMs through proof-of-concept applications in model pruning, hallucination detection, and LLM fingerprinting. Codes and data for the graph probing toolbox are available at <https://anonymous.4open.science/r/llm-graph-probing-71BD/>.

1 INTRODUCTION

Large language models (LLMs) exhibit remarkable generative capabilities (Wei et al., 2022; Touvron et al., 2023; GLM et al., 2024; Team et al., 2024; Guo et al., 2025), yet our understanding of how they succeed and what they have learned remains limited (Sharkey et al., 2025). **Probing, which extract interpretable features from neural activations (Alain & Bengio, 2017; Rogers et al., 2020; Hewitt & Liang, 2019; Voita & Titov, 2020; Pimentel et al., 2020), has emerged as a powerful approach for reverse-engineering LLMs (Belinkov, 2022; Gurnee & Tegmark, 2024).** For instance, Gurnee *et al.* (Gurnee & Tegmark, 2024) showed that LLMs encode a compact world model of space and time using linear regression probes. Unsupervised probing, such as sparse auto-encoders (Engels et al., 2025; Gao et al., 2025; Rajamanoharan et al., 2024; Lieberum et al., 2024; Mudide et al., 2025) and cross-layer transcoders (Dunefsky et al., 2024; Lindsey et al., 2025), have further revealed dictionaries of interpretable, mono-semantic concepts (Huben et al., 2024) and even causal circuits (Marks et al., 2025), corresponding to directions in neural latent space. While these advances shed light on the semantics of neural activations (Sharkey et al., 2025), much less is known about how neurons are functionally connected, *i.e.* the neural topology, which is believed to play an essential role in the emergence of intelligence (Rathi et al., 2025; Bassett & Sporns, 2017).

Recent studies have drawn compelling parallels between neurons in LLMs and those in the human brain (Toneva & Wehbe, 2019; Schrimpf et al., 2021; Caucheteux et al., 2023; Kumar et al., 2024; Rathi et al., 2025; Mischler et al., 2024; Tuckute et al., 2024; Bonnasse-Gahot & Pallier, 2024; Sun et al., 2024a; Liu et al., 2025), revealing shared properties such as spatial-functional organization (Kumar et al., 2024; Rathi et al., 2025) and left lateralization (Bonnasse-Gahot & Pallier, 2024). Neural activations at internal layers of LLMs have also been shown to reliably predict human

054 brain responses given the same linguistic stimuli (Schrimpf et al., 2021; Tuckute et al., 2024; Luo
 055 et al., 2022). However, these efforts primarily focus on static neural activations of LLMs, while
 056 overlooking the key aspect of temporal and functional neural topology that has been studied in
 057 neuroscience for decades (Bassett & Bullmore, 2006; Bassett & Sporns, 2017; Fotiadis et al., 2024).
 058 Moreover, although analogies between LLMs and human brains are insightful (Toneva & Wehbe,
 059 2019; Goldstein et al., 2022), few works explicitly connect these findings to the language generation
 060 performance, which is one of the primary indicators of an LLM’s intelligence.

061 **In this work, we introduce *graph probing*, a novel approach for investigating the functional connec-**
 062 **tivity of neurons in LLMs and its relationship to language generation and understanding performance.**
 063 By analyzing neural activity time series as LLMs process text token by token, we compute tem-
 064 poral and functional correlations between neurons to construct dynamic neural graphs. **Using this**
 065 **large-scale dataset of text-induced neural topology, we train probes to predict LLMs’ accuracy in**
 066 **auto-regressively generating the corresponding text, as well as how LLMs represent fundamental**
 067 **concepts like space and time. In essence, graph probing connects the micro-level topology of how**
 068 **neurons are connected given a token sequence, to the macro-level performance of how well LLMs**
 069 **understand and predict these tokens, offering a new lens to study the emergent capabilities of LLMs.**
 070 Our method is summarized in Figure 1 and described in Section 2.

071 We then apply our graph probing framework to comprehensively analyze the neural topology of
 072 LLMs through extensive experiments. **First, we demonstrate that language understanding and**
 073 **generation performance can be reliably predicted using only the neural connectivity graph.** This
 074 predictability holds universally across LLM families and scales, outperforming activation-based
 075 probing approaches by up to 130.4% even when preserving only 1% of neuron connections, with
 076 empirical results spanning GPT (Radford et al., 2019), Pythia (Biderman et al., 2023), and Qwen (Yang
 077 et al., 2024), ranging from millions to billions of parameters (Section 3). Next, we show causal
 078 evidence through interventional experiments and analysis on the MMLU (Hendrycks et al., 2021;
 079 Wang et al., 2024) benchmark, discovering stable default neural topology and hub neurons in LLMs,
 080 and more importantly, validating that LLMs actually utilize their internal topological information
 081 when generating text (Section 4). Finally, we offer three proof-of-concept applications of neural
 082 topology, showcasing its potential in model pruning, hallucination detection, and LLM fingerprinting
 083 (Section 5). While not without limitations, we expect graph probing to provide valuable insights into
 084 the inner workings of LLMs and to guide their future development in an interpretable and safe way.

085 2 GRAPH PROBING

087 **Neural Topology.** To construct neural graphs from LLMs, we draw inspiration from neuroscience
 088 where functional brain networks are derived from temporal correlations in fMRI or EEG activation
 089 signals (Bassett & Sporns, 2017; Vértes et al., 2012; Bullmore & Bassett, 2011), as shown in Figure 1.
 090 Given an LLM composed of stacked attention layers, the neural topology is constructed as follows:

091 **Neural Activity:** $\mathbf{H} = \text{HIDDEN_STATE}(\text{LLM}(X)) = [\mathbf{h}_0, \mathbf{h}_1, \dots, \mathbf{h}_t] \in \mathbb{R}^{n \times t}$, (1)

092 **Neural Topology:** $\mathbf{A} = (a_{ij}) \in \mathbb{R}^{n \times n}$, (2)

093
$$a_{ij} = \rho(\mathbf{H}_{i,:}, \mathbf{H}_{j,:}) = \frac{\sum_{k=0}^t (\mathbf{H}_{i,k} - \overline{\mathbf{H}}_{i,:})(\mathbf{H}_{j,k} - \overline{\mathbf{H}}_{j,:})}{\sqrt{\sum_{k=0}^t (\mathbf{H}_{i,k} - \overline{\mathbf{H}}_{i,:})^2} \sqrt{\sum_{k=0}^t (\mathbf{H}_{j,k} - \overline{\mathbf{H}}_{j,:})^2}},$$
 (3)

094 where neurons at each layer produce a time series of hidden states \mathbf{H} as the model processes a token
 095 sequence $X = [x_0, x_1, \dots, x_t]$ (**n and t represent the number of neurons and tokens respectively**),
 096 and the temporal co-activation patterns among neurons define their *functional connectivity*. We
 097 capture this through a complete $n \times n$ weighted connectivity matrix \mathbf{A} , where each node corresponds
 098 to a neuron and each edge weight a_{ij} represents the Pearson correlation coefficient between the
 099 activation time series of a pair of neurons i and j (Bassett & Sporns, 2017; Fotiadis et al., 2024).
 100
 101
 102
 103

104 **Probing on Neural Topology.** We propose *graph probing* to study the dependence between LLM
 105 performance and neural topology. Specifically, we adopt simple linear or multi-layer perceptrons
 106 (MLP) probes (Rumelhart et al., 1986) that take neural topology as input to predict its corresponding
 107 language generation performance, as illustrated in Figure 1. Given a connectivity matrix \mathbf{A} induced
 by feeding a tokenized sequence X to an LLM, where each element a_{ij} denotes the functional

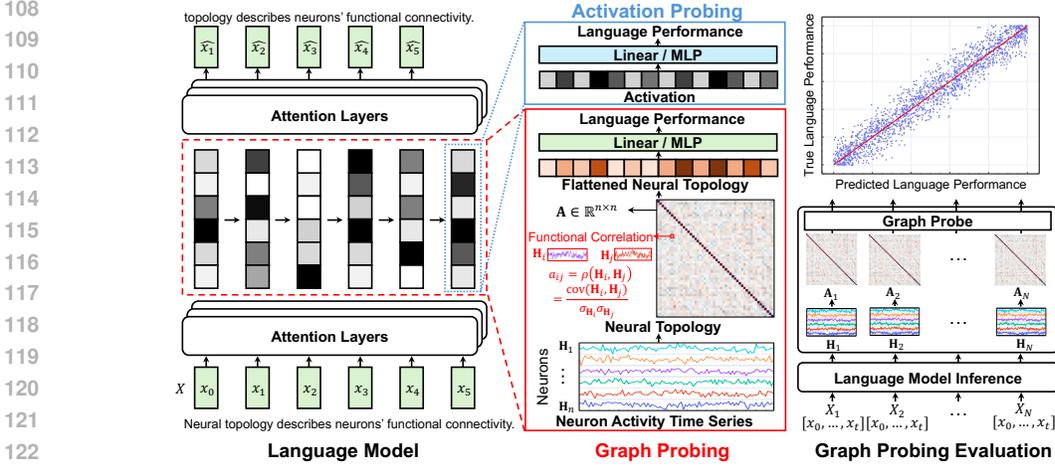


Figure 1: Overview of graph probing. We extract the neuron activity time series in an LLM as it processes text token by token. We then compute temporal and functional correlations between neural activations to obtain topological connectivity graphs of neurons. Unlike existing probing methods that take neural activation as input, we train linear or MLP probes on flattened neural topology to predict the language generation performance for the input token sequence.

connectivity (Pearson correlation coefficient) between neurons i and j , our probe produces the graph representation \mathbf{z} as follows:

$$\text{Linear} : \hat{p} = \mathbf{W}_1 \cdot \text{Flatten}(\mathbf{A})^T, \quad (4)$$

$$\text{MLP} : \hat{p} = \mathbf{W}_3 \cdot \text{ReLU}(\mathbf{W}_2 \cdot \text{Flatten}(\mathbf{A})^T), \quad (5)$$

where $\text{Flatten}(\mathbf{A}) \in \mathbb{R}^{(n \times n)}$ is the flattened topology matrix, $\mathbf{W}_1 \in \mathbb{R}^{1 \times (n \times n)}$, $\mathbf{W}_2 \in \mathbb{R}^{d \times (n \times n)}$, $\mathbf{W}_3 \in \mathbb{R}^{1 \times d}$ are weight matrices (d as a hyper-parameter), \hat{p} is the prediction of language performance. In essence, the core and only difference of our method lies at the input, where we probe on neural topology instead of neural activation by existing approaches (Belinkov, 2022; Gurnee et al., 2023). We will later show that neural topology contains orders of richer information than neural activation regarding LLMs’ intelligence.

