

# Do Large Language Models Think Like the Brain? Sentence-Level Evidences from Layer-Wise Embeddings and fMRI

Yu Lei<sup>1\*†</sup>, Xinyang Ge<sup>2,3\*</sup>, Yi Zhang<sup>4,5</sup>, Yiming Yang<sup>3,6†</sup>, Bolei Ma<sup>4,7†</sup>

<sup>1</sup>School of Artificial Intelligence, Beijing University of Posts and Telecommunications

<sup>2</sup>School of Literature, Shandong University

<sup>3</sup>School of Linguistic Sciences and Arts, Jiangsu Normal University

<sup>4</sup>Department of Statistics, LMU Munich

<sup>5</sup>FAU Erlangen-Nuremberg

<sup>6</sup>Collaborative Innovation Center for Language Ability, Jiangsu Normal University

<sup>7</sup>Munich Center for Machine Learning

leiyu@bupt.edu.cn, 202320175@mail.sdu.edu.cn, yi.zhang@fau.de, yangym@jsnu.edu.cn, bolei.ma@lmu.de

## Abstract

Understanding whether large language models (LLMs) and the human brain converge on similar computational principles remains a fundamental and important question in cognitive neuroscience and AI. Do the brain-like patterns observed in LLMs emerge simply from scaling, or do they reflect deeper alignment with the architecture of human language processing? This study focuses on the sentence-level neural mechanisms of language models, systematically investigating how layer-wise representations in LLMs align with the dynamic neural responses during human sentence comprehension. By comparing hierarchical embeddings from 14 publicly available LLMs with fMRI data collected from participants, who were exposed to a naturalistic narrative story, we constructed sentence-level neural prediction models to identify the model layers most significantly correlated with brain region activations. Results show that improvements in model performance drive the evolution of representational architectures toward brain-like hierarchies, particularly achieving stronger functional and anatomical correspondence at higher semantic abstraction levels. These findings advance our understanding of the computational parallels between LLMs and the human brain, highlighting the potential of LLMs as models for human language processing.

**Code** — <https://github.com/Lucasuuu02/LLM4Brain>

## Introduction

The intersection of artificial intelligence and neuroscience has emerged as a cutting-edge research frontier, particularly in understanding the parallels between large language models (LLMs) and human neural language processing (Toneva and Wehbe 2019; Liu et al. 2025; Schrimpf et al. 2021). Prior studies (Anderson et al. 2021; Caucheteux, Gramfort, and King 2021) have established evidence showing intriguing correlations between LLM-learned representations and neural responses during language processing (Zhao et al. 2025),

\*These authors contributed equally.

†Corresponding authors.

Copyright © 2026, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.




Figure 1: A brief presentation of the experimental design. We expose both LLMs and real humans to a narrative story (“The Little Prince”) and aim to compare the correlation between LLMs’ and human brains’ language processing.

particularly in feature extraction and representational similarity (Caucheteux and King 2022; Hosseini et al. 2024). These findings (Sun et al. 2021) suggest that both systems may utilize comparable linguistic features, as evidenced by the linear mappability of LLM representations to neural activity patterns. However, these observations lack mechanistic explanations for what critical properties enable LLMs to achieve brain-like processing capabilities.

Recent investigations (Goldstein et al. 2022; Caucheteux, Gramfort, and King 2023) have explored multimodal similarities between LLMs and neural processes. While some studies (Antonello, Vaidya, and Huth 2023; Antonello and Huth 2024) demonstrate stronger alignment between autoregressive LLMs (Ethayarajh 2019; Tenney, Das, and Pavlick 2019; Lei et al. 2024) and the predictive coding hypothesis of human language processing, others focus on metrics like language modeling performance, model scale, and representational generalization (Liu, Lei, and Wu 2025). Collectively, these works indicates that modeling quality critically determines brain-like representational capacities (Hickok and Poeppel 2007; Lerner et al. 2011). A fundamental question remains unresolved: *Does this similarity merely stem from increased model scale, or does it reflect deeper convergence in computational principles with the human speech processing pathway?* Resolving this dichotomy is crucial for advancing next-generation model architectures.

Current studies examining the relationship between

LLMs and brain decoding typically use publicly available datasets to evaluate model performance in open scenarios. However, such datasets may fail to accurately reflect the models' comprehension abilities in specific tasks. To address this limitation, we selected 14 publicly available pre-trained LLMs and designed a sentence understanding task to assess their contextual understanding using the same text materials used for human fMRI. Additionally, we utilized fMRI data from participants who listened to naturalistic text to construct decoding models, using the correlation metric to compare the relationship between LLMs and brain-activity patterns in relevant regions. A simple presentation of our experimental design is illustrated in Figure 1.

We summarize our main findings as follows:

- 1) **Instruction tuning boosts performance and brain alignment:** Instruction-tuned models consistently outperformed their base versions in both sentence comprehension and neural alignment, demonstrating statistically significant improvements.
- 2) **Intermediate layers better reflect brain activity:** All LLMs showed stronger correlations with brain responses at intermediate layers, and models with higher comprehension ability exhibited stronger neural alignment.
- 3) **Hemispheric asymmetry reveals functional specialization:** Left-dominant regions aligned with core language processing, while right-hemisphere regions reflected higher-level cognitive functions, both shaping model-brain correspondence.

## Related Work

**Neuroscientific Foundations of Language Comprehension.** Before the advent of the current large-scale LLMs, cognitive neuroscience had established that human language comprehension relies on complex hierarchical processing (Friederici 2011). Research centered on understanding how the brain integrates perceptual units (e.g., phonemes and graphemes) into meaningful structures at the lexical, sentential, and discourse levels (Price 2012). Early models, such as the Wernicke-Lichtheim-Geschwind model (Geschwind 1967), identified localized brain regions for language but struggled to explain sentence-level and discourse-level processing, especially information integration across sentences for coherence (Hickok and Poeppel 2004). Cognitive Neuroscience Methods like fMRI and ERP (Kutas and Hillyard 1984; Osterhout and Holcomb 1992) led to a shift from localized to distributed network models, such as the Memory-Unification-Control (MUC) framework (Hagoort 2016). However, early fMRI studies were limited by isolated stimulus presentation and signal averaging, which made capturing long-range linguistic integration difficult (Humphries et al. 2007). New naturalistic paradigms like narrative listening (Brennan 2016) advanced investigation into discourse comprehension, but approaches like representational similarity analysis (RSA) (Kriegeskorte, Mur, and Bandettini 2008) often underestimate discourse features like coherence and context-aware meaning (Zacks, Mar, and Calarco 2017; Messi and Pyllkanen 2025), emphasizing the need for multi-sentence integration analyses.

