

Far from the Shallow: Brain-Predictive Reasoning Embedding through Residual Disentanglement¹

Linyang He*

Tianjun Zhong*

Richard Antonello

Gavin Mischler

Micah Goldblum

Nima Mesgarani

Columbia University

LINYANG.HE@COLUMBIA.EDU

TIANJUN.ZHONG@COLUMBIA.EDU

RJA2163@COLUMBIA.EDU

GM2944@COLUMBIA.EDU

MICAH.G@COLUMBIA.EDU

NIMA@EE.COLUMBIA.EDU

Editors: List of editors' names

Abstract

Understanding how the human brain progresses from processing simple linguistic inputs to performing high-level reasoning is a fundamental challenge in neuroscience. While modern large language models (LLMs) are increasingly used to model neural responses to language, their internal representations are highly "entangled," mixing information about lexicon, syntax, meaning, and reasoning. This entanglement biases conventional brain encoding analyses toward linguistically shallow features (e.g., lexicon and syntax), making it difficult to isolate the neural substrates of cognitively deeper processes. Here, we introduce a residual disentanglement method that computationally isolates these components. By first probing an LM to identify feature-specific layers, our method iteratively regresses out lower-level representations to produce four nearly orthogonal embeddings for lexicon, syntax, meaning, and, critically, reasoning. We used these disentangled embeddings to model intracranial (ECoG) brain recordings from neurosurgical patients listening to natural speech. We show that: 1) This isolated reasoning embedding exhibits unique predictive power, accounting for variance in neural activity not explained by other linguistic features and even extending to the recruitment of visual regions beyond classical language areas. 2) The neural signature for reasoning is temporally distinct, peaking later (~350-400ms) than signals related to lexicon, syntax, and meaning, consistent with its position atop a processing hierarchy. 3) Standard, non-disentangled LLM embeddings can be misleading, as their predictive success is primarily attributable to linguistically shallow features, masking the more subtle contributions of deeper cognitive processing. Our work provides compelling neural evidence for an abstract reasoning computation during language comprehension and offers a robust framework for mapping distinct cognitive functions from artificial models to the human brain².

Keywords: Disentangled Representations, Encoding Models, Reasoning, Language

* Equal Contribution.

1. This work is also accepted at NeurIPS 2025 (Main Track). This shorter version is dedicated to discussions at the NeurRepS Workshop. For the full conference paper, please refer to [arXiv:2510.22860](https://arxiv.org/abs/2510.22860).

2. Code available [here](https://github.com/HeLinyang/ReasoningEmbedding).

1. Introduction

A growing body of work has shown that LLMs exhibit strong representational alignment with human brain activity during language comprehension. These studies employ *language encoding models*, powerful predictors of measured brain activity in response to a language stimulus (Wehbe et al., 2014; Huth et al., 2016; de Heer et al., 2017; Jain and Huth, 2018; Toneva and Wehbe, 2019; Brennan et al., 2020; Goldstein et al., 2022; Oota et al., 2022; Heilbron et al., 2022; Chen et al., 2023; Li et al., 2023; Antonello et al., 2023; Aw and Toneva, 2023; Chen et al., 2024a,b). Typically, these models are created by inputting a language stimulus into an LLM, to create a contextual embedding of the stimulus using the models’ internal hidden state. A linear mapping is then learned between these hidden states and the measured brain response to that stimulus. Modern language encoding models represent some of the most advanced computational approaches for studying neural processing across sensory modalities, and they provide a unique window into the semantic processing landscape of human cortex.

However, most work that has studied this brain-LLM alignment has focused on semantics and low-level phonological relationships (Antonello et al., 2021; Vaidya et al., 2022; Mischler et al., 2024): little is known about the degree to which LLMs and the brain align in their higher-level reasoning processes. One reason is historical: reasoning capabilities have only recently emerged robustly in modern LLMs (Ke et al., 2025). Another issue is methodological. Most prior studies have largely treated language encoding models as monolithic black boxes that map entire hidden states to brain responses without disentangling the specific linguistic or cognitive functions represented within them, such as lexical features, syntactic structure, semantic meaning, or reasoning processes.

Recent findings (Lampinen et al., 2024) show that when models are trained to perform multiple tasks of varying complexity, their internal representations disproportionately favor simpler and more linearly extractable features, even when accuracy on complex tasks is equally high. This suggests that in brain alignment studies, unsegmented LLM features may reflect a bias toward lexical or syntactic components simply because they are computationally easier to represent.

To move beyond these biases and probe the deeper relationship between machine and human reasoning, we introduce a novel approach: residual reasoning embeddings. Rather than comparing raw LLM activations to brain activations, we isolate the reasoning-specific component of LLM hidden states by separating them from discrete lexical, syntactic, and meaning features. This residual embedding captures the unique portion of variance in the initial embedding that is attributable to high-level reasoning information.