Probing Target. By its definition, graph probing can be utilized to predict almost any fact of interest, and in this work we test it on various performance-related metrics, including perplexity, hallucination, and functional specialization, as well as fundamental semantics like space and time. Here we introduce the case of perplexity as it is a fundamental metric directly reflecting LLMs’ language generation performance, while details of other cases are provided in later sections. Specifically, given a token sequence, the LLM’s auto-regressive generation performance is commonly measured by perplexity (Bengio et al., 2003) which corresponds to the exponentiated average negative log-likelihood over the token sequence:

$$\text{Perplexity: } \text{PPL}(X) = \exp \left\{ -\frac{1}{t} \sum_{i=1}^t \log p_{\theta}(x_i | x_{<i}) \right\}. \quad (6)$$

The neural topology is dynamically induced by the specific token sequence, and our goal is to investigate whether the text-responsive neural topology is linked to how well the model predicts the text. Towards this end, we train the graph probe to minimize the mean squared error (MSE) between predicted and true perplexities over a dataset of tokenized sequences $\mathbf{X} = \{X_1, \dots, X_N\}$:

$$\mathcal{L}(\mathbf{X}) = \frac{1}{N} \sum_{i=1}^N (\hat{p}_i - \text{PPL}(X_i))^2, \quad (7)$$

where the prediction \hat{p}_i is calculated via the graph probe by Equation (4) or (5). Details of hyper-parameters and computational configurations are provided in Appendix A.2.

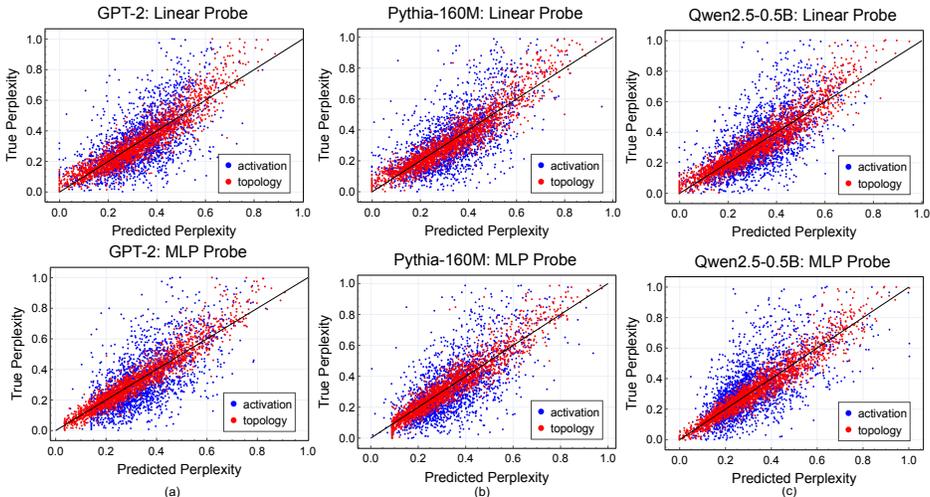


Figure 2: Out-of-sample performance of linear and MLP probing on the test set for (a) GPT-2 (b) Pythia-160M (c) Qwen2.5-0.5B. We compare activation-based probing and our topology-based probing. The correlation between the perplexity predicted by probing and the ground-truth perplexity reflects how well LLM performance can be inferred from neural activation or topology.

3 RESULTS

LLMs. In our experiments, we train graph probes on neural topology derived from three families of LLMs, each spanning across different sizes. Specifically, we evaluate GPT2 (Radford et al., 2019) (GPT2, GPT2-large), Pythia (Biderman et al., 2023) (160M, 1.4B, 2.8B), and Qwen2.5 (Yang et al., 2024) (0.5B, 3B, 7B, 14B). Details of the experimented LLMs are provided in Appendix A.3.

Datasets. To enable our study, we construct neural topology using the OpenWebText dataset (Gokaslan et al., 2019). To ensure consistent temporal resolution, we control the length of neural activity time series to fall between 256 and 1024 tokens by merging consecutive sentences as needed. For each token sequence, we perform LLM inference to compute its perplexity and simultaneously extract hidden state time series to generate the corresponding neural topology. For each model, we construct a probing dataset comprising about 10,000 graph-perplexity pairs. Further details on dataset construction are provided in Appendix A.4.

Evaluation. We split the dataset into training and test sets using an 8:2 ratio. Having learned graph probes on the training set, we evaluate their out-of-sample prediction performance on the test set, which reveals the extent to which micro-level neural topology is predictive of macro-level language generation ability. To quantify the effectiveness of graph probing, we report standard regression metrics on our test data, including mean squared error (MSE), mean absolute error (MAE), coefficient of determination (R^2), Pearson correlation (ρ_p), and Spearman rank correlation (ρ_s). We compare with existing activation based baselines using both linear and MLP probes (Alain & Bengio, 2017; Belinkov, 2022; Gurnee & Tegmark, 2024; Tuckute et al., 2024).

We visualize the predicted and groundtruth perplexity in Figure 2 and summarize the results in Table 1, where graph probing consistently outperforms activation probing across all three LLM families for both linear and MLP probes. The improvements are strikingly substantial, with the maximum progress in R^2 exceeding 130.4%. For example, perplexity is barely predictable from neural activation, with R^2 less than 0.45 for all models, while that of graph probing all close to or even larger than 0.90. The enormous gain of graph probing validates the hypothesis that neural topology contains much richer information of LLMs’ language generation performance than neural activation, which can be easily extracted using simple linear or MLP probes.

Sparsity and scalability. Probing on complete graphs, *i.e.*, dense $n \times n$ connectivity matrices that capture pairwise functional correlations between all neurons, can become computationally prohibitive as the LLM size increases, due to the quadratic number of edges that directly impacts the computational cost in both time and memory. For instance, while complete graph probing is feasible for Pythia-160M with 768 neurons and 0.6M edges per layer, the number of edges in

Table 1: Performance of different probing methods (* indicates p -value < 0.05).

LLM	Probe	MSE ↓	MAE ↓	R^2 ↑	ρ_p ↑	ρ_s ↑
GPT-2	Activation (Linear)	0.0199	0.1067	0.3987	0.6352	0.6497
	Activation (MLP)	0.0201	0.1056	0.3930	0.6370	0.6426
	Graph (Linear)	0.0049*	0.0526*	0.8517*	0.9229*	0.9320*
	Graph (MLP)	0.0031*	0.0399*	0.9057*	0.9534*	0.9577*
Pythia-160M	Activation (Linear)	0.0199	0.1061	0.4151	0.6517	0.6504
	Activation (MLP)	0.0194	0.1030	0.4304	0.6618	0.6519
	Graph (Linear)	0.0047*	0.0518*	0.8607*	0.9278*	0.9364*
	Graph (MLP)	0.0036*	0.0432*	0.8930*	0.9458*	0.9499*
Qwen2.5-0.5B	Activation (Linear)	0.0190	0.1028	0.4225	0.6536	0.6632
	Activation (MLP)	0.0196	0.1045	0.4044	0.6496	0.6457
	Graph (Linear)	0.0045*	0.0496*	0.8640*	0.9296*	0.9422*
	Graph (MLP)	0.0030*	0.0396*	0.9095*	0.9538*	0.9583*

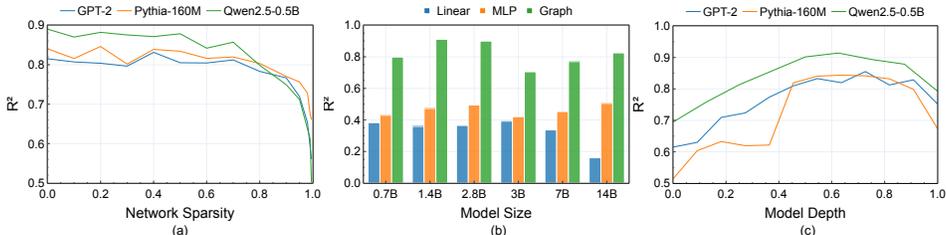


Figure 3: (a) Out-of-sample graph probing performance on neural topology of different sparsity levels. (b) Out-of-sample probing performance on LLMs of different sizes. (c) Out-of-sample performance of graph probing on different layers of LLMs.

Qwen2.5-14B—comprising 5,120 neurons per layer—explodes to over 26M per graph. To address this, on the one hand, we investigate whether perplexity can still be predicted from sparse graphs with weakly correlated edges pruned out by thresholding, which is commonly employed in human brain network construction (Bassett & Sporns, 2017) and such sparsity is frequently observed in artificial neural networks (Vig et al., 2020; Timkey & van Schijndel, 2021; Dettmers et al., 2022); on the other hand, we develop graph neural network probe (Kipf & Welling, 2017) that reduces the number of probe parameters from $\mathcal{O}(n^2)$ to $\mathcal{O}(n \cdot d)$ by weight sharing (details in Appendix A.5). We train graph probes on neural topology with varying levels of sparsification (Figure 3(a)). Surprisingly, the predictive performance remains remarkably stable even after removing up to 90% of the edges, with minimal degradation. Notably, even under extreme sparsity where only 1% of the original edges are retained, the neural topology still enables effective prediction of perplexity, achieving above 0.6 R^2 that is still higher than activation probing.

The above experiments suggest that most of the predictive signal resides in a small subset of strong connections, making it possible to significantly reduce the number of edges while preserving nearly all critical topological information. Leveraging this insight, we scale up graph probing to much larger models by operating on sparsified neural topology. While our earlier results focused on models with fewer than 0.5B parameters, we now train probes on sparse graphs derived from LLMs with up to 14B parameters (GPT2-large, Pythia-{1.4B, 2.8B}, Qwen2.5-{3B, 7B, 14B}). As shown in Figure 3(b), graph probing continues to exhibit strong regression performance across all six large models, achieving a maximum R^2 of over 0.91, providing compelling evidence that the relationship between neural topology and language modeling performance is universal across model sizes. Particularly, the gap between baselines and graph probing remains as substantial as 92.6%, confirming the informativeness of neural topology.

Probing different layers. We further train probes on neural topology derived from different layers of GPT-2, Pythia-160M, and Qwen2.5-0.5B models. Figure 3(c) reveals that while topologies from all layers are predictive of LLM performance, the middle layers are consistently the most informative across all three models, with R^2 values exceeding 0.8. This finding aligns with prior work identifying the middle layers of LLMs as the hub of semantic processing, whereas the initial and final layers are more specialized for {de, re}-tokenization (Gurnee & Tegmark, 2024; Tuckute et al., 2024).

Table 2: Probing results on Arts (time) and World Places (space) datasets. We compare Average Activation and our Graph Probing method (* indicates p -value < 0.05).