**Leveraging LLMs for Brain-Language Mapping.** Recent studies (Luo, Xu, and Xiong 2022; Yu et al. 2024; Fu et al. 2025) use the strong semantic capabilities of pretrained language models to examine brain-language mappings and decode neural processes. For instance, Ren et al. (2025); Yu et al. (2024) applied Dynamic Similarity Analysis (DSA) to compare text embeddings with fMRI signals, constructing Representational Dissimilarity Matrices (RDMs) using measures like Pearson correlation. Other works (Mischler et al. 2024; Bonnasse-Gahot and Pallier 2024) aligned layer-wise activations of language models with averaged fMRI activity maps using ridge regression. Moreover, Tuckute et al. (2024) trained encoding models on fMRI data from participants exposed to diverse sentences, optimizing GPT-2 XL embeddings for neural alignment. There is also evidence that high-level visual representations in the human brain are aligned with LLMs (Doerig et al. 2025). These studies showcase the potential of LLMs to reveal insights into the neural mechanisms of language comprehension, offering new tools and methodologies for cognitive neuroscience.

## Methodology

To examine the alignment between LLM representations and brain activity during naturalistic language comprehension, we implement a multi-stage pipeline (Figure 2). We first preprocess the fMRI data and use a General Linear Model (GLM) to estimate sentence-level neural responses. Language-related ROIs are then extracted for analysis. We compute sentence embeddings from different LLM layers and evaluate their correspondence with brain activity using voxel-wise encoding models.

## Data and Stimuli

We used the existing data from Li et al. (2022) which provides the “The Little Prince” multilingual naturalistic fMRI corpus. The dataset included 35 healthy, right-handed, native Mandarin speakers (15 female; mean age 19.3, range 18–25), but one participant was excluded due to incomplete data. The stimuli used are parallel corpora of the diary “The Little Prince” in both English and Chinese. The texts are segmented into 1,577 sentences in Chinese, with corresponding English sentences aligned for each one.

## fMRI Acquisition and Preprocessing

Stimuli were presented via PsychoPy 2 and MRI-compatible headphones. Participants listened to a 99-min Chinese audiobook (nine 10-min segments) and answered four questions after each (36 total) via head-coil mirror and button box. Total session 2.5 h.

Data acquisition used a 3T GE Discovery MR750 scanner (32-channel coil). Structural images: T1-weighted MP-RAGE sequence; Functional images: multi-echo planar imaging (ME-EPI) sequence (TR = 2000 ms; TEs = 12.8, 27.5, 43 ms; voxel size = 3.75 × 3.75 × 3.8 mm). Preprocessing was performed using AFNI 16 (Cox, 1996), including removal of the first 4 time points, multi-echo independent component analysis (ME-ICA) for denoising, and spatial normalization to the MNI standard space (resampled to 2 mm<sup>3</sup> voxels).

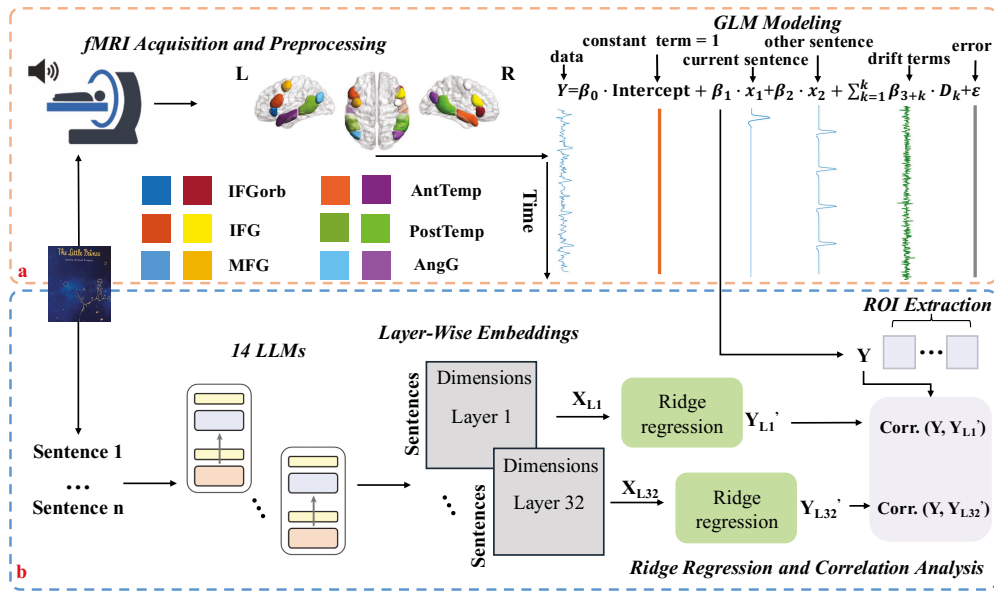


Figure 2: Multi-stage pipeline to analyze the alignment between LLM representations and neural responses during naturalistic language comprehension. The methodology includes auditory stimulus presentation, layer-wise embedding extraction from LLMs, voxel-wise regression modeling, and region-of-interest (ROI)-based brain-model alignment analysis. Panel (a) outlines the neuroimaging data acquisition and preprocessing steps, while panel (b) describes the brain-LLM alignment analysis.

### General Linear Model (GLM)

The general linear model (GLM) serves as the standard analytical framework in fMRI neuroscience research, employing statistical modeling to accurately isolate task-evoked Blood oxygenation level dependent (BOLD) signals from physiological and system noise. The core model is:

$$\mathbf{Y} = \mathbf{X}\boldsymbol{\beta} + \boldsymbol{\epsilon}, \quad \boldsymbol{\epsilon} \sim \mathcal{N}(0, \sigma^2\mathbf{I}) \quad (1)$$

where  $\mathbf{Y}$  represents continuous BOLD time series sampled at each TR,  $\mathbf{X}$  is the design matrix encoding experimental conditions and nuisances,  $\boldsymbol{\beta}$  contains regression coefficients, and  $\boldsymbol{\epsilon}$  is residual noise.