We then use these residual reasoning embeddings to build encoding models that predict human electrocorticographic (ECOG) brain activity during tasks that require inference. We demonstrate that these reasoning representations align with distinct spatiotemporal neural patterns, peaking later in time and engaging more frontal and visual regions than lexical or syntactic signals, suggesting a shared computational hierarchy between LLMs and the human brain. Furthermore, we show that full LLM embeddings are biased toward lower-level features, and that only through disentanglement can reasoning-related brain activity be meaningfully isolated. Together, our results offer a new lens on both model interpretability and the neural basis of abstract linguistic reasoning.

2. Methods

Our main experiments use the Qwen2.5-14B model. We apply minimal pair probing (He et al., 2024a,b) to identify the LLM layers most specialized for syntactic, meaning, and reasoning features. For each probing dataset, we extract sentence representations from each layer by feeding the sentence in isolation and taking the hidden state of its final token. We then use these sentence representations to train a logistic regression classifier to predict the correct item in each minimal pair. We define the saturation layer as the earliest layer where performance plateaus. Building on the saturation layers identified via minimal pair probing, we train ridge regression models on an expanded podcast corpus to isolate feature-specific residuals. We then apply these models to the ECoG-aligned transcript to extract feature-specific embeddings, which are used to predict neural responses. See Appendix A for more details.

3. Probing Results

Applying our probing pipeline to Qwen2.5-14B, we observe a clear progression in the emergence of representational features. Syntax saturates earliest at layer 6, followed by meaning at layer 20, and reasoning only at the deeper layer 30. This ordering highlights that low-level linguistic structure is captured quickly, whereas high-level reasoning requires substantially more depth.

To verify that this pattern is not specific to a single model, we extended the probing pipeline across multiple Qwen families spanning different sizes and generations. The same emergence order was consistently observed, with only one exception in Qwen-1.8B, where syntax and meaning saturated at the same layer. Full cross-model analyses are reported in Appendix C.

4. Disentanglement Validation

Matrix-Level Orthogonality of Residual Embeddings We justify the approximate orthogonality of the feature-specific embeddings at the level of their embedding matrices: the n -dimensional feature columns of one residual embedding matrix are approximately uncorrelated with those of another (Appendix E.1). This sample-axis orthogonality ensures that residual embedding matrices contribute non-overlapping signals across the dataset, which is precisely the property required for the subsequent brain encoding analysis.

Token-Level Independence via Cosine Similarity We also test independence at the level of individual tokens (Appendix E.2). Figure 2a compares the pairwise mean absolute cosine similarity among raw hidden states at the four saturation layers (top) and among the corresponding residual embeddings (bottom). While the hidden states show substantial overlap, especially between meaning and reasoning layers ($\bar{C} = 0.751$), the residual embeddings exhibit near-zero off-diagonal similarity (all ≤ 0.045) across all pairs.

Feature Probing on Residual Embeddings Each residual should also preserve information relevant to its intended linguistic feature. We evaluate this by reapplying the same probing tasks used to define feature emergence. Results are shown in Figure 2b. Each residual achieves the highest performance on its own task while performing worse on unrelated

tasks. These findings demonstrate that the residual embeddings are meaningfully disentangled and capture feature-specific information rather than reflecting general task complexity. See Appendix E.3 for more details.

5. Encoding Results

Shallow features dominate neural prediction, but lexical activations are sparse. Encoding models based on shallow linguistic features, such as lexical identity and syntax, showed the strongest neural correlations (Figure 3). Electrodes responsive to these features exhibited higher peak correlations than those for meaning or reasoning ($p < 0.001$). Despite strong correlations, the lexicon embedding activated the fewest electrodes, suggesting specialized cortical sites. Deeper features engaged broader but weaker populations. Thus, shallow features dominate localized alignment, while deeper ones distribute across cortex. Without disentanglement, lexical and syntactic peaks mask slower reasoning dynamics, underscoring the need for residual separation. See Appendix G.1 for more details.

Reasoning embedding shows different temporal pattern compared to Lexicon, Syntax and Meaning. The temporal correlation profiles of Figure 4 illustrate a clear processing hierarchy among linguistic representations. Syntactic signals rise early and peak slightly before onset, suggesting syntactically driven prediction. Lexical features peak immediately after onset, reflecting rapid mapping between sound and word identity. Meaning representations peak later, indicating semantic integration. In contrast, reasoning signals rise only after onset and peak around 362 ms, reflecting delayed high-level computation beyond linguistic parsing. Together, these patterns reveal a progression from syntax to semantics to reasoning in the human cortex. See Appendix G.2 for more details.

Reasoning recruits more than language area compared to low-level aspects. Shallow features like lexicon show high correlations but remain confined to classical language regions such as IFG and STG. In contrast, syntax, meaning, and especially reasoning engage broader cortical areas. Reasoning-selective electrodes extend into SFG and visual cortex with stronger activation ($p < 0.001$; Figure 6). This expanded recruitment suggests reasoning involves higher-order, multimodal cognition beyond linguistic processing. See Appendix G.3 for more details.