LLM	Probe	Arts (time)			World Places (space)		
		MSE	MAE	R^2	MSE	MAE	R^2
GPT-2	Activation (Linear)	0.0406	0.1596	0.2945	0.0360	0.1957	0.5106
	Activation (MLP)	0.0382	0.1511	0.3359	0.0311	0.1727	0.5733
	Graph (Linear)	0.0326*	0.1371*	0.4341*	0.0317*	0.1765*	0.5689*
	Graph (MLP)	0.0267*	0.1230*	0.5361*	0.0258*	0.1514*	0.6486*
Pythia-160M	Activation (Linear)	0.0399	0.1580	0.3071	0.0360	0.1968	0.5116
	Activation (MLP)	0.0391	0.1560	0.3211	0.0306	0.1705	0.5812
	Graph (Linear)	0.0282*	0.1275*	0.5100*	0.0267*	0.1583*	0.6335*
	Graph (MLP)	0.0266*	0.1234*	0.5387*	0.0246*	0.1447*	0.6636*
Qwen2.5-0.5B	Activation (Linear)	0.0325	0.1406	0.4359	0.0296	0.1747	0.6002
	Activation (MLP)	0.0314	0.1368	0.4555	0.0268	0.1574	0.6367
	Graph (Linear)	0.0258*	0.1221*	0.5524*	0.0215*	0.1364*	0.7080*
	Graph (MLP)	0.0268*	0.1254*	0.5341*	0.0223*	0.1398*	0.6953*

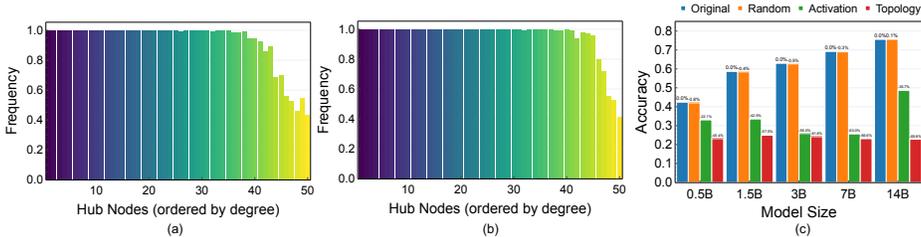


Figure 4: (a-b) Occurrence frequency of hub nodes in (a) Qwen2.5-0.5B and (b) Qwen2.5-1.5B on MMLU benchmark. (c) Accuracy on MMLU benchmark of Qwen2.5 models (0.5B, 1.5B, 3B, 7B, 14B) under different interventions of top 1% neurons.

Probing space and time semantics. Besides predicting perplexity, we also extend our analysis to include a wider range of semantics, probing neural topology to predict time and space context of the input text. we adopt the space and time datasets utilized by Gurnee & Tegmark (2024). We construct the neural topology for input texts regarding historical artworks and world places, and trained probes to predict: (1) time: the release year of historical artworks; (2) space: the longitude and latitude of world places. The results are summarized in Table 2, which serve two critical purposes. Firstly, they validate findings from previous work Gurnee & Tegmark (2024) showing that LLMs possess internal representations of time and space. Secondly and crucially, they confirm that neural topology contains significantly richer information than neural activation for these semantic tasks. The substantial performance margin of as much as 67.77% suggests that the structure of computation (how neurons connect) is more predictive than the magnitude of computation (activations) for extracting high-level concepts.

4 TOPOLOGICAL ANALYSIS AND CAUSAL INTERVENTION

While our probing experiments discover the correlation between neural topology and language generation performance, the specific underlying topological structures remain unclear. In neuroscience, human brains exhibit a stable intrinsic network—a core connectivity pattern that persists to great extent across different tasks and stimuli (Fotiadis et al., 2024; Cole et al., 2014). Inspired by this, we investigated whether a similar default network of highly connected neurons exists within LLMs. To test this hypothesis, we first calculated an average neural topology for the Qwen2.5-0.5B and Qwen2.5-1.5B models over the entire OpenWebText dataset. From this average, we identified the top 50 “hub” neurons (those with the highest degree) and then measured their occurrence frequency, *i.e.*, how often these same neurons ranked as hubs in the topology of each individual data sample. The results, shown in Figure 4(a-b), are striking. A core set of nearly 40 neurons remain hubs in 100% of the topologies across the entire dataset. This remarkable stability confirms the existence of

a default network in LLMs, where a fixed set of hub neurons play dominant roles regardless of the input, suggesting a fundamental organizing principle of their internal structure.

To determine if LLMs rely on their neural topology, we conducted a series of interventional experiments. We disabled a selected 1% of neurons in the middle layer of Qwen2.5 models (0.5B, 1.5B, 3B, 7B, 14B) by pinning their activations to zero at all tokens and measured the impact on the MMLU benchmark (Hendrycks et al., 2021). We compared three distinct neuron selection strategies: random selection, selecting neurons with the highest activation, and selecting neurons with the highest topological degree. As shown in the Figure 4(c), the results demonstrate a clear hierarchy of neuronal importance. While disabling random neurons caused a negligible accuracy drop (<1.0%), targeting neurons by either activation or topology incurred a substantial performance collapse of at least 20%. Critically, the topology-based intervention was the most detrimental. Disabling the top 1% of hub neurons resulted in a catastrophic accuracy drop of 45.4%, 57.5%, 61.6%, 66.6%, and 69.8% for the five tested models, consistently outperforming the activation-based strategy across all scales. Particularly, for the largest 14B model, the relative performance drop against activation based intervention is over 95.5%. These experiments provide causal evidence that LLMs actively utilize their underlying topological structure, particularly their hub neurons, for computation.

Our analysis also revealed a notable phenomenon of functional specialization (Cole et al., 2014), where neural topology displays distinct patterns for different subjects. We found a striking example by identifying certain neuron that consistently specializes in STEM subjects, while remaining almost entirely dormant when processing texts from the social science subjects. This suggests a subject-based organization within LLMs, and the details are provided in Appendix A.6.

5 APPLICATIONS

In this section, we explore the practical implications of neural topology through three proof-of-concept applications: model pruning for efficiency, hallucination detection for reliability, and model fingerprinting for security. Though preliminary in their current forms, these case studies showcase the versatile utility of neural topology and lay the groundwork for future work to translate these concepts into robust, system-level solutions.

5.1 MODEL PRUNING

Model pruning is a standard technique for optimizing Large Language Models (LLMs) in real-world deployments that prioritize low latency and computational efficiency (Sun et al., 2024b; Ma et al., 2023; Gao et al., 2024; Muralidharan et al., 2024; Xia et al., 2024; Xu et al., 2024b). The finding that high-degree neurons are functionally dominant (Section 4) naturally suggests a pruning strategy: disabling low-degree neurons to reduce computational cost. To test this, we prune these neurons by zeroing out their activations during inference. We apply this technique to a middle layer of Qwen2.5-0.5B and Qwen2.5-1.5B and evaluate the performance on the MMLU benchmark (Hendrycks et al., 2021). We compare our topology based pruning approach with a state-of-the-art LLM pruning method WANDA (Sun et al., 2024b) which is based on neural activation. The results, presented in Figure 5(a), confirm the models’ robustness. Performance degradation is minimal even with substantial pruning; for the 1.5B model, accuracy drops by just 2.86% when 10% of neurons are pruned, and by 11.73% when 25% are removed. Strikingly, the models maintain above-random-guess accuracy until half of the low-degree neurons are pruned. In addition, our topology based pruning approach consistently outperforms WANDA, particularly in the 1.5B model where our method retains over 51% accuracy compared to WANDA’s 30% accuracy when pruning 25% of the neurons. These results highlight the potential for developing more sophisticated pruning methods based on neural topology.

5.2 HALLUCINATION DETECTION

Building on our finding that subtle topological patterns in LLMs can predict text generation performance, we now investigate if these patterns can also detect hallucination, a critical challenge that has been extensively explored through various probing techniques (Burns et al., 2023; Du et al., 2024; Hou et al., 2025; Su et al., 2024; Sriramanan et al., 2024; Orgad et al., 2025; Farquhar et al., 2024; Niu et al., 2025). To create distinct *genuine* and *hallucinating* states, we construct inputs from the TruthfulQA dataset (Lin et al., 2022) by concatenating each of its 817 question with its corresponding

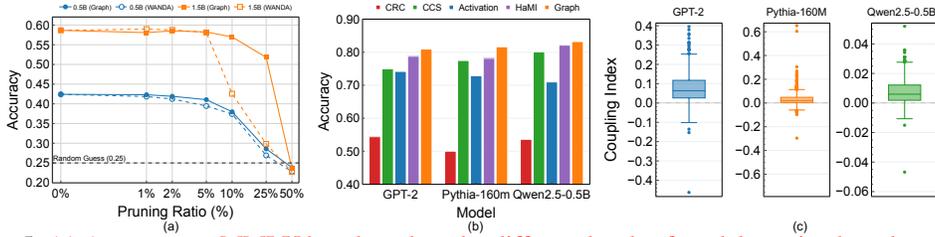


Figure 5: (a) Accuracy on MMLU benchmark under different levels of model pruning based on neural topology and activation (WANDA). (b) Accuracy of hallucination detection of different approaches on TruthfulQA dataset. (c) Coupling index of neural topology for hallucination on TruthfulQA dataset.

true and false answers, respectively, resulting in a dataset of 5918 samples (true/false classification). We then extract the neural topology from the LLM as it processes these inputs (Equations 1-3). For this task, we adapt probing from regression to binary classification by modifying the MLP probe layer to have two output channels ($\mathbf{W}_1 \in \mathbb{R}^{2 \times (n \times n)}$ and $\mathbf{W}_3 \in \mathbb{R}^{2 \times d}$ in Equations (4-5)) and replacing the MSE loss with a cross-entropy loss:

$$\text{Loss} : \mathcal{L}(\mathbf{X}) = \frac{1}{N} \sum_{i=1}^N \text{CROSS_ENTROPY}(\hat{y}_i, y_i), \quad (8)$$

where \hat{y} and y are the prediction and ground-truth for hallucination (0/1). We compare against MLP probes trained on the neural activation at the last-token position. We also include the state-of-the-art HaMI (Niu et al., 2025) and CCS (Burns et al., 2023) hallucination detection approach, as well as CCS’s variant CRC (Burns et al., 2023) for comparison.

We split the dataset into training and test sets using an 8:2 ratio to evaluate our hallucination detection framework. As shown in Figure 5(b), probing neural topology substantially outperforms all baselines probing neural activation for all three tested models (GPT-2, Pythia-160M, and Qwen2.5-0.5B), with accuracy gains of up to 9.73%. This superior performance suggests that distinct topological patterns emerge when an LLM is generating a factual response versus hallucinating. To validate this hypothesis directly, we introduce a neural topology coupling index,

$$C_{TT} = \text{AVG}(\{\rho(A_i, A_j) | A_i, A_j \in A_T\}), C_{HH} = \text{AVG}(\{\rho(A_i, A_j) | A_i, A_j \in A_H\}), \quad (9)$$

$$C_{TH} = \text{AVG}(\{\rho(A_i, A_j) | A_i \in A_T, A_j \in A_H\}), \quad (10)$$

$$C = C_{TT} + C_{HH} - 2C_{TH}, \quad (11)$$

where the index C measures the difference between intra-group similarity (C_{TT} , C_{HH}) and inter-group similarity (C_{TH}). Here A_T and A_H represent the sets of neural topologies for truthful and hallucinated responses, respectively, and similarity is measured by the Pearson correlation (ρ) between each pair of the flattened adjacency matrices. We calculated this index for each question in the dataset, and the distribution in Figure 5(c) shows that over 80% of samples have a positive coupling index. This confirms that topologies are indeed more similar within the same state (truthful-to-truthful, hallucinated-to-hallucinated) than across different states (truthful-to-hallucinated), implying that neural topology serves as a promising and reliable signature for detecting LLM hallucinations and paving the way for future work on improving model reliability.