The design matrix  $\mathbf{X}$  is constructed by convolving event onsets with a canonical hemodynamic response function (HRF). Crucially, this convolution process uses the precise onset time of each sentence, not the TR, thereby resolving the temporal mismatch between sentence boundaries and the discrete TR sampling intervals. Each condition’s regressor  $\mathbf{x}_j$  includes the trial onsets and nuisance confounds (e.g., motion, drift):

$$\mathbf{x}_j(t) = \sum_{k=1}^{K_j} \text{HRF}(t - t_{\text{onset}}^{(j,k)}) + \sum_{m=1}^M \gamma_m \mathbf{n}_m(t) \quad (2)$$

We adopt the Least-Squares Separate (LS-S) approach (Mumford, Davis, and Poldrack 2014) to isolate BOLD responses for each sentence-level trial by modeling each target sentence independently. This method reduces collinearity and improves sensitivity to transient, event-related activity compared to block-based or condition-averaged models. LS-S enables precise sentence-level neural activation patterns under naturalistic stimuli.

Voxel-wise parameter estimates are obtained as:

$$\hat{\boldsymbol{\beta}} = (\mathbf{X}^\top \mathbf{X})^{-1} \mathbf{X}^\top \mathbf{Y} \quad (3)$$

$t$ -statistics are computed for testing contrasts  $\mathbf{c}$ :

$$t_v = \frac{\mathbf{c}^\top \hat{\boldsymbol{\beta}}_v}{\sqrt{\hat{\sigma}_v^2 \cdot \mathbf{c}^\top (\mathbf{X}^\top \mathbf{X})^{-1} \mathbf{c}}}, \quad \hat{\sigma}_v^2 = \frac{\|\mathbf{Y}_v - \mathbf{X} \hat{\boldsymbol{\beta}}_v\|^2}{T - \text{rank}(\mathbf{X})} \quad (4)$$

This enables inference on condition-specific activations while accounting for noise and drift, improving sensitivity and interpretability of neural responses to linguistic stimuli.

### Region of Interest (ROI) Extraction

We then extract the ROIs using the language network proposed by (Fedorenko et al. 2010), which serves as the core analytical framework for our investigation into the neurobiological distinctions between human brains and LLMs during natural language comprehension. This network has been widely validated as the neural substrate for representing linguistic knowledge and supporting key language processes (Tuckute et al. 2024).

### Layer-Wise Embeddings of Sentences

We then introduce the LLM processing of “The Little Prince”, as shown in Figure 2(b). The text is segmented into sentences and then fed into the LLMs.

We extract layer-wise embeddings from 14 language models of varying depths. Each sentence is processed through model-specific sliding windows, capturing all hidden states at terminal positions. Segment representations are averaged into layer  $\times$  dimension tensors, preserving native dimensionality per model for neural decoding.

## Ridge Regression and Correlation Analysis

In this framework, we employ ridge regression to quantify how semantic representations from different layers of a neural network relate to brain activity patterns captured via fMRI. By integrating deep learning and neuroscience, we systematically evaluate how well neural embeddings predict fMRI signals for a specific region of interest (ROI). Given fMRI response vectors  $\mathbf{y} \in \mathbb{R}^N$  and neural embeddings  $\mathbf{X} \in \mathbb{R}^{N \times L \times D}$  (where  $T$  is the number of time points,  $L$  is the number of layers, and  $D$  denotes embedding dimensionality), the predictive performance for each layer is computed as the average Pearson correlation ( $\rho$ ) across  $K$  cross-validation folds:

$$\rho_l = \frac{1}{K} \sum_{k=1}^K \text{corr}(\mathbf{y}_{\text{test}}^{(k)}, \mathbf{X}_{\text{test}}^{(l,k)} \hat{\boldsymbol{\beta}}^{(l,k)}) \quad (5)$$

The model uses ridge regression to estimate these predictions, balancing data fidelity and regularization to avoid overfitting. Regression weights ( $\boldsymbol{\beta}$ ) are optimized with:

$$\hat{\boldsymbol{\beta}}^{(l,k)} = (\mathbf{X}_{\text{train}}^{(l,k)\top} \mathbf{X}_{\text{train}}^{(l,k)} + \alpha \mathbf{I})^{-1} \mathbf{X}_{\text{train}}^{(l,k)\top} \mathbf{y}_{\text{train}} \quad (6)$$

Here,  $\alpha$  controls the regularization strength, selected via grid search in a nested cross-validation framework. All data are z-scored to ensure compatibility between fMRI and neural embeddings.

**Feature Normalization and Parallelization.** Both  $\mathbf{y}$  and  $\mathbf{X}$  are standardized to zero mean and unit variance before regression to improve numerical stability. The analysis pipeline parallelizes computations across subjects, ROIs, and layers, training  $\mathcal{O}(SRL)$  models, where  $S$  is the number of subjects,  $R$  is the number of ROIs, and  $L$  is the number of layers. This ensures efficient scalability in the layer-wise evaluation of embeddings.

This approach provides layer-wise correlation metrics ( $\rho_l$ ), which can offer insights into how neural network representations align with ROI-specific brain activity.

## Experiments and Results

Now we present our three main experiments and results to evaluate the cross-lingual semantic alignment capabilities of LLMs, their correspondence with brain activity, and their sensitivity to hemispheric differences in neural responses.

### Performance of Large Language Models

To evaluate the contextual understanding capabilities of LLMs under cross-lingual settings, we first design a multiple-choice test based on semantic alignment between Chinese source sentences and their English translations, using the stimuli sentences of ‘‘The Little Prince’’ used for the human fMRI data. Drawing from recent work on evaluating the robustness of LLMs in multiple choice setups (Zheng et al. 2024; Wang et al. 2024), we adopt several perturbations to test the LLMs’ performance. For each original Chinese sentence from the stimuli, we generate five English options: (A) correct English translation, (B) word order scrambled, (C) part-of-speech substitution, (D) sentence structure transformation, (E) information insertion/deletion. Each of

these English sentence options is then input into the LLM to obtain their corresponding embeddings.

To quantitatively evaluate the cross-lingual semantic understanding capability of large language models, we propose the **Cross-lingual Semantic Alignment Accuracy (CSAA)** metric. This metric is defined as the proportion of cases in which the model correctly identifies the true translation (option A) as the most semantically similar candidate to the original Chinese sentence, as measured by cosine similarity in the embedding space.