6. Discussion

Using residual embeddings, we disentangled the contributions of lexicon, syntax, meaning, and reasoning to neural encoding of speech. This hierarchical framework localizes feature-emergent layers and isolates higher-order reasoning representations. Lower-level features show stronger, earlier, and more localized activations, while reasoning appears later and extends beyond language regions into SFG and visual cortex. These findings reveal a spatiotemporal hierarchy from syntax to reasoning and show that reasoning engages broader, multimodal cortical systems. Our framework provides a principled way to separate linguistic and cognitive components in LLMs and align them with the human brain. See Appendix I for more details.

Acknowledgments

This work is funded by the National Institutes of Health (NIH- NIDCD) and a grant from Marie-Josée and Henry R. Kravis.

References

Richard Antonello, Javier S Turek, Vy Vo, and Alexander Huth. Low-dimensional structure in the space of language representations is reflected in brain responses. *Advances in Neural Information Processing Systems*, 34, 2021.

Richard Antonello, Aditya Vaidya, and Alexander G. Huth. Scaling laws for language encoding models in fmri, 2023.

Khai Loong Aw and Mariya Toneva. Training language models to summarize narratives improves brain alignment, 2023.

Jinze Bai, Shuai Bai, Yunfei Chu, Zeyu Cui, Kai Dang, Xiaodong Deng, Yang Fan, Wenbin Ge, Yu Han, Fei Huang, et al. Qwen technical report. *arXiv preprint arXiv:2309.16609*, 2023.

Jonathan R Brennan, Chris Dyer, Adhiguna Kuncoro, and John T Hale. Localizing syntactic predictions using recurrent neural network grammars. *Neuropsychologia*, 146:107479, 2020.

Charlotte Caucheteux and Jean-Rémi King. Brains and algorithms partially converge in natural language processing. *Communications biology*, 5(1):134, 2022.

Catherine Chen, Tom Dupré la Tour, Jack Gallant, Daniel Klein, and Fatma Deniz. The cortical representation of language timescales is shared between reading and listening. *bioRxiv*, pages 2023–01, 2023.

Peili Chen, Linyang He, Li Fu, Lu Fan, Edward F Chang, and Yuanning Li. Do self-supervised speech and language models extract similar representations as human brain? In *ICASSP 2024-2024 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 2225–2229. IEEE, 2024a.

Peili Chen, Shiji Xiang, Linyang He, Edward F Chang, and Yuanning Li. Convergent representations and spatiotemporal dynamics of speech and language in brain and deep neural networks. *bioRxiv*, pages 2024–12, 2024b.

Wendy A de Heer, Alexander G Huth, Thomas L Griffiths, Jack L Gallant, and Frédéric E Theunissen. The hierarchical cortical organization of human speech processing. *Journal of Neuroscience*, 37(27):6539–6557, 2017.

Richard Futrell, Ethan Wilcox, Takashi Morita, Peng Qian, Miguel Ballesteros, and Roger Levy. Neural language models as psycholinguistic subjects: Representations of syntactic state. *arXiv preprint arXiv:1903.03260*, 2019.

Giorgio Ganis, William L. Thompson, and Stephen M. Kosslyn. Brain areas underlying visual mental imagery and perception: An fMRI study. *Cognitive Brain Research*, 20(2):226–241, 2004. doi: 10.1016/j.cogbrainres.2004.02.012.

Ariel Goldstein, Zaid Zada, Eliav Buchnik, Mariano Schain, Amy Price, Bobbi Aubrey, Samuel A Nastase, Amir Feder, Dotan Emanuel, Alon Cohen, et al. Shared computational principles for language processing in humans and deep language models. *Nature neuroscience*, 25(3):369–380, 2022.

Linyang He, Peili Chen, Ercong Nie, Yuanning Li, and Jonathan R Brennan. Decoding probing: Revealing internal linguistic structures in neural language models using minimal pairs. In *Proceedings of the 2024 Joint International Conference on Computational Linguistics, Language Resources and Evaluation (LREC-COLING 2024)*, pages 4488–4497, 2024a.

Linyang He, Ercong Nie, Helmut Schmid, Hinrich Schütze, Nima Mesgarani, and Jonathan Brennan. Large language models as neurolinguistic subjects: Identifying internal representations for form and meaning. *arXiv preprint arXiv:2411.07533*, 2024b.

Linyang He, Qiaolin Wang, Xilin Jiang, and Nima Mesgarani. Layer-wise minimal pair probing reveals contextual grammatical-conceptual hierarchy in speech representations. *arXiv preprint arXiv:2509.15655*, 2025.

Micha Heilbron, Kristijan Armeni, Jan-Mathijs Schoffelen, Peter Hagoort, and Floris P De Lange. A hierarchy of linguistic predictions during natural language comprehension. *Proceedings of the National Academy of Sciences*, 119(32):e2201968119, 2022.

John Hewitt and Christopher D Manning. A structural probe for finding syntax in word representations. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4129–4138, 2019.

Alexander G Huth, Wendy A De Heer, Thomas L Griffiths, Frédéric E Theunissen, and Jack L Gallant. Natural speech reveals the semantic maps that tile human cerebral cortex. *Nature*, 532(7600):453–458, 2016.