5.3 MODEL MATCHING AND FINGERPRINTING

In this section, we further investigate potential topological similarity across different LLMs to answer the following question: can we detect the genetic relationships of LLMs (e.g. same model family, or finetuning) from their internal neural topology? To this end, we extend graph probing with contrastive learning to perform *graph matching*, as illustrated in Figure 6. Specifically, suppose we feed a batch of B token sequences into two LLMs, Ω and Γ . We compute the corresponding neural connectivity graphs and use two GNN probes to encode them into representations $\mathbf{Z}_\Omega = [\mathbf{z}_1^\Omega, \dots, \mathbf{z}_B^\Omega]$ and $\mathbf{Z}_\Gamma = [\mathbf{z}_1^\Gamma, \dots, \mathbf{z}_B^\Gamma]$. Matching is implemented using a contrastive cross-entropy loss that encourages

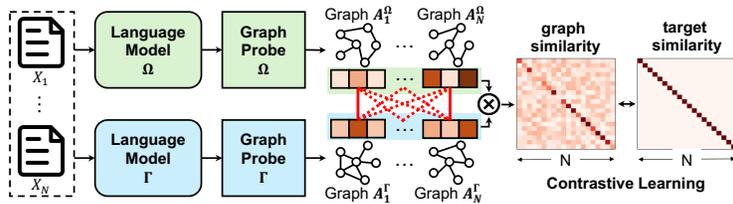


Figure 6: An overview of graph matching. We learn representations of neural topologies derived from two different LLMs processing the same text dataset. We then perform contrastive learning on the graph representations such that matching pairs are more similar by inner product.

alignment between graph representations from the same input texts:

$$\mathcal{S} = \text{MAT_MUL}(\mathbf{Z}_\Omega^T, \mathbf{Z}_\Gamma), \quad \mathcal{T} = \text{IDENTITY}(B), \tag{12}$$

$$\mathcal{L} = \sum_{i=1}^B \text{CROSS_ENTROPY}(\mathcal{S}_{i,:}, \mathcal{T}_{i,:}) + \sum_{j=1}^B \text{CROSS_ENTROPY}(\mathcal{S}_{:,j}, \mathcal{T}_{:,j}), \tag{13}$$

where \mathcal{S} is the similarity matrix by taking inner product of graph representations and IDENTITY indicates identity matrix which is the target \mathcal{T} for graph matching. After training the graph probes contrastively on a shared set of training texts, the out-of-sample graph matching performance serves as an indicator of neural topology similarity between two LLMs. To evaluate this, we adopt the commonly used GAUC metric (Li et al., 2019) (see Appendix A.7 for details).

Table 3 presents the graph matching results. As a sanity check, we first perform *self-matching* using the same LLM, and the results indeed show that GAUC is close to 1.0, validating the rationality of our methodology. We then incorporate two configurations: (1) LLMs within the same family but from different generations, (2) LLMs across different families. Given the reduced architectural and training data differences within the same model family, cross-generation LLMs are *genetically* closer than cross-family LLMs. As expected, cross-generation matching significantly outperforms cross-family matching with the average and maximum GAUC gap up to 4.84% and 9.97%, confirming the effectiveness of graph matching in detecting genetic closeness of LLMs.

Table 3: Graph matching performance (GAUC $\times 100$) between different LLMs.

Matching	LLM Ω	LLM Γ	GAUC
Self	GPT-2	GPT-2	98.64
	Pythia	Pythia	96.92
	Qwen2.5	Qwen2.5	99.24
Generation	Qwen2.5	Qwen2	93.27
	Qwen2.5	Qwen1.5	96.10
	Qwen2	Qwen1.5	94.21
Family	GPT-2	Pythia	92.00
	GPT-2	Qwen2.5	91.11
	Pythia	Qwen2.5	87.39

The graph matching extension enables a direct application in LLM fingerprinting (Xu et al., 2024a; Russinovich & Salem, 2024), which is crucial for protecting intellectual property. To test this, we perform graph matching on the Pythia-160M model, using checkpoints at 1k, 2k, 4k, 8k, and 16k steps of continued training from a base checkpoint (127k). Figure 7 illustrates the graph matching accuracy measured by GAUC, as well as the relative differences of weight values. Despite significant parameter drift, where the weight difference after 16k steps is 5.57 times greater than after 1k steps, the topological signature remains nearly unchanged, with our method achieving a graph matching GAUC above 0.96 in all cases. This suggests that neural topology can serve as a robust fingerprint, resilient to weight modifications from finetuning, which we propose as a significant avenue for future work.

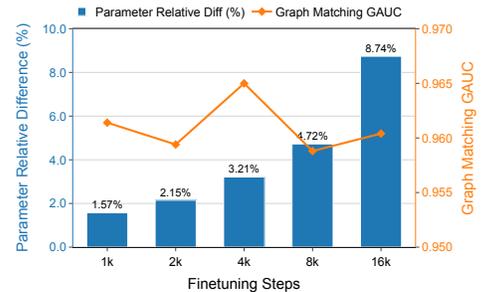


Figure 7: Relative differences of parameters (blue) and graph matching accuracy (orange) of Pythia-160M at different checkpoints.

6 RELATED WORK

Probing LLMs. Growing concerns over the transparency and steerability of LLMs have driven recent advances in reverse-engineering LLMs by extracting interpretable features from their neural activations through probes (Sharkey et al., 2025; Alain & Bengio, 2017; Rogers et al., 2020; Hewitt & Liang, 2019; Voita & Titov, 2020; Pimentel et al., 2020). Supervised probing typically maps neuron activations to interpretable semantics through regression or classification (Gurnee et al., 2023; Gurnee & Tegmark, 2024; Jin & Rinard, 2024; Ju et al., 2024; Dong et al., 2023; Kissane et al., 2024; Templeton et al., 2024; Belinkov, 2022). For example, Gurnee *et al.* (Gurnee & Tegmark, 2024) predicted the time and location of input entities from LLM activations. Unsupervised probing, by contrast, aims to learn a dictionary of disentangled features related to more abstract concepts (Engels et al., 2025; Gao et al., 2025; Rajamanoharan et al., 2024; Lieberum et al., 2024; Mudide et al., 2025; Engels et al., 2025). A famous example is the *Golden Gate Bridge* feature identified in the Claude 3 Sonnet model (Templeton et al., 2024). While prior work focused on connecting LLM activations to external semantics, our work studies the *functional topology* of neurons in LLMs, and relates this internal structure directly to language generation performance via *graph probing*.

Network Neuroscience. The study of functional networks in the human brain has been a central topic in neuroscience for decades (Bassett & Bullmore, 2006; Bassett & Sporns, 2017; Fotiadis et al., 2024; Medaglia et al., 2015) which motivates this research. Brain networks are typically constructed by correlating fMRI or EEG signals across different neural regions, and then analyzed using tools from network science (Barabási, 2013), which has revealed a range of structural and functional properties, such as small-worldness (Bassett & Bullmore, 2006), economical wiring (Bullmore & Sporns, 2012), and functional specialization (Fotiadis et al., 2024). More recently, several studies have drawn parallels between LLM activations and human brain activity (Toneva & Wehbe, 2019; Caucheteux et al., 2023; Kumar et al., 2024; Rathi et al., 2025; Mischler et al., 2024; Tuckute et al., 2024; Bonnasse-Gahot & Pallier, 2024; Sun et al., 2024a; Liu et al., 2025). For instance, Tuckute *et al.* (Tuckute et al., 2024) used GPT-2 activations to identify sentence stimuli that drive or suppress human brain responses. However, while these efforts focus on representational similarities, the *functional topology* of neurons within LLMs and its relationship to the model’s language generation capabilities remain largely unexplored.

7 DISCUSSION

Neurons in LLMs are connected both structurally through the model’s architecture and functionally through their dynamic responses to input linguistic stimuli. **In this work, we focus on the latter and demonstrate that the language understanding and generation performance of LLMs can be reliably predicted from their functional neural topologies using our proposed graph probing approach, implying that LLMs develop intricate and consistent topological structures among their neurons that are fundamental to their emergent linguistic ability.** Besides causal intervention on benchmarks validating LLMs actually leverage their internal neural topology, we also offer practical applications of neural topology in model pruning, hallucination detection, and LLM fingerprinting. While we have empirically shown a stable *default* neural topology and hub neurons in LLMs regardless of input, we have not yet identified more nuanced structures such as motifs, or physical metrics like small-worldness and modularity within these graphs. It remains an open question whether such properties exist in LLMs’ neural topology and play a causal role in shaping their intelligence. Additionally, this paper evaluates LLMs with up to 14B parameters, while leaving even larger models for future work.

Our graph probing results raise many interesting directions for future research. **While we have linked neural topology to general linguistic ability, and discovered functional specialization and subject-specific neurons (Appendix A.6), it remains unclear how these specialized structures emerge during LLMs’ learning process.** Additionally, recent advances in enhancing LLMs’ reasoning abilities (Guo et al., 2025) raise a natural question: does reasoning alter, or is it constrained by, neural topology? Finally, graph probing is model-agnostic and can be extended to models other than LLMs, for example to vision-language models (VLMs). We provide preliminary results probing neural topology of VLMs in Appendix A.8-A.9 with causal intervention as well, while further efforts are required to achieve deeper insights into their multi-modal understanding and generation capabilities. In all, we believe graph probing offers a promising lens for understanding AI models and ultimately guiding their improvement in a reliable and safe way.

ETHICS STATEMENT

The research presented in this paper aims to advance the scientific understanding and interpretability of large language models. Our work relies exclusively on publicly available, pre-trained models (e.g., GPT-2, Pythia, Qwen2.5) and standard academic benchmarks (e.g., MMLU, TruthfulQA, OpenWebText) to ensure reproducibility and avoid the use of sensitive or private data.

We believe our findings have several positive societal implications. The methods for hallucination detection directly contribute to the development of more reliable, truthful, and safe AI systems. Similarly, our work on model pruning promotes computational efficiency, which can reduce the environmental impact of AI and make powerful models more accessible.

REPRODUCIBILITY STATEMENT

To ensure full reproducibility of our findings, we commit to making our research process transparent and accessible. The complete source code, written in Python using the PyTorch and Hugging Face libraries, are released at <https://anonymous.4open.science/r/llm-graph-probing-71BD/>. Our study is built exclusively on publicly available models (GPT-2, Pythia, Qwen2.5) and standard academic benchmarks (MMLU, TruthfulQA, OpenWebText), all of which are accessible through the Hugging Face Hub. All hyper-parameters and experimental settings are detailed in the Appendix.