We define an indicator variable  $\delta_i$  for each sample  $i$  as follows:

$$\delta_i = \begin{cases} 1, & \text{if } \operatorname{argmax}_x (\cos(\mathbf{v}_{c_i}, \mathbf{v}_{e_{i,x}})) = A \\ 0, & \text{otherwise} \end{cases} \quad (7)$$

The CSAA is then defined as:

$$\text{CSAA} = \frac{1}{N} \sum_{i=1}^N \delta_i \quad (8)$$

where  $N$  is the total number of Chinese sentences,  $\mathbf{v}_{c_i}$  denotes the embedding of the  $i$ -th Chinese sentence, and  $\mathbf{v}_{e_{i,x}}$  denotes the embedding of the  $x$ -th English candidate for the  $i$ -th sentence. A higher CSAA indicates stronger cross-lingual semantic alignment ability of the model.

The results in Figure 3 show that the Llama-3.1-Instruct version leads with 31.4 points, followed by two Gemma variants (30.7 and 30.3). A sharp performance drop occurs after the top three models, with mid-tier scores (22.7-15.1) dominated by Baichuan2 and DeepSeek variants, while glm-4.9b (7.9) and Qwen2.5-7B (6.4) anchor the lower end. The red-to-green gradient visually reinforces score disparities, highlighting a >28-point gap between best and worst performers. Consistently, we notice that the instruction-tuned models always perform better than the base models in each model family.

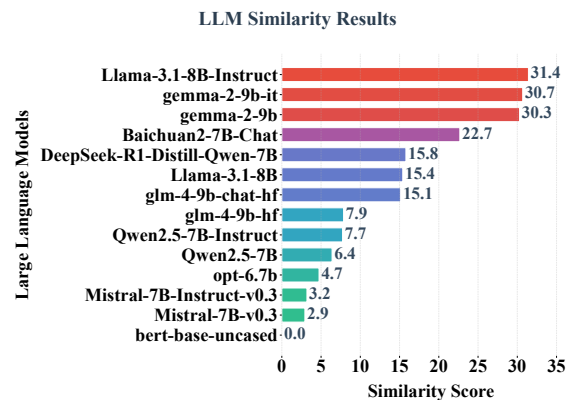


Figure 3: Performance comparison of 14 LLMs

This experiment demonstrates that LLMs exhibit varying degrees of cross-lingual semantic alignment capabilities. Instruction tuning consistently improves model performance, showing the importance of fine-tuning and alignment strategies in enhancing semantic understanding. These findings establish the foundation for further investigations into how semantic alignment translates into neural representations.

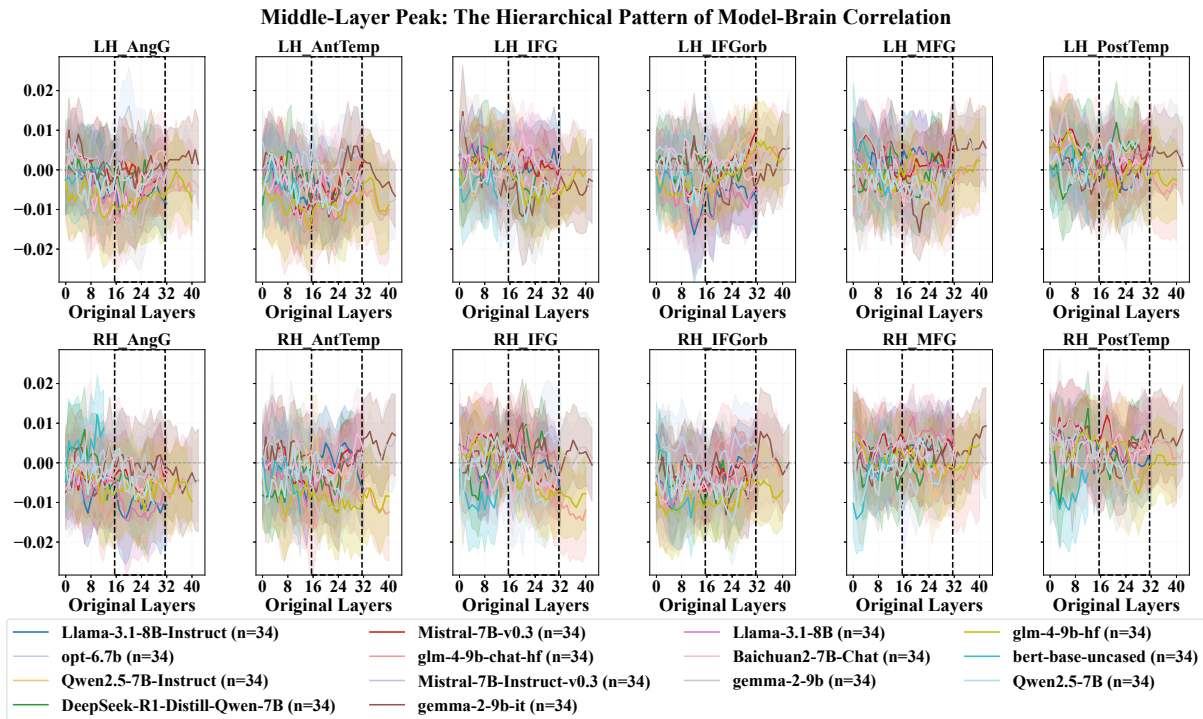


Figure 4: Correlation between model predictions and brain activity across layers. Shaded areas are 95% confidence intervals.

### Model Performance and Activation Correlation

Next, to comprehensively assess the correspondence between models and brain activity, we compared model performance across multiple brain regions and layers.

The comparative evaluation of various computational models across all brain regions is summarized in Figure 4, which displays the layer-wise correlation coefficients together with their 95% confidence intervals. Each curve represents the average performance trajectory of a given model, with shaded regions indicating confidence bounds. Notably, we found that in most cases, models achieve peak predictive performance at their intermediate layers, not the final layer, consistent with previous findings (Mischler et al. 2024).




Figure 5: Average correlation of LLMs across 12 ROIs.

To further compare models, we selected the optimal layer exhibiting peak performance as the model’s output representation and examined the average correlation across 12 ROIs,

as shown in Figure 5. Our results show that LLMs consistently yield higher correlation metrics than the base BERT model across brain regions.