Shailee Jain and Alexander Huth. Incorporating context into language encoding models for fmri. In S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 31. Curran Associates, Inc., 2018. URL <https://proceedings.neurips.cc/paper/2018/file/f471223d1a1614b58a7dc45c9d01df19-Paper.pdf>.

Zixuan Ke, Fangkai Jiao, Yifei Ming, Xuan-Phi Nguyen, Austin Xu, Do Xuan Long, Minzhi Li, Chengwei Qin, Peifeng Wang, Silvio Savarese, et al. A survey of frontiers in llm reasoning: Inference scaling, learning to reason, and agentic systems. *arXiv preprint arXiv:2504.09037*, 2025.

Sakaguchi Keisuke, Le Bras Ronan, Bhagavatula Chandra, and Choi Yejin. Winogrande: An adversarial winograd schema challenge at scale. *arXiv preprint arXiv:1907.10641*, 2019.

Menoua Keshishian, Serdar Akkol, Jose Herrero, Stephan Bickel, Ashesh D Mehta, and Nima Mesgarani. Joint, distributed and hierarchically organized encoding of linguistic features in the human auditory cortex. *Nature human behaviour*, 7(5):740–753, 2023.

Menoua Keshishian, Gavin Mischler, Samuel Thomas, Brian Kingsbury, Stephan Bickel, Ashesh D Mehta, and Nima Mesgarani. Parallel hierarchical encoding of linguistic representations in the human auditory cortex and recurrent automatic speech recognition systems. *bioRxiv*, 2025.

Stephen M. Kosslyn, William L. Thompson, and Nathaniel M. Alpert. Neural systems shared by visual imagery and visual perception: A PET study. *NeuroImage*, 6(4):320–334, 1997. doi: 10.1006/nimg.1997.0295.

Andrew Kyle Lampinen, Stephanie C. Y. Chan, and Katherine Hermann. Learned feature representations are biased by complexity, learning order, position, and more. *Transactions on Machine Learning Research*, September 2024. URL <https://openreview.net/forum?id=aY2nsgE97a>.

Yuanning Li, Gopala K Anumanchipalli, Abdelrahman Mohamed, Peili Chen, Laurel H Carney, Junfeng Lu, Jinsong Wu, and Edward F Chang. Dissecting neural computations in the human auditory pathway using deep neural networks for speech. *Nature Neuroscience*, 26(12):2213–2225, 2023.

Tal Linzen, Emmanuel Dupoux, and Yoav Goldberg. Assessing the ability of lstms to learn syntax-sensitive dependencies. *Transactions of the Association for Computational Linguistics*, 4:521–535, 2016.

Christopher D Manning, Kevin Clark, John Hewitt, Urvashi Khandelwal, and Omer Levy. Emergent linguistic structure in artificial neural networks trained by self-supervision. *Proceedings of the National Academy of Sciences*, 117(48):30046–30054, 2020.

Huanru Henry Mao, Shuyang Li, Julian McAuley, and Garrison Cottrell. Speech recognition and multi-speaker diarization of long conversations. *arXiv preprint arXiv:2005.08072*, 2020.

Rebecca Marvin and Tal Linzen. Targeted syntactic evaluation of language models. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 1192–1202, Brussels, Belgium, October–November 2018. Association for Computational Linguistics. doi: 10.18653/v1/D18-1151. URL <https://aclanthology.org/D18-1151>.

Gavin Mischler, Yinghao Aaron Li, Stephan Bickel, Ashesh D Mehta, and Nima Mesgarani. Contextual feature extraction hierarchies converge in large language models and the brain. *Nature Machine Intelligence*, pages 1–11, 2024.

Kanishka Misra, Julia Rayz, and Allyson Ettinger. Comps: Conceptual minimal pair sentences for testing robust property knowledge and its inheritance in pre-trained language models. In *Proceedings of the 17th Conference of the European Chapter of the Association for Computational Linguistics*, pages 2920–2941, 2023.

Subba Reddy Oota, Manish Gupta, and Mariya Toneva. Joint processing of linguistic properties in brains and language models. *arXiv preprint arXiv:2212.08094*, 2022.

Abulhair Saparov and He He. Language models are greedy reasoners: A systematic formal analysis of chain-of-thought. *arXiv preprint arXiv:2210.01240*, 2023.

Qwen Team. Qwen2.5 technical report. *arXiv preprint arXiv:2412.15115*, 2025a. URL <https://arxiv.org/abs/2412.15115>.

Qwen Team. Qwen3 technical report, 2025b. URL <https://arxiv.org/abs/2505.09388>.

Mariya Toneva and Leila Wehbe. Interpreting and improving natural-language processing (in machines) with natural language-processing (in the brain). In H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc., 2019. URL <https://proceedings.neurips.cc/paper/2019/file/749a8e6c231831ef7756db230b4359c8-Paper.pdf>.

Aditya R Vaidya, Shailee Jain, and Alexander G Huth. Self-supervised models of audio effectively explain human cortical responses to speech. *arXiv preprint arXiv:2205.14252*, 2022.