REFERENCES

- Guillaume Alain and Yoshua Bengio. Understanding intermediate layers using linear classifier probes. In *5th International Conference on Learning Representations, ICLR 2017, Toulon, France, April 24-26, 2017, Workshop Track Proceedings*. OpenReview.net, 2017. URL <https://openreview.net/forum?id=HJ4-rAVtL>.
- Albert-László Barabási. Network science. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 371(1987):20120375, 2013.
- Danielle S Bassett and Olaf Sporns. Network neuroscience. *Nature neuroscience*, 20(3):353–364, 2017.
- Danielle Smith Bassett and ED Bullmore. Small-world brain networks. *The neuroscientist*, 12(6): 512–523, 2006.
- Yonatan Belinkov. Probing classifiers: Promises, shortcomings, and advances. *Computational Linguistics*, 48(1):207–219, 2022.
- Yoshua Bengio, Réjean Ducharme, Pascal Vincent, and Christian Jauvin. A neural probabilistic language model. *Journal of machine learning research*, 3(Feb):1137–1155, 2003.
- Stella Biderman, Hailey Schoelkopf, Quentin Gregory Anthony, Herbie Bradley, Kyle O’Brien, Eric Hallahan, Mohammad Aflah Khan, Shivanshu Purohit, USVSN Sai Prashanth, Edward Raff, et al. Pythia: A suite for analyzing large language models across training and scaling. In *International Conference on Machine Learning*, pp. 2397–2430. PMLR, 2023.
- Laurent Bonnasse-Gahot and Christophe Pallier. fmri predictors based on language models of increasing complexity recover brain left lateralization. In Amir Globersons, Lester Mackey, Danielle Belgrave, Angela Fan, Ulrich Paquet, Jakub M. Tomczak, and Cheng Zhang (eds.), *Advances in Neural Information Processing Systems 38: Annual Conference on Neural Information Processing Systems 2024, NeurIPS 2024, Vancouver, BC, Canada, December 10 - 15, 2024*, 2024. URL http://papers.nips.cc/paper_files/paper/2024/hash/e28e19d00b23fe0265f433fa05a96b06-Abstract-Conference.html.
- Ed Bullmore and Olaf Sporns. The economy of brain network organization. *Nature reviews neuroscience*, 13(5):336–349, 2012.

- 594 Edward T Bullmore and Danielle S Bassett. Brain graphs: graphical models of the human brain
595 connectome. *Annual review of clinical psychology*, 7(1):113–140, 2011.
596
- 597 Collin Burns, Haotian Ye, Dan Klein, and Jacob Steinhardt. Discovering latent knowledge in
598 language models without supervision. In *The Eleventh International Conference on Learning*
599 *Representations, ICLR 2023, Kigali, Rwanda, May 1-5, 2023*. OpenReview.net, 2023. URL
600 <https://openreview.net/forum?id=ETKGuby0hcs>.
- 601 Charlotte Caucheteux, Alexandre Gramfort, and Jean-Rémi King. Evidence of a predictive coding
602 hierarchy in the human brain listening to speech. *Nature human behaviour*, 7(3):430–441, 2023.
603
- 604 Michael W Cole, Danielle S Bassett, Jonathan D Power, Todd S Braver, and Steven E Petersen.
605 Intrinsic and task-evoked network architectures of the human brain. *Neuron*, 83(1):238–251, 2014.
606
- 607 Tim Dettmers, Mike Lewis, Younes Belkada, and Luke Zettlemoyer. Gpt3. int8 (): 8-bit matrix
608 multiplication for transformers at scale. *Advances in neural information processing systems*, 35:
609 30318–30332, 2022.
- 610 Xiangjue Dong, YiBo Wang, Philip S Yu, and James Caverlee. Probing explicit and implicit gender
611 bias through llm conditional text generation. *arXiv preprint arXiv:2311.00306*, 2023.
612
- 613 Xuefeng Du, Chaowei Xiao, and Sharon Li. Haloscope: Harnessing unlabeled llm generations for
614 hallucination detection. *Advances in Neural Information Processing Systems*, 37:102948–102972,
615 2024.
- 616 Jacob Dunefsky, Philippe Chlenski, and Neel Nanda. Transcoders find interpretable llm feature
617 circuits. *Advances in Neural Information Processing Systems*, 37:24375–24410, 2024.
618
- 619 Joshua Engels, Isaac Liao, Eric J Michaud, Wes Gurnee, and Max Tegmark. Not all language model
620 features are linear. In *The Thirteenth International Conference on Learning Representations,*
621 *ICLR 2025, Singapore, Apr 24-28, 2025*. OpenReview.net, 2025. URL <https://openreview.net/forum?id=d63a4AM4hb>.
622
- 623 Sebastian Farquhar, Jannik Kossen, Lorenz Kuhn, and Yarin Gal. Detecting hallucinations in large
624 language models using semantic entropy. *Nature*, 630(8017):625–630, 2024.
625
- 626 Matthias Fey and Jan E. Lenssen. Fast graph representation learning with PyTorch Geometric. In
627 *ICLR Workshop on Representation Learning on Graphs and Manifolds*, 2019.
628
- 629 Panagiotis Fotiadis, Linden Parkes, Kathryn A Davis, Theodore D Satterthwaite, Russell T Shinohara,
630 and Dani S Bassett. Structure–function coupling in macroscale human brain networks. *Nature*
631 *Reviews Neuroscience*, 25(10):688–704, 2024.
- 632 Kunihiko Fukushima. Cognitron: A self-organizing multilayered neural network. *Biological*
633 *cybernetics*, 20(3):121–136, 1975.
634
- 635 Leo Gao, Tom Dupré la Tour, Henk Tillman, Gabriel Goh, Rajan Troll, Alec Radford, Ilya Sutskever,
636 Jan Leike, and Jeffrey Wu. Scaling and evaluating sparse autoencoders. In *The Thirteenth*
637 *International Conference on Learning Representations, ICLR 2025, Singapore, Apr 24-28, 2025*.
638 OpenReview.net, 2025. URL <https://openreview.net/forum?id=tcsZt9ZNKD>.
639
- 640 Shangqian Gao, Chi-Heng Lin, Ting Hua, Zheng Tang, Yilin Shen, Hongxia Jin, and Yen-Chang
641 Hsu. Disp-llm: Dimension-independent structural pruning for large language models. *Advances in*
642 *Neural Information Processing Systems*, 37:72219–72244, 2024.
- 643 Team GLM, Aohan Zeng, Bin Xu, Bowen Wang, Chenhui Zhang, Da Yin, Dan Zhang, Diego Rojas,
644 Guanyu Feng, Hanlin Zhao, et al. Chatglm: A family of large language models from glm-130b to
645 glm-4 all tools. *arXiv preprint arXiv:2406.12793*, 2024.
646
- 647 Aaron Gokaslan, Vanya Cohen, Ellie Pavlick, and Stefanie Tellex. Openwebtext corpus. [http://](http://Skyllion007.github.io/OpenWebTextCorpus)
Skyllion007.github.io/OpenWebTextCorpus, 2019.

- 648 Ariel Goldstein, Zaid Zada, Eliav Buchnik, Mariano Schain, Amy Price, Bobbi Aubrey, Samuel A
649 Nastase, Amir Feder, Dotan Emanuel, Alon Cohen, et al. Shared computational principles for
650 language processing in humans and deep language models. *Nature neuroscience*, 25(3):369–380,
651 2022.
- 652 Daya Guo, Dejian Yang, Haowei Zhang, Junxiao Song, Ruoyu Zhang, Runxin Xu, Qihao Zhu,
653 Shirong Ma, Peiyi Wang, Xiao Bi, et al. Deepseek-r1: Incentivizing reasoning capability in llms
654 via reinforcement learning. *arXiv preprint arXiv:2501.12948*, 2025.
- 655 Wes Gurnee and Max Tegmark. Language models represent space and time. In *The Twelfth Inter-
656 national Conference on Learning Representations, ICLR 2024, Vienna, Austria, May 7-11, 2024*.
657 OpenReview.net, 2024. URL <https://openreview.net/forum?id=jE8xbmvFin>.
- 659 Wes Gurnee, Neel Nanda, Matthew Pauly, Katherine Harvey, Dmitrii Troitskii, and Dimitris Bertsimas.
660 Finding neurons in a haystack: Case studies with sparse probing. *Trans. Mach. Learn. Res.*, 2023,
661 2023. URL <https://openreview.net/forum?id=JYs1R9IMJr>.
- 662 Dan Hendrycks, Collin Burns, Steven Basart, Andy Zou, Mantas Mazeika, Dawn Song, and Jacob
663 Steinhardt. Measuring massive multitask language understanding. In *9th International Conference
664 on Learning Representations, ICLR 2021, Virtual Event, Austria, May 3-7, 2021*. OpenReview.net,
665 2021. URL <https://openreview.net/forum?id=d7KBjmI3GmQ>.
- 667 John Hewitt and Percy Liang. Designing and interpreting probes with control tasks. In Kentaro
668 Inui, Jing Jiang, Vincent Ng, and Xiaojun Wan (eds.), *Proceedings of the 2019 Conference on
669 Empirical Methods in Natural Language Processing and the 9th International Joint Conference on
670 Natural Language Processing (EMNLP-IJCNLP)*, pp. 2733–2743, Hong Kong, China, November
671 2019. Association for Computational Linguistics. doi: 10.18653/v1/D19-1275. URL <https://aclanthology.org/D19-1275/>.
- 672
673 Bairu Hou, Yang Zhang, Jacob Andreas, and Shiyu Chang. A probabilistic framework for llm
674 hallucination detection via belief tree propagation. In *Proceedings of the 2025 Conference of
675 the Nations of the Americas Chapter of the Association for Computational Linguistics: Human
676 Language Technologies (Volume 1: Long Papers)*, pp. 3076–3099, 2025.
- 677 Robert Huben, Hoagy Cunningham, Logan Riggs, Aidan Ewart, and Lee Sharkey. Sparse au-
678 toencoders find highly interpretable features in language models. In *The Twelfth International
679 Conference on Learning Representations, ICLR 2024, Vienna, Austria, May 7-11, 2024*. OpenRe-
680 view.net, 2024. URL <https://openreview.net/forum?id=F76bwRSLeK>.
- 681 Charles Jin and Martin Rinard. Emergent representations of program semantics in language models
682 trained on programs. In *Forty-first International Conference on Machine Learning*, 2024.
- 683
684 Justin Johnson, Bharath Hariharan, Laurens van der Maaten, Li Fei-Fei, C. Lawrence Zitnick, and
685 Ross B. Girshick. CLEVR: A diagnostic dataset for compositional language and elementary visual
686 reasoning. *CoRR*, abs/1612.06890, 2016. URL <http://arxiv.org/abs/1612.06890>.
- 687 Tianjie Ju, Weiwei Sun, Wei Du, Xinwei Yuan, Zhaochun Ren, and Gongshen Liu. How large
688 language models encode context knowledge? a layer-wise probing study. In *Proceedings of the
689 2024 Joint International Conference on Computational Linguistics, Language Resources and
690 Evaluation (LREC-COLING 2024)*, pp. 8235–8246, 2024.
- 691 Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In Yoshua
692 Bengio and Yann LeCun (eds.), *3rd International Conference on Learning Representations, ICLR
693 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*, 2015. URL <http://arxiv.org/abs/1412.6980>.
- 694
695 Thomas N. Kipf and Max Welling. Semi-supervised classification with graph convolutional networks.
696 In *5th International Conference on Learning Representations, ICLR 2017, Toulon, France, April 24-
697 26, 2017, Conference Track Proceedings*. OpenReview.net, 2017. URL <https://openreview.net/forum?id=SJU4ayYgl>.
- 698
699 Connor Kissane, Robert Krzyzanowski, Joseph Isaac Bloom, Arthur Conmy, and Neel Nanda.
700 Interpreting attention layer outputs with sparse autoencoders. *arXiv preprint arXiv:2406.17759*,
701 2024.