From the 14 LLMs, we selected 5 instruction-tuned models (which have their corresponding base versions) along with their 5 corresponding base versions for comparative analysis. The selection was motivated by the intent to directly compare the impact of instruction tuning. Figure 6(a) presents the comparison of activation correlations, while Figure 6(b) shows the performance differences between base and instruction-tuned versions. We notice that instruction-tuned models exhibit performance improvements in both Correlation and Performance metrics when compared with the base models. Specifically, the p-value for the Correlation Change indicates a trend toward significance, while the Performance Change (permutation test,  $p = 0.03125$ ) demonstrates that the observed performance gains are not only substantial in magnitude but also consistent enough to reach statistical significance. These results underscore the potential of instruction-tuning in enhancing model behavior.

We then present the relationship between model performance and layer activation correlations in Figure 7, where higher performance scores show weak positive correlation with activation values across 6.7B to 9B parameter models (Pearson correlation:  $r = 0.601^*$ ,  $p = 0.030$ ). Notable examples include Llama-3-1-8B-Instruct (6.7B) and gemma-2-9b (9B), with activation patterns demonstrating parameter-scale dependent clustering (pink-to-purple gradient). The white grid and light blue confidence intervals facilitate model comparison.




Figure 6: Comparison of instruction-tuned versus base models across performance and activation correlations

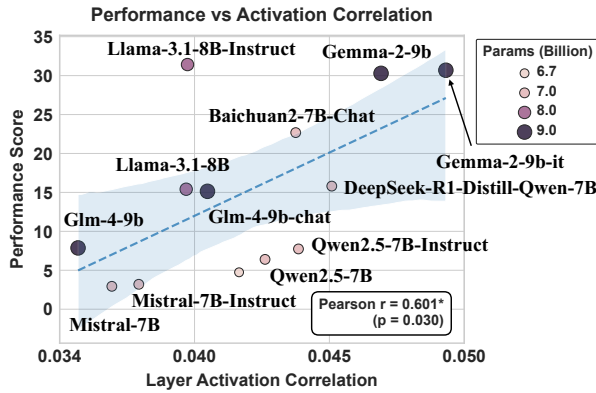


Figure 7: Correlation between model performance and activation patterns across LLMs

Taken together, these results reveal a clear relationship between model performance and their neural activation correlations across brain regions, with intermediate layers showing peak correspondence. Instruction tuning not only boosts performance but also strengthens alignment with brain activity, suggesting that model optimization can bring artificial and biological language processing closer together.

### Left-Right Hemispheric Asymmetry

In the third experiment, we investigate the lateralization patterns of neural activity associated with LLM processing across key language-related brain regions, focusing on the asymmetry between left and right hemispheres and its relationship to model performance. In Figure 8(a), we visualize the localization of ROIs in both hemispheres.

Then, we observed left-right hemispheric asymmetry in neural activity correlations between specific brain ROIs and LLMs. In Figure 8(b), the inferior frontal gyrus (IFG) and posterior temporal (PostTemp) regions showed stronger left-hemispheric dominance in LLM-related neural correlations, consistent with their roles in core language functions like production and comprehension (Hu et al. 2023). Conversely, the middle frontal gyrus (MFG) and anterior temporal (AntTemp) regions exhibited stronger right-hemispheric involvement, possibly reflecting specialization in tasks such as metaphor processing, contextual integration, and cross-modal semantics. Right MFG correlations may relate to LLMs’ demands for cognitive control in complex narra-

tives (Japee et al. 2015), while right AntTemp activity likely supports multimodal semantic representations (One-Sample t-test: IFG:  $p = 0.025$ ; PostTemp:  $p = 0.007$ ; MFG:  $p = 0.005$ ; AntTemp:  $p = 0.001$ ). These findings align with the classical language lateralization hypothesis of left-hemisphere dominance for syntax/semantics and emerging evidence of right-hemisphere contributions to cognitive control (Vigneau et al. 2006; Menon and D’Esposito 2022).

To examine the relationship between hemispheric asymmetry and model performance, Figure 8(c) shows correlations between LH-RH differences and performance metrics. Among the six ROIs analyzed, IFG and MFG lateralization were potentially linked to performance, showing positive relationships (Pearson correlation: IFG:  $r = 0.54$ ,  $p = 0.055$ ; MFG:  $r = 0.50$ ,  $p = 0.084$ ). No significant correlations were found for other regions (Pearson correlation: AntTemp:  $r = 0.02$ ,  $p = 0.941$ ; PostTemp:  $r = 0.09$ ,  $p = 0.770$ ; AngG:  $r = 0.09$ ,  $p = 0.687$ ; IFGorb:  $r = -0.06$ ,  $p = 0.845$ ). The regulatory role of the prefrontal cortex over distributed brain regions may underlie this phenomenon (Badre and Nee 2018), where its lateralized functional specialization could enhance top-down coordination of cognitive resources, mirroring the efficiency optimization observed in computational models.

This experiment confirms that hemispheric lateralization plays a significant role in the neural representation of language-related processes as captured by LLMs. The observed left-hemisphere dominance in core language areas and right-hemisphere contributions to higher-level cognitive functions align well with existing neurocognitive theories.

## Discussion and Conclusion

Understanding how LLMs align with human brain activity during sentence processing offers important insights into the parallels and divergences between artificial and biological language systems (Mahowald et al. 2024; Zhou et al. 2024). Our study contributes to this growing body of work through three core findings: (1) intermediate-layer alignment with neural activity, (2) the effects of instruction tuning and semantic comprehension on brain alignment, and (3) hemispheric asymmetries in neural–LLM correspondence. Below, we discuss each in turn.

### Intermediate layers better align with brain activity.

Across all evaluated models, we found that intermediate layers exhibit higher brain–model correlation than final layers, particularly across language-selective brain regions (e.g., IFG, PostTemp). This aligns with previous studies using intracranial EEG and MEG (Mischler et al. 2024; Zhou et al. 2024), and confirms that middle-layer representations in LLMs carry rich linguistic information that most closely mirrors human sentence-level processing.

Our use of fMRI complements these prior modalities by providing broader spatial coverage across regions and hemispheres. These results emphasize the importance of considering layer-wise architecture when interpreting the cognitive fidelity of LLMs, reinforcing the notion that later layers, often optimized for downstream tasks, may abstract away from biologically grounded representations.