Alex Warstadt, Alicia Parrish, Haokun Liu, Anhad Mohananey, Wei Peng, Sheng-Fu Wang, and Samuel R Bowman. Blimp: The benchmark of linguistic minimal pairs for english. *Transactions of the Association for Computational Linguistics*, 8:377–392, 2020.

Leila Wehbe, Brian Murphy, Partha Talukdar, Alona Fyshe, Aaditya Ramdas, and Tom Mitchell. Simultaneously uncovering the patterns of brain regions involved in different story reading subprocesses. *PloS one*, 9(11):e112575, 2014.

An Yang, Baosong Yang, Binyuan Hui, Bo Zheng, Bowen Yu, Chang Zhou, Chengpeng Li, Chengyuan Li, Dayiheng Liu, Fei Huang, Guanting Dong, Haoran Wei, Huan Lin, Jialong Tang, Jialin Wang, Jian Yang, Jianhong Tu, Jianwei Zhang, Jianxin Ma, Jianxin Yang, Jin Xu, Jingren Zhou, Jinze Bai, Jinzheng He, Junyang Lin, Kai Dang, Keming Lu, Keqin Chen, Kexin Yang, Mei Li, Mingfeng Xue, Na Ni, Pei Zhang, Peng Wang, Ru Peng, Rui Men, Ruize Gao, Runji Lin, Shijie Wang, Shuai Bai, Sinan Tan, Tianhang Zhu, Tianhao Li, Tianyu Liu, Wenbin Ge, Xiaodong Deng, Xiaohuan Zhou, Xingzhang Ren, Xinyu Zhang, Xipin Wei, Xuancheng Ren, Xuejing Liu, Yang Fan, Yang Yao, Yichang Zhang, Yu Wan, Yunfei Chu, Yuqiong Liu, Zeyu Cui, Zhenru Zhang, Zhifang Guo, and Zhihao Fan. Qwen2 technical report. *arXiv preprint arXiv:2407.10671*, 2024. URL <https://arxiv.org/abs/2407.10671>.

Zaid Zada, Samuel A Nastase, Bobbi Aubrey, Itamar Jalon, Sebastian Michelmann, Haocheng Wang, Liat Hasenfratz, Werner Doyle, Daniel Friedman, Patricia Dugan, Lucia Melloni, Sasha Devore, Adeen Flinker, Orrin Devinsky, Ariel Goldstein, and Uri Hasson. The “podcast” ecog dataset for modeling neural activity during natural language comprehension. *bioRxiv*, 2025. doi: 10.1101/2025.02.14.638352. URL <https://doi.org/10.1101/2025.02.14.638352>.