- 702 Sreejan Kumar, Theodore R Sumers, Takateru Yamakoshi, Ariel Goldstein, Uri Hasson, Kenneth A
703 Norman, Thomas L Griffiths, Robert D Hawkins, and Samuel A Nastase. Shared functional
704 specialization in transformer-based language models and the human brain. *Nature communications*,
705 15(1):5523, 2024.
- 706 Yujia Li, Chenjie Gu, Thomas Dullien, Oriol Vinyals, and Pushmeet Kohli. Graph matching networks
707 for learning the similarity of graph structured objects. In *International conference on machine*
708 *learning*, pp. 3835–3845. PMLR, 2019.
- 709 Tom Lieberum, Senthoooran Rajamanoharan, Arthur Conmy, Lewis Smith, Nicolas Sonnerat, Vikrant
710 Varma, János Kramár, Anca Dragan, Rohin Shah, and Neel Nanda. Gemma scope: Open sparse
711 autoencoders everywhere all at once on gemma 2. In *Proceedings of the 7th BlackboxNLP*
712 *Workshop: Analyzing and Interpreting Neural Networks for NLP*, pp. 278–300, 2024.
- 713 Stephanie Lin, Jacob Hilton, and Owain Evans. Truthfulqa: Measuring how models mimic human
714 falsehoods. In *Proceedings of the 60th Annual Meeting of the Association for Computational*
715 *Linguistics (Volume 1: Long Papers)*, pp. 3214–3252, 2022.
- 716 Tsung-Yi Lin, Michael Maire, Serge J. Belongie, Lubomir D. Bourdev, Ross B. Girshick, James
717 Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C. Lawrence Zitnick. Microsoft COCO:
718 common objects in context. *CoRR*, abs/1405.0312, 2014. URL [http://arxiv.org/abs/](http://arxiv.org/abs/1405.0312)
719 [1405.0312](http://arxiv.org/abs/1405.0312).
- 720 Jack Lindsey, Wes Gurnee, Emmanuel Ameisen, Brian Chen, Adam Pearce, Nicholas L. Turner,
721 Craig Citro, David Abrahams, Shan Carter, Basil Hosmer, Jonathan Marcus, Michael Sklar, Adly
722 Templeton, Trenton Bricken, Callum McDougall, Hoagy Cunningham, Thomas Henighan, Adam
723 Jermyn, Andy Jones, Andrew Persic, Zhenyi Qi, T. Ben Thompson, Sam Zimmerman, Kelley
724 Rivoire, Thomas Conerly, Chris Olah, and Joshua Batson. On the biology of a large language
725 model. *Transformer Circuits Thread*, 2025. URL [https://transformer-circuits.](https://transformer-circuits.pub/2025/attribution-graphs/biology.html)
726 [pub/2025/attribution-graphs/biology.html](https://transformer-circuits.pub/2025/attribution-graphs/biology.html).
- 727 Haotian Liu, Chunyuan Li, Yuheng Li, and Yong Jae Lee. Improved baselines with visual instruction
728 tuning. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*,
729 pp. 26296–26306, 2024.
- 730 Yiheng Liu, Xiaohui Gao, Haiyang Sun, Bao Ge, Tianming Liu, Junwei Han, and Xintao Hu. Brain-
731 inspired exploration of functional networks and key neurons in large language models. *arXiv*
732 *preprint arXiv:2502.20408*, 2025.
- 733 Zixiang Luo, Kaining Peng, Zhichao Liang, Shengyuan Cai, Chenyu Xu, Dan Li, Yu Hu, Changsong
734 Zhou, and Quanying Liu. Mapping effective connectivity by virtually perturbing a surrogate brain.
735 *arXiv preprint arXiv:2301.00148*, 2022.
- 736 Xinyin Ma, Gongfan Fang, and Xinchao Wang. Llm-pruner: On the structural pruning of large
737 language models. *Advances in neural information processing systems*, 36:21702–21720, 2023.
- 738 Samuel Marks, Can Rager, Eric J Michaud, Yonatan Belinkov, David Bau, and Aaron Mueller.
739 Sparse feature circuits: Discovering and editing interpretable causal graphs in language models.
740 In *The Thirteenth International Conference on Learning Representations, ICLR 2025, Singapore,*
741 *Apr 24-28, 2025*. OpenReview.net, 2025. URL [https://openreview.net/forum?id=](https://openreview.net/forum?id=I4e82CIDxv)
742 [I4e82CIDxv](https://openreview.net/forum?id=I4e82CIDxv).
- 743 John D Medaglia, Mary-Ellen Lynall, and Danielle S Bassett. Cognitive network neuroscience.
744 *Journal of cognitive neuroscience*, 27(8):1471–1491, 2015.
- 745 Gavin Mischler, Yinghao Aaron Li, Stephan Bickel, Ashesh D Mehta, and Nima Mesgarani. Con-
746 textual feature extraction hierarchies converge in large language models and the brain. *Nature*
747 *Machine Intelligence*, pp. 1–11, 2024.
- 748 Anish Mudide, Joshua Engels, Eric J Michaud, Max Tegmark, and Christian Schroeder de Witt.
749 Efficient dictionary learning with switch sparse autoencoders. In *The Thirteenth International Con-*
750 *ference on Learning Representations, ICLR 2025, Singapore, April 24-28, 2025*. OpenReview.net,
751 2025. URL <https://openreview.net/forum?id=k2ZVAzVeMP>.

- 756 Saurav Muralidharan, Sharath Turuvekere Sreenivas, Raviraj Joshi, Marcin Chochowski, Mostofa
757 Patwary, Mohammad Shoeybi, Bryan Catanzaro, Jan Kautz, and Pavlo Molchanov. Compact language
758 models via pruning and knowledge distillation. *Advances in Neural Information Processing
759 Systems*, 37:41076–41102, 2024.
- 760 Mengjia Niu, Hamed Haddadi, and Guansong Pang. Robust hallucination detection in llms via
761 adaptive token selection. *arXiv preprint arXiv:2504.07863*, 2025.
- 762 Hadas Orgad, Michael Toker, Zorik Gekhman, Roi Reichart, Idan Szpektor, Hadas Kotek, and
763 Yonatan Belinkov. Llms know more than they show: On the intrinsic representation of LLM
764 hallucinations. In *The Thirteenth International Conference on Learning Representations, ICLR
765 2025, Singapore, April 24-28, 2025*. OpenReview.net, 2025. URL [https://openreview.
766 net/forum?id=KRnsX5Em3W](https://openreview.net/forum?id=KRnsX5Em3W).
- 767 Tiago Pimentel, Josef Valvoda, Rowan Hall Maudslay, Ran Zmigrod, Adina Williams, and Ryan
768 Cotterell. Information-theoretic probing for linguistic structure. In Dan Jurafsky, Joyce
769 Chai, Natalie Schluter, and Joel Tetreault (eds.), *Proceedings of the 58th Annual Meeting
770 of the Association for Computational Linguistics*, pp. 4609–4622, Online, July 2020. Asso-
771 ciation for Computational Linguistics. doi: 10.18653/v1/2020.acl-main.420. URL [https:
772 //aclanthology.org/2020.acl-main.420/](https://aclanthology.org/2020.acl-main.420/).
- 773 Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, Ilya Sutskever, et al. Language
774 models are unsupervised multitask learners. *OpenAI blog*, 1(8):9, 2019.
- 775 Senthoooran Rajamanoharan, Arthur Conmy, Lewis Smith, Tom Lieberum, Vikrant Varma, János
776 Kramár, Rohin Shah, and Neel Nanda. Improving dictionary learning with gated sparse autoen-
777 coders. *arXiv preprint arXiv:2404.16014*, 2024.
- 778 Neil Rathi, Johannes Mehrer, Badr AlKhamissi, Taha Osama A Binhuraib, Nicholas Blauch, and
779 Martin Schrimpf. Topolm: brain-like spatio-functional organization in a topographic language
780 model. In *The Thirteenth International Conference on Learning Representations, ICLR 2025,
781 Singapore, Apr 24-28, 2025*. OpenReview.net, 2025. URL [https://openreview.net/
782 forum?id=aWXnKanInf](https://openreview.net/forum?id=aWXnKanInf).
- 783 Anna Rogers, Olga Kovaleva, and Anna Rumshisky. A primer in BERTology: What we know about
784 how BERT works. *Transactions of the Association for Computational Linguistics*, 8:842–866, 2020.
785 doi: 10.1162/tacl_a_00349. URL <https://aclanthology.org/2020.tacl-1.54/>.
- 786 David E Rumelhart, Geoffrey E Hinton, and Ronald J Williams. Learning representations by
787 back-propagating errors. *nature*, 323(6088):533–536, 1986.
- 788 Mark Russinovich and Ahmed Salem. Hey, that’s my model! introducing chain & hash, an llm
789 fingerprinting technique. *arXiv preprint arXiv:2407.10887*, 2024.
- 790 Martin Schrimpf, Idan Asher Blank, Greta Tuckute, Carina Kauf, Eghbal A Hosseini, Nancy Kan-
791 wisher, Joshua B Tenenbaum, and Evelina Fedorenko. The neural architecture of language:
792 Integrative modeling converges on predictive processing. *Proceedings of the National Academy of
793 Sciences*, 118(45):e2105646118, 2021.
- 794 Lee Sharkey, Bilal Chughtai, Joshua Batson, Jack Lindsey, Jeff Wu, Lucius Bushnaq, Nicholas
795 Goldowsky-Dill, Stefan Heimersheim, Alejandro Ortega, Joseph Bloom, et al. Open problems in
796 mechanistic interpretability. *arXiv preprint arXiv:2501.16496*, 2025.
- 797 Gaurang Sriramanan, Siddhant Bharti, Vinu Sankar Sadasivan, Shoumik Saha, Priyatham Kattakinda,
798 and Soheil Feizi. Llm-check: Investigating detection of hallucinations in large language models.
799 *Advances in Neural Information Processing Systems*, 37:34188–34216, 2024.
- 800 Weihang Su, Changyue Wang, Qingyao Ai, Yiran Hu, Zhijing Wu, Yujia Zhou, and Yiqun Liu.
801 Unsupervised real-time hallucination detection based on the internal states of large language
802 models. In Lun-Wei Ku, Andre Martins, and Vivek Srikumar (eds.), *Findings of the Association
803 for Computational Linguistics, ACL 2024, Bangkok, Thailand and virtual meeting, August 11-16,
804 2024*, pp. 14379–14391. Association for Computational Linguistics, 2024. doi: 10.18653/v1/2024.
805 FINDINGS-ACL.854. URL [https://doi.org/10.18653/v1/2024.findings-acl.
806 854](https://doi.org/10.18653/v1/2024.findings-acl.854).