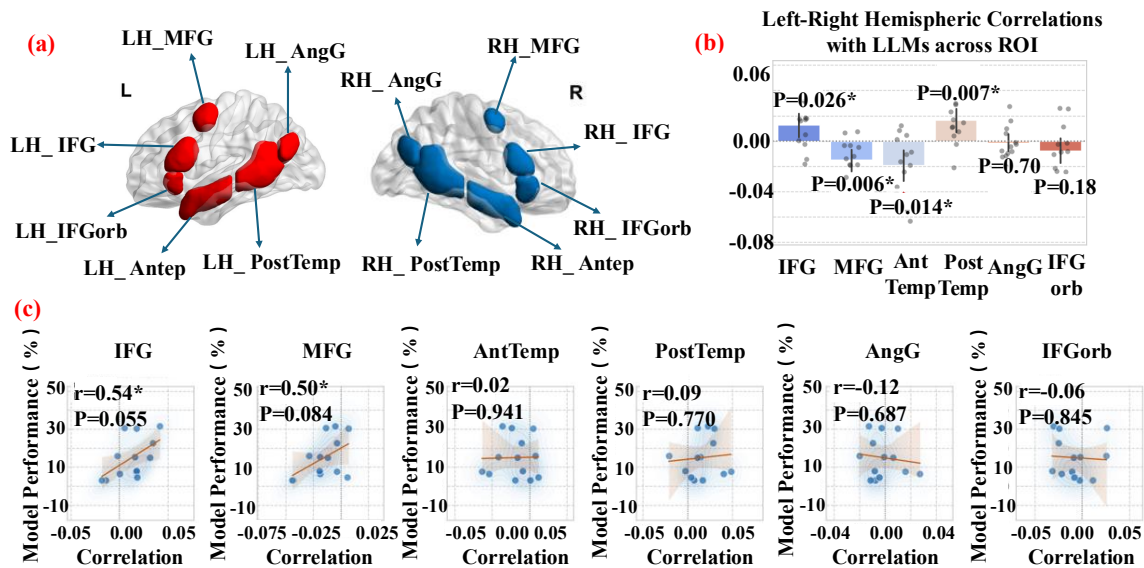


Figure 8: (a) illustrates the localization of ROIs in both hemispheres; (b) displays the left-minus-right (LH-RH) correlation differences in ROI-LLM associations; (c) examines the relationship between LH-RH asymmetry and model performance.

**Instruction tuning enhances brain alignment and comprehension.** We compared instruction-tuned models with their base versions across both semantic performance and brain activation correlation. Instruction-tuned variants consistently outperformed the base models, with improvements observed in both CSAA metric and neural alignment measures (Figure 6). These gains were statistically significant in performance and showed marginal significance in brain correlation, echoing trends reported in Ren et al. (2025).

Importantly, we shift the interpretive focus away from model size, a dominant narrative in prior work (Bonnasse-Gahot and Pallier 2024), toward semantic comprehension ability as a more meaningful predictor of neural alignment. Our results reveal a scaling law: within the 6.7B–9B parameter range, models with stronger comprehension capabilities (higher CSAA scores) also exhibit significantly higher similarity to human brain activation patterns (Pearson  $r = 0.601$ ,  $p = 0.030$ ; Figure 7).

This suggests that optimization techniques that improve semantic understanding (such as instruction tuning) not only enhance task performance but also produce representational structures more consistent with human neural activity. Future model evaluation efforts may benefit from incorporating neural alignment as an additional axis of assessment.

**Hemispheric asymmetry reflects functional specialization** Our third finding addresses the hemispheric organization of LLM–brain alignment. Consistent with classical neurocognitive theories, we observed left-hemispheric dominance in the inferior frontal gyrus (IFG) and posterior temporal (PostTemp) regions, areas central to syntactic processing and semantic integration (Hu et al. 2023; Vigneau et al. 2006). Conversely, the right hemisphere showed stronger correlations in the middle frontal gyrus (MFG) and anterior temporal (AntTemp), suggesting involvement in metaphor,

cross-modal semantics, and cognitive control (Japee et al. 2015; Menon and D’Esposito 2022).

These asymmetries were not only statistically robust (e.g., IFG:  $p = 0.025$ , AntTemp:  $p = 0.001$ ; Figure 8b), but also showed positive trends with model performance in specific ROIs. Notably, greater left–right differences in IFG and MFG were associated with higher CSAA scores (Pearson  $r = 0.54$  and  $0.50$ , respectively). This suggests that lateralized neural specialization, particularly in prefrontal regions, may support more efficient language representations, paralleling the efficiency gains observed in well-optimized LLMs (Badre and Nee 2018).

Interestingly, other regions (e.g., angular gyrus, IFGorb) showed no significant lateralization–performance correlation, possibly reflecting their roles in bilateral, domain-general processes such as semantic integration or the default mode network (Kuhnke et al. 2023).

**Implications and Future Directions.** Together, our findings paint a nuanced picture of how current LLMs reflect, and diverge from, human language systems. Instruction tuning and semantic comprehension capacity, not merely parameter size, are key drivers of neural similarity. Moreover, the correspondence is anatomically and functionally specific: intermediate-layer embeddings align best with neural signals, and this alignment respects known hemispheric specializations in the brain.

These insights raise exciting opportunities for future research. For instance, should training procedures explicitly target alignment with brain data to produce more cognitively plausible models? Could hybrid neuro-symbolic LLMs better reflect distributed and lateralized processing patterns observed in the brain? Our results lay the groundwork for these explorations by bridging computational performance with neurobiological plausibility in future work.