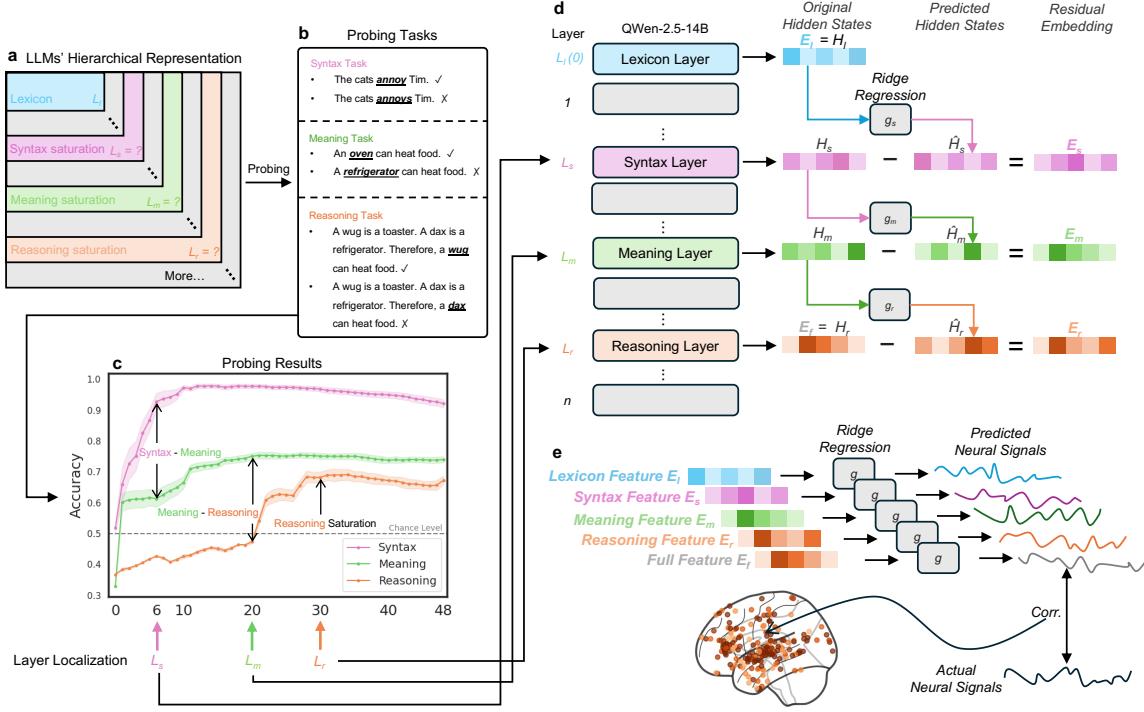


Figure 1: **a)** Hierarchical representations in an LLM. Transformer layers accumulate information in order: lexical features emerge first, followed by syntax, contextual meaning and eventually higher-order reasoning, with still-richer knowledge continuing in later layers. **b)** Minimal-pair probing tasks. Three diagnostic sentence sets separately test syntax, concept meaning and multi-premise reasoning. **c)** Layer localization from probing curves. We define L_s – the earliest layer where syntax performance saturates while meaning is still low; L_m – layer where meaning saturates but reasoning has not yet emerged; and L_r – layer where reasoning performance plateaus. These identified layers through probing will be used in later analyses. **d)** Feature disentangling across layers. Starting from the localized layers, we iteratively regress lower-level features out of higher ones. Details of residual disentangling constructions could be found in Algorithm 1. Residual disentangling yields four orthogonal embeddings that isolate lexicon, syntax, meaning and reasoning information. **e)** Brain encoding with purified features. Each residual feature is fed into a ridge encoder to predict high-gamma ECoG responses of Podcast listening. Comparing predicted and actual neural signals reveals the spatiotemporal distribution of cortical activity uniquely associated with lexicon, syntax, meaning and reasoning representations.

Appendix A. Detailed Methodology

A.1. Probing Datasets

To identify the LLM layers most specialized for encoding syntactic, semantic, and reasoning features, we apply feature-specific probing using two diagnostic datasets. For syntactic probing, we use the Benchmark of Linguistic Minimal Pairs (BLiMP) (Warstadt et al., 2020), which evaluates grammatical sensitivity across 67 controlled paradigms. By training classifiers on hidden states from each layer of LLMs, we determine where syntactic competence emerges and stabilizes. For meaning and reasoning probing, we use the Conceptual Minimal Pair Sentences (COMPS) dataset (Misra et al., 2023), which tests conceptual understanding and property inheritance. Its controlled sentence-pair design allows us to distinguish between surface-level semantic association and inference-based reasoning. We further extend reasoning evaluation with ProntoQA (Saparov and He, 2023) and Wino-Grande (Keisuke et al., 2019) to assess multi-hop deductive and commonsense reasoning abilities respectively. These datasets together enable a hierarchical dissection of linguistic representations. Detailed dataset structures, task formulations, and examples are provided in Appendix B.

A.2. Language Model

Our main experiments use the Qwen2.5-14B model, a 14.7B-parameter transformer with 48 layers (Team, 2025a). Each layer outputs hidden states of dimension 5120. The model supports a maximum context length of 131k tokens, though our experiments only use context sizes of 50 tokens. We use the base model without any instruction tuning or task-specific fine-tuning, and employ it both to probe the layer-wise emergence of linguistic and reasoning features and to perform brain encoding analyses.

To assess the robustness of our probing pipeline, we additionally examined other models across the Qwen family, spanning sizes from 1.8B to 14B parameters and multiple generations (Qwen1, Qwen1.5, Qwen2, Qwen2.5, and Qwen3) (Bai et al., 2023; Yang et al., 2024; Team, 2025a,b). We evaluated all models and found that Qwen2.5-14B exhibited the strongest reasoning capability, which is why we selected it as our primary model for subsequent analyses (see Appendix D.2).

A.3. Minimal Pair Probing

We apply minimal pair probing (He et al., 2024a,b) to identify the LLM layers most specialized for syntactic, meaning, and reasoning features. For each probing dataset, we extract sentence representations from each layer by feeding the sentence in isolation and taking the hidden state of its final token. We then use these sentence representations to train a logistic regression classifier to predict the correct item in each minimal pair. Model performance at a given layer is measured by the normalized F1 score of this classifier.

Let $H_l \in \mathbb{R}^{n \times d}$ denote the matrix of sentence embeddings at layer l , where each row corresponds to the final-token hidden state of one sentence in the minimal pair dataset. To determine the most specialized layer for a given feature, we define the saturation layer as the earliest layer where performance plateaus:

$$L_x := \min \{l \mid \forall l' > l, \text{Acc}^{\mathcal{D}_x}(H_{l'}) - \text{Acc}^{\mathcal{D}_x}(H_l) < \varepsilon\}, \quad x \in \{s, m, r\}$$

where $\mathcal{D}_s = \text{BLiMP}$, $\mathcal{D}_m = \text{COMPS-BASE}$, and $\mathcal{D}_r = \text{COMPS-WUGS-DIST}$, and ε is a small threshold representing the tolerance for marginal improvement. The hidden states at the identified saturation layers, denoted as $H_s := H_{L_s}$, $H_m := H_{L_m}$, and $H_r := H_{L_r}$, are used to construct feature-specific embeddings in the next stage.