- 810 Haiyang Sun, Lin Zhao, Zihao Wu, Xiaohui Gao, Yutao Hu, Mengfei Zuo, Wei Zhang, Junwei Han,
811 Tianming Liu, and Xintao Hu. Brain-like functional organization within large language models.
812 *arXiv preprint arXiv:2410.19542*, 2024a.
- 813 Mingjie Sun, Zhuang Liu, Anna Bair, and J. Zico Kolter. A simple and effective pruning approach for
814 large language models. In *The Twelfth International Conference on Learning Representations, ICLR*
815 *2024, Vienna, Austria, May 7-11, 2024*. OpenReview.net, 2024b. URL <https://openreview.net/forum?id=PxoFut3dWW>.
- 816 Gemma Team, Thomas Mesnard, Cassidy Hardin, Robert Dadashi, Surya Bhupatiraju, Shreya Pathak,
817 Laurent Sifre, Morgane Rivière, Mihir Sanjay Kale, Juliette Love, et al. Gemma: Open models
818 based on gemini research and technology. *arXiv preprint arXiv:2403.08295*, 2024.
- 819 Adly Templeton, Tom Conerly, Jonathan Marcus, Jack Lindsey, Trenton Bricken, Brian Chen,
820 Adam Pearce, Craig Citro, Emmanuel Ameisen, Andy Jones, Hoagy Cunningham, Nicholas L
821 Turner, Callum McDougall, Monte MacDiarmid, C. Daniel Freeman, Theodore R. Sumers,
822 Edward Rees, Joshua Batson, Adam Jermyn, Shan Carter, Chris Olah, and Tom Henighan.
823 Scaling monosemanticity: Extracting interpretable features from claude 3 sonnet. *Trans-*
824 *former Circuits Thread*, 2024. URL <https://transformer-circuits.pub/2024/scaling-monosemanticity/index.html>.
- 825 William Timkey and Marten van Schijndel. All bark and no bite: Rogue dimensions in transformer
826 language models obscure representational quality. In Marie-Francine Moens, Xuanjing Huang,
827 Lucia Specia, and Scott Wen-tau Yih (eds.), *Proceedings of the 2021 Conference on Empirical*
828 *Methods in Natural Language Processing*, pp. 4527–4546, Online and Punta Cana, Dominican
829 Republic, November 2021. Association for Computational Linguistics. doi: 10.18653/v1/2021.
830 emnlp-main.372. URL <https://aclanthology.org/2021.emnlp-main.372/>.
- 831 Mariya Toneva and Leila Wehbe. Interpreting and improving natural-language processing (in
832 machines) with natural language-processing (in the brain). *Advances in neural information*
833 *processing systems*, 32, 2019.
- 834 Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timothée
835 Lacroix, Baptiste Rozière, Naman Goyal, Eric Hambro, Faisal Azhar, et al. Llama: Open and
836 efficient foundation language models. *arXiv preprint arXiv:2302.13971*, 2023.
- 837 Greta Tuckute, Aalok Sathe, Shashank Srikant, Maya Taliaferro, Mingye Wang, Martin Schrimpf,
838 Kendrick Kay, and Evelina Fedorenko. Driving and suppressing the human language network
839 using large language models. *Nature Human Behaviour*, 8(3):544–561, 2024.
- 840 Petra E Vértes, Aaron F Alexander-Bloch, Nitin Gogtay, Jay N Giedd, Judith L Rapoport, and
841 Edward T Bullmore. Simple models of human brain functional networks. *Proceedings of the*
842 *National Academy of Sciences*, 109(15):5868–5873, 2012.
- 843 Jesse Vig, Sebastian Gehrmann, Yonatan Belinkov, Sharon Qian, Daniel Nevo, Yaron Singer, and
844 Stuart Shieber. Investigating gender bias in language models using causal mediation analysis.
845 *Advances in neural information processing systems*, 33:12388–12401, 2020.
- 846 Elena Voita and Ivan Titov. Information-theoretic probing with minimum description length. In
847 Bonnie Webber, Trevor Cohn, Yulan He, and Yang Liu (eds.), *Proceedings of the 2020 Conference*
848 *on Empirical Methods in Natural Language Processing (EMNLP)*, pp. 183–196, Online, November
849 2020. Association for Computational Linguistics. doi: 10.18653/v1/2020.emnlp-main.14. URL
850 <https://aclanthology.org/2020.emnlp-main.14/>.
- 851 Yubo Wang, Xueguang Ma, Ge Zhang, Yuansheng Ni, Abhranil Chandra, Shiguang Guo, Weiming
852 Ren, Aaran Arulraj, Xuan He, Ziyang Jiang, et al. Mmlu-pro: A more robust and challenging multi-
853 task language understanding benchmark. *Advances in Neural Information Processing Systems*, 37:
854 95266–95290, 2024.
- 855 Jason Wei, Yi Tay, Rishi Bommasani, Colin Raffel, Barret Zoph, Sebastian Borgeaud, Dani Yo-
856 gatama, Maarten Bosma, Denny Zhou, Donald Metzler, Ed H. Chi, Tatsunori Hashimoto, Oriol
857 Vinyals, Percy Liang, Jeff Dean, and William Fedus. Emergent abilities of large language mod-
858 els. *Trans. Mach. Learn. Res.*, 2022, 2022. URL <https://openreview.net/forum?id=yzkSU5zdWd>.

864 Mengzhou Xia, Tianyu Gao, Zhiyuan Zeng, and Danqi Chen. Sheared llama: Accelerating language
865 model pre-training via structured pruning. In *The Twelfth International Conference on Learning*
866 *Representations, ICLR 2024, Vienna, Austria, May 7-11, 2024*. OpenReview.net, 2024. URL
867 <https://openreview.net/forum?id=09iOdaeOzp>.
868

869 Jiashu Xu, Fei Wang, Mingyu Ma, Pang Wei Koh, Chaowei Xiao, and Muhao Chen. Instructional
870 fingerprinting of large language models. In *Proceedings of the 2024 Conference of the North Amer-*
871 *ican Chapter of the Association for Computational Linguistics: Human Language Technologies*
872 *(Volume 1: Long Papers)*, pp. 3277–3306, 2024a.

873 Peng Xu, Wenqi Shao, Mengzhao Chen, Shitao Tang, Kaipeng Zhang, Peng Gao, Fengwei An,
874 Yu Qiao, and Ping Luo. BESA: pruning large language models with blockwise parameter-efficient
875 sparsity allocation. In *The Twelfth International Conference on Learning Representations, ICLR*
876 *2024, Vienna, Austria, May 7-11, 2024*. OpenReview.net, 2024b. URL [https://openreview.](https://openreview.net/forum?id=gC6JTEU3jl)
877 [net/forum?id=gC6JTEU3jl](https://openreview.net/forum?id=gC6JTEU3jl).

878 An Yang, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chengyuan Li,
879 Dayiheng Liu, Fei Huang, Haoran Wei, et al. Qwen2. 5 technical report. *arXiv preprint*
880 *arXiv:2412.15115*, 2024.
881
882
883
884
885
886
887
888
889
890
891
892
893
894
895
896
897
898
899
900
901
902
903
904
905
906
907
908
909
910
911
912
913
914
915
916
917

A APPENDIX

A.1 USE OF LLMs

LLMs are used for polishing and proofreading the text. All core ideas were independently conceived by the authors.

A.2 GRAPH PROBING CONFIGURATION

Hyperparameters. We train graph probes using the Adam optimizer (Kingma & Ba, 2015) with mean squared error (MSE) loss, as defined in Equation (7). The initial learning rate is set to 0.00001, with a batch size of 16. We set the hidden dimension for MLP probes d as 32. We apply a learning rate decay strategy, reducing the rate by a factor of 0.1 if the loss does not improve for 5 consecutive epochs. Each model is trained for up to 100 epochs, with early stopping triggered if no improvement is observed for 20 epochs. Dropout is not used, as preliminary experiments showed no significant impact on regression performance.

Computational Resources. LLM inference for computing neural topologies and perplexity scores requires GPUs with large memory. All experiments were conducted on a Linux server equipped with 8 NVIDIA A100 GPUs (80GB memory each). In contrast, training graph probes is relatively lightweight and can be performed on a single GPU with 16GB memory in less than 1 hour.

A.3 EXPERIMENTED LLMs

We run graph probing experiments on a diverse range of LLMs across three different families, with the number of parameters ranging from 124M to 14B. Basic information of these experimented LLMs is summarized in Table 4.

Table 4: Basic information of the experimented LLMs.

LLM family	#params	#layers	#neurons per layer	experimented layer id
GPT-2	124M	12	768	1-12
	774M	36	1280	18
Pythia	160M	12	768	1-12
	1.4B	24	2048	12
	2.8B	32	2560	16
Qwen2.5	0.5B	24	896	1-12
	3B	36	2048	18
	7B	28	3584	14
	14B	48	5120	24

A.4 DATASETS

We conduct graph probing experiments using the OpenWebText dataset (Gokaslan et al., 2019). For each dataset, we randomly sample 10,000 text sequences to construct neural connectivity graphs. Each sample is generated by merging and tokenizing raw text until it reaches a length between 256 and 1024 tokens, which defines the length of the corresponding neural activity time series used for computing pairwise correlations. We then construct a text-responsive neural connectivity graph for each sample and compute its associated perplexity score. To remove outliers that distort the distribution, we filter out the top 1% and bottom 1% of samples based on perplexity. Finally, we normalize all perplexity values to the range $[0, 1]$ by subtracting the minimum perplexity and dividing by the observed range. Summary statistics for the constructed datasets are provided in Table 5.