## References

- Anderson, A. J.; Kiela, D.; Binder, J. R.; Fernandino, L.; Humphries, C. J.; Conant, L. L.; Raizada, R. D. S.; Grimm, S.; and Lalor, E. C. 2021. Deep Artificial Neural Networks Reveal a Distributed Cortical Network Encoding Propositional Sentence-Level Meaning. *Journal of Neuroscience*, 41(18): 4100–4119.
- Antonello, R.; and Huth, A. 2024. Predictive Coding or Just Feature Discovery? An Alternative Account of Why Language Models Fit Brain Data. *Neurobiology of Language*, 5(1): 64–79.
- Antonello, R.; Vaidya, A.; and Huth, A. 2023. Scaling laws for language encoding models in fMRI. In Oh, A.; Naumann, T.; Globerson, A.; Saenko, K.; Hardt, M.; and Levine, S., eds., *Advances in Neural Information Processing Systems*, volume 36, 21895–21907. Curran Associates, Inc.
- Badre, D.; and Nee, D. E. 2018. Frontal cortex and the hierarchical control of behavior. *Trends in cognitive sciences*, 22(2): 170–188.
- Bonnasse-Gahot, L.; and Pallier, C. 2024. fMRI predictors based on language models of increasing complexity recover brain left lateralization. *arXiv preprint arXiv:2405.17992*.
- Brennan, J. 2016. Naturalistic sentence comprehension in the brain. *Language and Linguistics Compass*, 10(7): 299–313.
- Caucheteux, C.; Gramfort, A.; and King, J.-R. 2021. Disentangling syntax and semantics in the brain with deep networks. In Meila, M.; and Zhang, T., eds., *Proceedings of the 38th International Conference on Machine Learning*, volume 139 of *Proceedings of Machine Learning Research*, 1336–1348. PMLR.
- Caucheteux, C.; Gramfort, A.; and King, J.-R. 2023. Evidence of a predictive coding hierarchy in the human brain listening to speech. *Nature Human Behaviour*, 7(3): 430–441.
- Caucheteux, C.; and King, J.-R. 2022. Brains and algorithms partially converge in natural language processing. *Communications Biology*, 5(1): 134. © 2022. The Author(s). Research Support, Non-U.S. Gov't.
- Doerig, A.; Kietzmann, T. C.; Allen, E.; Wu, Y.; Naselaris, T.; Kay, K.; and Charest, I. 2025. High-level visual representations in the human brain are aligned with large language models. *Nature Machine Intelligence*, 7(8): 1220–1234.
- Ethayarajh, K. 2019. How Contextual are Contextualized Word Representations? Comparing the Geometry of BERT, ELMo, and GPT-2 Embeddings. In *Proceedings of the 2019 EMNLP-IJCNLP*, 55–65. Hong Kong, China: Association for Computational Linguistics.
- Fedorenko, E.; Hsieh, P.-J.; Nieto-Castañón, A.; Whitfield-Gabrieli, S.; and Kanwisher, N. 2010. New Method for fMRI Investigations of Language: Defining ROIs Functionally in Individual Subjects. *Journal of Neurophysiology*, 104(2): 1177–1194. PMID: 20410363.
- Friederici, A. D. 2011. The brain basis of language processing: from structure to function. *Physiological reviews*, 91(4): 1357–1392.
- Fu, T.; Xu, X.; Xu, W.; Chen, J.; Ren, R.; Deng, B.; Zhao, X.; Cao, J.; and Cao, X. 2025. Two Heads are Better than One: Distilling Large Language Model Features Into Small Models with Feature Decomposition and Mixture. arXiv:2511.07110.
- Geschwind, N. 1967. Wernicke's Contribution to the Study of Aphasia. *Cortex*, 3(4): 449–463.
- Goldstein, A.; Zada, Z.; Buchnik, E.; Schain, M.; Price, A.; Aubrey, B.; Nastase, S. A.; Feder, A.; Emanuel, D.; Cohen, A.; Jansen, A.; Gazula, H.; Choe, G.; Rao, A.; Kim, C.; Casto, C.; Fanda, L.; Doyle, W.; Friedman, D.; Dugan, P.; Melloni, L.; Reichart, R.; Devore, S.; Flinker, A.; Hasenfratz, L.; Levy, O.; Hassidim, A.; Brenner, M.; Matias, Y.; Norman, K. A.; Devinsky, O.; and Hasson, U. 2022. Shared computational principles for language processing in humans and deep language models. *Nature Neuroscience*, 25(3): 369–380.
- Hagoort, P. 2016. MUC (Memory, Unification, Control): A model on the neurobiology of language beyond single word processing. In *Neurobiology of language*, 339–347. Elsevier.
- Hickok, G.; and Poeppel, D. 2004. Dorsal and ventral streams: a framework for understanding aspects of the functional anatomy of language. *Cognition*, 92(1-2): 67–99.
- Hickok, G.; and Poeppel, D. 2007. The cortical organization of speech processing. *Nature Reviews Neuroscience*, 8(5): 393–402.
- Hosseini, E. A.; Schrimpf, M.; Zhang, Y.; Bowman, S.; Zaslavsky, N.; and Fedorenko, E. 2024. Artificial Neural Network Language Models Predict Human Brain Responses to Language Even After a Developmentally Realistic Amount of Training. *Neurobiology of Language*, 5(1): 43–63.
- Hu, J.; Small, H.; Kean, H.; Takahashi, A.; Zekelman, L.; Kleinman, D.; Ryan, E.; Nieto-Castañón, A.; Ferreira, V.; and Fedorenko, E. 2023. Precision fMRI reveals that the language-selective network supports both phrase-structure building and lexical access during language production. *Cerebral Cortex*, 33(8): 4384–4404.
- Humphries, C.; Binder, J. R.; Medler, D. A.; and Liebenthal, E. 2007. Time course of semantic processes during sentence comprehension: an fMRI study. *Neuroimage*, 36(3): 924–932.
- Japee, S.; Holiday, K.; Satyshur, M. D.; Mukai, I.; and Ungerleider, L. G. 2015. A role of right middle frontal gyrus in reorienting of attention: a case study. *Frontiers in systems neuroscience*, 9: 23.
- Kriegeskorte, N.; Mur, M.; and Bandettini, P. A. 2008. Representational similarity analysis-connecting the branches of systems neuroscience. *Frontiers in systems neuroscience*, 2: 249.
- Kuhnke, P.; Chapman, C. A.; Cheung, V. K.; Turker, S.; Graessner, A.; Martin, S.; Williams, K. A.; and Hartwigsen, G. 2023. The role of the angular gyrus in semantic cognition: a synthesis of five functional neuroimaging studies. *Brain Structure and Function*, 228(1): 273–291.
- Kutas, M.; and Hillyard, S. A. 1984. Brain potentials during reading reflect word expectancy and semantic association. *Nature*, 307(5947): 161–163.