A.4. Feature-Specific Embeddings

Building on the saturation layers identified via minimal pair probing, we construct four feature-specific embeddings to isolate distinct types of linguistic information: lexical, syntactic, meaning, and reasoning. Since hidden states at various layers contain overlapping linguistic information, we remove contributions from earlier representations to disentangle the targeted feature.

Lexical embedding. We define the lexical embedding $E_l \in \mathbb{R}^d$ as the hidden state at layer 0 of the LLM. As this layer follows token embeddings directly, it reflects uncontextualized lexical properties.

Residual embeddings for syntax, meaning, and reasoning. For the remaining features, syntactic, meaning, and reasoning, we construct residual embeddings by removing lower-level contributions from higher-layer representations. For instance, to isolate reasoning information, we remove the meaning contribution from the layer where reasoning saturates:

$$E_r := H_r - g_r(H_m), \quad \text{where } g_r = \arg \min_W \|H_r - WH_m\|_F^2 + \alpha \|W\|_F^2,$$

where g_r is a ridge regression trained via 4-fold cross validation on a podcast corpus described in the next paragraph (A.4). The same procedure applies to compute E_s and E_m . Specifically:

$$\begin{aligned} E_m &= H_m - g_s(H_s) && \text{(meaning minus syntactic)} \\ E_s &= H_s - g_l(H_l) && \text{(syntax minus lexical)} \end{aligned}$$

This yields feature-specific representations that are aligned with the linguistic hierarchy (Appendix C) and minimally confounded by lower-level signals.

Dataset for residual regression training. To extract feature-specific residual embeddings, we train ridge regression models that require a large number of training samples. However, the transcript that will later be used for neural alignment, which is drawn from a single 30-minute podcast episode, is too short to support stable regression. To address this, we train the models on an expanded corpus of 16 transcribed episodes from the same podcast series, including the episode used in the alignment analysis (Mao et al., 2020). This 160k-token dataset enables regression without PCA, preserving richer structure in the hidden states. We then apply the trained models to the target transcript to extract residuals for encoding analysis.

A.5. Encoding Model

ECoG dataset. After constructing feature-specific embeddings, we assess their neural alignment using the Podcast ECoG dataset (Zada et al., 2025). This dataset contains high-gamma band (70–200 Hz) intracranial recordings from nine participants as they listened

Algorithm 1: Construction of Feature-Specific Residual Embeddings

Input: LLM hidden states $\{H_L\}_{L=0}^{L_{\max}}$ for each token; probing datasets $\mathcal{D}_s, \mathcal{D}_m, \mathcal{D}_r$
Output: Feature-specific embeddings E_l, E_s, E_m, E_r

- 1 Perform probing with $\mathcal{D}_s, \mathcal{D}_m, \mathcal{D}_r$ to find saturation layers:;
- 2 $L_s \leftarrow$ syntax saturation layer from \mathcal{D}_s ;
- 3 $L_m \leftarrow$ meaning saturation layer from \mathcal{D}_m ;
- 4 $L_r \leftarrow$ reasoning saturation layer from \mathcal{D}_r ;
- 5 $L_l \leftarrow 0$;
- 6 Define lexical embedding: $E_l \leftarrow H_{L_l}$;
- 7 **for** each (L_{low}, L_{high}) in $\{(L_l, L_s), (L_s, L_m), (L_m, L_r)\}$ **do**
- 8 Train ridge regression g to predict $H_{L_{high}}$ from $H_{L_{low}}$;
- 9 Compute residual embedding: $E \leftarrow H_{L_{high}} - g(H_{L_{low}})$;
- 10 Assign E_s, E_m, E_r accordingly;
- 11 **end**

to a 30-minute narrative podcast. It includes 1,330 electrodes and a time-aligned word-level transcript, making it ideal for testing how lexical, syntactic, meaning, and reasoning embeddings align with neural activity during language comprehension.

As described in Section A.4, we train ridge regression models on an expanded podcast corpus to isolate feature-specific residuals. We then apply these models to the ECoG-aligned transcript to extract feature-specific embeddings, which are used to predict neural responses. Each embedding is aligned to individual word onsets, and for each word, neural signals are epoched in a ± 2 s window and downsampled to 32 Hz, yielding $t = 128$ time bins per event. Let $X \in \mathbb{R}^{n \times d}$ be the matrix of input embeddings (lexical, syntactic, meaning, or reasoning) across n word-aligned tokens, and $Y \in \mathbb{R}^{n \times (c \cdot t)}$ the corresponding ECoG response matrix across c electrodes and t time lags. We fit:

$$W^* = \arg \min_W \|Y - XW\|_F^2 + \alpha \|W\|_F^2,$$

with α selected via 5-fold cross-validation over a log-spaced grid, and $b = 5$ bootstrap resamples per fold using contiguous chunks of length $l = 32$. Model performance is quantified by Pearson correlation between predicted and actual signals. Temporal profiles are obtained by averaging over channels at each lag; spatial maps visualize per-channel peak correlations on 3D brain coordinates.