A.5 GRAPH NEURAL NETWORK PROBE

To reduce the number of probe parameters, we adopt a graph neural network (GNN)-based probe that encodes each node by aggregating neighborhood information through convolutional message passing on the graph (Kipf & Welling, 2017; Fey & Lenssen, 2019). We employ the ReLU activation function (Fukushima, 1975) between graph convolution layers and use average and maximum pooling

Table 5: Basic information of constructed graph probing datasets.

LLM family	#tokens	#graphs	#training graphs	#test graphs
GPT-2	7,020,215	10,384	8,308	2,076
Pythia	6,798,668	10,441	8,353	2,088
Qwen2.5	7,935,555	11,452	9,162	2,290

to summarize node-level embeddings into a graph-level representation. Given a connectivity matrix \mathbf{A} induced by feeding a tokenized sequence X to an LLM, where each element a_{ij} denotes the functional connectivity (Pearson correlation coefficient) between neurons i and j , our probe produces the graph representation \mathbf{z} as follows:

$$\Phi^0 \in \mathbb{R}^{n \times d}, \tag{14}$$

$$\Phi^l = \text{ReLU}(\mathbf{A} \Phi^{l-1} \Theta^l), l = 1, \dots, L, \tag{15}$$

$$\mathbf{z} = \text{AVG_POOLING}\{\Phi_{1,:}^L, \dots, \Phi_{n,:}^L\} \parallel \text{MAX_POOLING}\{\Phi_{1,:}^L, \dots, \Phi_{n,:}^L\}, \tag{16}$$

where $\Phi^0 \in \mathbb{R}^{n \times d}$ denotes the initial learnable node embeddings, $\Theta^l \in \mathbb{R}^{d \times d}$ is the weight matrix of the l -th layer in the GNN with L total layers, and d is a hidden dimensionality hyperparameter. We then feed the graph representation $\mathbf{z} \in \mathbb{R}^{2d}$ into a multi-layer perceptron (MLP) (Rumelhart et al., 1986) to predict the perplexity associated with the input tokenized sequence X :

$$\hat{p} = \mathbf{W}_2 \cdot \text{ReLU}(\mathbf{W}_1 \cdot \mathbf{z}^T), \tag{17}$$

where \hat{p} is the predicted perplexity, and $\mathbf{W}_1 \in \mathbb{R}^{d \times 2d}$, $\mathbf{W}_2 \in \mathbb{R}^{1 \times d}$ are learnable weights of the MLP.

A.6 NEURAL TOPOLOGY ON TEXTS OF DIFFERENT SUBJECTS.

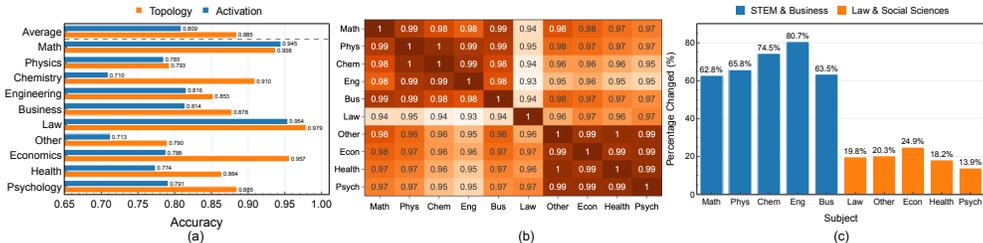


Figure 8: (a) Subject classification accuracy probed from neural topology and activation of Qwen2.5-0.5B model on MMLU benchmark. (b) Correlation of neural topology of Qwen2.5-0.5B on different subjects. (c) Percentage of changed questions in each subject by intervening neuron #894 in layer 12 of Qwen2.5-0.5B.

Beyond the intrinsic default network, we investigated whether LLMs form specialized networks for different knowledge domains, a plausible consequence of their multi-disciplinary pretraining. Such task-invoked networks have been observed in human brains (Cole et al., 2014). To test this, we leveraged the MMLU-Pro benchmark (Wang et al., 2024) to see if a query’s subject can be predicted from its corresponding neural topology. We framed this as a 10-way classification problem using the 10 subjects with the most samples. For computational efficiency, we applied a simple linear probe to the flattened adjacency matrix, optimized with a cross-entropy loss, similar to Section 5.2. The results of Qwen2.5-0.5B model, shown in Figure 8(a), reveal that neural topology is highly indicative of the subject matter. The topology-based linear probe outperformed a baseline probe on neural activations by 9.39% on average. This advantage held true for nearly every individual subject, with the maximum performance gap exceeding 28.16%. These findings strongly suggest that LLMs develop distinct and linearly separable topological patterns for different knowledge domains, allowing for easy extraction of the context or subject being processed.

To provide intuitive evidence for subject-invoked topologies, we calculated the average neural topology for each of the 10 subjects and visualized their pairwise correlations in Figure 8(b). The results are striking: the correlation matrix largely mirrors the conceptual similarities between these

1026 academic fields. While all pairwise correlations exceed 0.9 (reaffirming the existence of a strong
 1027 default network), the variations reveal an intriguing structure. For example, the topologies for Math,
 1028 Physics, and Engineering are highly inter-correlated (>0.98), and all share a much lower correlation
 1029 (<0.95) with Law. More broadly, we observe a clear clustering that separates STEM and social science
 1030 subjects, consistent with commonsense knowledge. To seek causal evidence for this phenomenon,
 1031 we performed a sweeping intervention on individual neurons. We systematically pinned a neuron’s
 1032 activation to a fixed value in $\{-2, -1, 0, 1, 2\}$ and measured the resulting change in the model’s
 1033 output on queries from different subjects. This allowed us to identify neuron #894 in the 12th
 1034 layer of Qwen2.5-0.5B, a "STEM neuron" whose activation state has a disproportionate impact on
 1035 STEM-related subjects. As illustrated in Figure A.6(c), intervening on this single neuron induced
 1036 3.57 times more change in model outputs for STEM queries than for social science queries. While
 1037 preliminary, these correlational and causal findings strongly suggest that functional specialization
 1038 is an emergent property of LLMs, hinting at shared organizational principles between artificial and
 1039 biological intelligence.

1040 A.7 GRAPH MATCHING METRICS

1041
 1042 Given the predicted similarity matrix $\mathcal{S} \in \mathbb{R}^{N \times N}$ and the target similarity matrix $\mathcal{T} = \text{IDENTITY}(N)$,
 1043 we calculate the following metrics for graph matching (Li et al., 2019):

$$1044 \text{GAUC} = \frac{1}{2N} \sum_{i=1}^N (\text{AREA_UNDER_ROC}(\mathcal{S}_{i,:}, \mathcal{T}_{i,:}) + \text{AREA_UNDER_ROC}(\mathcal{S}_{:,i}, \mathcal{T}_{:,i})). \quad (18)$$

1048 A.8 GRAPH PROBING ON VISION LANGUAGE MODELS

1049
 1050 Unlike LLMs, which contain only textual hidden states, Vision-Language Models (VLMs) jointly
 1051 encode image and text features within a shared transformer backbone, producing multimodal hidden
 1052 states that remain underexplored. To study their internal topology, we extract the hidden representa-
 1053 tions at each layer and compute correlations as in Equations (1-3), yielding a connectivity matrix
 1054 whose co-activation time series span both modalities.

1055 We evaluated neural topology in VLMs by probing different sizes of LLaVA-v1.5 (Liu et al., 2024).
 1056 For this setup, we adapted the CLEVR dataset (Johnson et al., 2016) into an object-counting clas-
 1057 sification task of 10,000 randomly sampled images with labels corresponding to the number of
 1058 objects (integers 3-10). GNN-based probes (Equations (14-16)) were trained with cross-entropy loss
 1059 on neural topology graphs constructed at a fixed density of 0.01. As shown in Figure 9(a), graph
 1060 probing consistently outperformed linear activation probing across both the 7B and 13B models,
 1061 validating that neural topology is more informative than neural activation in VLMs regarding visual
 1062 understanding capabilities.

1063 To further evaluate the role of neural topology, we conducted intervention experiments on VLMs on
 1064 the same dataset. Instead of restricting the probe to classification, we asked the model to generate
 1065 numeric outputs and then ablated the top 1% of nodes, selected by degree, activation, or at random,
 1066 by setting their values to zero. Figure 9(b) demonstrates that ablating top nodes by either activation
 1067 or topology reduced accuracy far below random ablation and the original baseline, with topology
 1068 producing the largest drop.

1069 A.9 GRAPH MATCHING ON VISION LANGUAGE MODELS

1070
 1071 We employ graph matching on VLMs to evaluate topological similarity across modalities and to
 1072 test whether multimodal training reorganizes neural topology relative to unimodal language models.
 1073 We use the MS-COCO dataset (Lin et al., 2014), where paired images and captions describe the
 1074 same content, providing a natural basis for structural alignment. Text graphs are constructed by
 1075 masking visual tokens, and image graphs by isolating patch embeddings, and multimodal graphs
 1076 by combining both. We fix graph density at 0.01 and use the same contrastive loss as in the LLM
 1077 matching experiments.

1078 Table 6 reports graph matching scores. As a baseline, multimodal LLaVA graphs compared against
 1079 themselves yield near-perfect alignment (95.98). Matching LLaVA text-only against image-only
 produces a notable drop (81.88), suggesting that while text and image graphs share broad structural

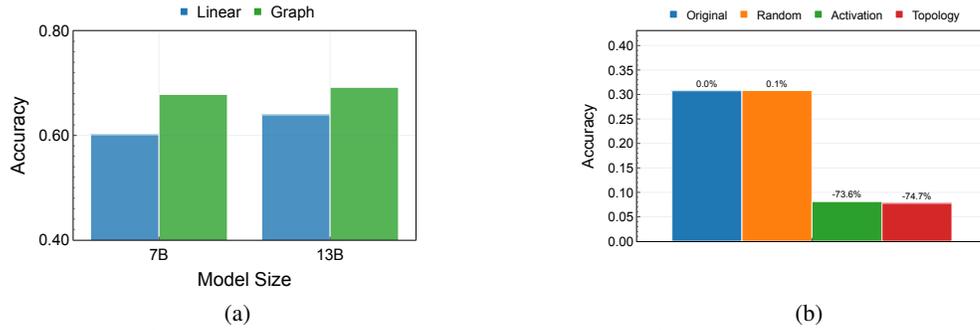


Figure 9: (a) Out-of-sample probing performance on LLaVA-v1.5 of different sizes. (b) Accuracy on CLEVR benchmark of LLaVA-v1.5-7b under different interventions of top 1% neurons.

Table 6: Graph matching performance (GAUC $\times 100$) for different VLMs

VLM Ω Modality	VLM Γ Modality	GAUC
LLaVA Image+Text	LLaVA Image+Text	95.98
LLaVA Text	LLaVA Image	81.88
LLaMA	LLaVA Text	68.03

similarities, each retains specialized organization patterns that multimodal training does not fully unify. In contrast, matching LLaMA text against LLaVA text yields a lower score (68.03), indicating that multimodal training reshapes internal topology beyond unimodal pretraining.