- Lei, Y.; Liu, H.; Xie, C.; Liu, S.; Yin, Z.; Chen, C.; Li, G.; Torr, P.; and Wu, Z. 2024. Fairmindsim: Alignment of behavior, emotion, and belief in humans and llm agents amid ethical dilemmas. *arXiv preprint arXiv:2410.10398*.
- Lerner, Y.; Honey, C. J.; Silbert, L. J.; and Hasson, U. 2011. Topographic Mapping of a Hierarchy of Temporal Receptive Windows Using a Narrated Story. *Journal of Neuroscience*, 31(8): 2906–2915.
- Li, J.; Bhattasali, S.; Zhang, S.; Franzluebbers, B.; Luh, W.-M.; Spreng, R. N.; Brennan, J. R.; Yang, Y.; Pallier, C.; and Hale, J. 2022. Le Petit Prince multilingual naturalistic fMRI corpus. *Scientific Data*, 9(1): 530.
- Liu, H.; Dai, Y.; Tan, H.; Lei, Y.; Zhou, Y.; and Wu, Z. 2025. Outraged AI: Large language models prioritise emotion over cost in fairness enforcement. *arXiv preprint arXiv:2510.17880*.
- Liu, H.; Lei, Y.; and Wu, Z. 2025. Prosocial behavior in Large Language Models: Value alignment and affective mechanisms. *Science China Technological Sciences*, 68(8): 1820403.
- Luo, Y.; Xu, M.; and Xiong, D. 2022. CogTaskonomy: Cognitively Inspired Task Taxonomy Is Beneficial to Transfer Learning in NLP. In Muresan, S.; Nakov, P.; and Villavicencio, A., eds., *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, 904–920. Dublin, Ireland: Association for Computational Linguistics.
- Mahowald, K.; Ivanova, A. A.; Blank, I. A.; Kanwisher, N.; Tenenbaum, J. B.; and Fedorenko, E. 2024. Dissociating language and thought in large language models. *Trends in cognitive sciences*.
- Menon, V.; and D’Esposito, M. 2022. The role of PFC networks in cognitive control and executive function. *Neuropsychopharmacology*, 47(1): 90–103.
- Messi, A.-P.; and Pyllkanen, L. 2025. Tracking neural correlates of contextualized meanings with representational similarity analysis. *Journal of Neuroscience*, 45(19).
- Mischler, G.; Li, Y. A.; Bickel, S.; Mehta, A. D.; and Mesgarani, N. 2024. Contextual feature extraction hierarchies converge in large language models and the brain. *Nature Machine Intelligence*, 1–11.
- Mumford, J. A.; Davis, T.; and Poldrack, R. A. 2014. The impact of study design on pattern estimation for single-trial multivariate pattern analysis. *NeuroImage*, 103: 130–138.
- Osterhout, L.; and Holcomb, P. J. 1992. Event-related brain potentials elicited by syntactic anomaly. *Journal of memory and language*, 31(6): 785–806.
- Price, C. J. 2012. A review and synthesis of the first 20 years of PET and fMRI studies of heard speech, spoken language and reading. *Neuroimage*, 62(2): 816–847.
- Ren, Y.; Jin, R.; Zhang, T.; and Xiong, D. 2025. Do Large Language Models Mirror Cognitive Language Processing? In Rambow, O.; Wanner, L.; Apidianaki, M.; Al-Khalifa, H.; Eugenio, B. D.; and Schockaert, S., eds., *Proceedings of the 31st International Conference on Computational Linguistics*, 2988–3001. Abu Dhabi, UAE: Association for Computational Linguistics.
- Schrimpf, M.; Blank, I. A.; Tuckute, G.; Kauf, C.; Hosseini, E. A.; Kanwisher, N.; Tenenbaum, J. B.; and Fedorenko, E. 2021. The neural architecture of language: Integrative modeling converges on predictive processing. *Proceedings of the National Academy of Sciences*, 118(45): e2105646118.
- Sun, J.; Wang, S.; Zhang, J.; and Zong, C. 2021. Neural Encoding and Decoding With Distributed Sentence Representations. *IEEE Transactions on Neural Networks and Learning Systems*, 32(2): 589–603.
- Tenney, I.; Das, D.; and Pavlick, E. 2019. BERT Rediscovered the Classical NLP Pipeline. In Korhonen, A.; Traum, D.; and Màrquez, L., eds., *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, 4593–4601. Florence, Italy: Association for Computational Linguistics.
- Toneva, M.; and Wehbe, L. 2019. Interpreting and improving natural-language processing (in machines) with natural language-processing (in the brain). In Wallach, H.; Larochelle, H.; Beygelzimer, A.; d’Alché-Buc, F.; Fox, E.; and Garnett, R., eds., *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc.
- Tuckute, G.; Sathe, A.; Srikant, S.; Taliaferro, M.; Wang, M.; Schrimpf, M.; Kay, K.; and Fedorenko, E. 2024. Driving and suppressing the human language network using large language models. *Nature Human Behaviour*, 8(3): 544–561.
- Vigneau, M.; Beaucoisin, V.; Hervé, P.-Y.; Duffau, H.; Crivello, F.; Houde, O.; Mazoyer, B.; and Tzourio-Mazoyer, N. 2006. Meta-analyzing left hemisphere language areas: phonology, semantics, and sentence processing. *Neuroimage*, 30(4): 1414–1432.
- Wang, X.; Hu, C.; Ma, B.; Rottger, P.; and Plank, B. 2024. Look at the Text: Instruction-Tuned Language Models are More Robust Multiple Choice Selectors than You Think. In *First Conference on Language Modeling*.
- Yu, S.; Gu, C.; Huang, K.; and Li, P. 2024. Predicting the next sentence (not word) in large language models: What model-brain alignment tells us about discourse comprehension. *Science advances*, 10(21): eadn7744.
- Zacks, J. M.; Mar, R. A.; and Calarco, N. 2017. The cognitive neuroscience of discourse: Covered ground and new directions. In *The Routledge handbook of discourse processes*, 269–294. Routledge.
- Zhao, X.; Zhao, S.; Yu, H.; Zhang, L.; and Li, Q. 2025. AgentCDM: Enhancing Multi-Agent Collaborative Decision-Making via ACH-Inspired Structured Reasoning. *arXiv preprint arXiv:2508.11995*.
- Zheng, C.; Zhou, H.; Meng, F.; Zhou, J.; and Huang, M. 2024. Large Language Models Are Not Robust Multiple Choice Selectors. In *The Twelfth International Conference on Learning Representations*.
- Zhou, Y.; Liu, E.; Neubig, G.; Tarr, M. J.; and Wehbe, L. 2024. Divergences between Language Models and Human Brains. In Globerson, A.; Mackey, L.; Belgrave, D.; Fan, A.; Paquet, U.; Tomczak, J.; and Zhang, C., eds., *Advances in Neural Information Processing Systems*, volume 37, 137999–138031. Curran Associates, Inc.