Word-rate feature regress out. We controlled for generic acoustic onset responses by adding a two-column word-rate covariate (word onsets and syllable-rate) to every ridge model. All variance-partitioning steps therefore quantify the variance explained beyond that attributable to mere word onsets. Let R_{full} denote the cross-validated prediction correlation achieved using the combined feature set (embedding + word rate), and let R_{wc} denote the correlation using only the word rate features. Assuming approximate orthogonality between the two feature sets, we estimated the unique contribution of the embedding features as: $R_{\text{embed}} = \text{sign}(R_{\text{full}}) \cdot \sqrt{\max(0, R_{\text{full}}^2 - R_{\text{wr}}^2)}$. This operation projects the full correlation vector onto the embedding-only axis, removing variance explained by word rate features. The

assumption of orthogonality is approximately satisfied due to the preprocessing pipeline, and helps to prevent over-attribution of shared variance.

Null distribution and responsiveness criterion. Considering that different channels and features have varying signal-to-noise ratios (SNRs), we constructed a subject–electrode–specific null distribution to assess whether a feature block explains neural activity beyond chance and to enable cross-electrode analysis. This was done by shuffling the feature rows 500 times while keeping the word-onset covariates fixed. For every shuffle we refit the ridge model and recorded the peak correlation R . Because Pearson correlations are bounded and skewed near ± 1 , we applied the standard Fisher z -transform (`atanh`) to all correlations, computed the shuffle mean and s.d. in z -space, and converted the true correlation to a z -score. Electrodes with $z > 3.95$ (one-tailed $\alpha = .05$, Bonferroni-corrected across $N = 1268$ electrodes) were deemed responsive, corresponding to values exceeding 3.95 standard deviations above the shuffle mean.

Appendix B. Detailed Description of Probing Datasets

BLiMP for syntactic probing. To construct feature-specific representations for brain alignment, we first identify LLM layers that specialize in syntactic encoding. Minimal pairs and probing techniques have been widely used in NLP field to access LLM’s internal representation (Marvin and Linzen, 2018; Linzen et al., 2016; Futrell et al., 2019; Hewitt and Manning, 2019; Manning et al., 2020; He et al., 2025). Here we use the Benchmark of Linguistic Minimal Pairs (BLiMP) (Warstadt et al., 2020), which evaluates syntactic knowledge through controlled minimal pairs differing in grammaticality across 67 paradigms. Each pair isolates a specific grammatical phenomenon such as subject–verb agreement, reflexive anaphora, or negative polarity. By training a classifier to distinguish acceptable from unacceptable sentences using hidden states from each LLM layer, we detect where syntactic competence emerges and peaks. For example, in a task targeting subject–verb agreement:

- a. The cats annoy Tim. (Correct)
- b. The cats annoys Tim. (Incorrect)

We average classifier performance across all paradigms to determine the saturation layer where syntactic features are stably encoded.

COMPS for Meaning and Reasoning Probing. For meaning and reasoning probing, we use the Conceptual Minimal Pair Sentences (COMPS) dataset (Misra et al., 2023). Sentence pairs in COMPS differ in whether the subject plausibly inherits a property. By probing whether each layer favors the correct concept-property association, we detect the layers where meaning and reasoning abilities peak.

COMPS-BASE evaluates surface-level understanding without requiring inference. For example, given the property “can heat food”:

- a. An oven can heat food. (Correct)
- b. A refrigerator can heat food. (Incorrect)

COMPS-WUGS-DIST increases reasoning complexity by replacing known concepts with nonsense words (e.g., “wug”, “dax”) and inserting distractor sentences that introduce interference. This design prevents the model from relying on surface-level co-occurrence or positional cues: distractors can be reordered. The model must infer the identity of the nonsense word from context to generalize property inheritance.

- a. A wug is an oven. A dax is a refrigerator. Therefore, a wug can heat food. (Correct)
- b. A wug is an oven. A dax is a refrigerator. Therefore, a dax can heat food. (Incorrect)

This progressive design increases reasoning difficulty and allows us to localize where models transition from semantic to reasoning-level representation.

Additional reasoning benchmarks. Beyond COMPS, we incorporate two additional reasoning benchmarks to validate robustness across task formats: ProntoQA (Saparov and He, 2023), which evaluates multi-hop deductive reasoning, and WinoGrande (Keisuke et al., 2019), which targets commonsense coreference resolution. For ProntoQA, we only use examples that require 5-hop deductive inferences to ensure high difficulty, where solving the task requires maintaining consistent symbolic relations across multiple premises and performing compositional logical deduction. Two illustrative examples are:

- a. Dogs are cats. Each dog is sour. Vertebrates are dull. Felines are dogs. Felines are not dull. Cows are felines. Each cow is aggressive. Snakes are cows. Snakes are orange. Animals are snakes. Every animal is not luminous. Mammals are animals. Each mammal is hot. Fae is a mammal. Therefore, Fae is not dull. (Correct)
- b. Every carnivore is a dog. Carnivores are angry. Every mammal is a carnivore. Each mammal is not red. Each snake is a mammal. Each snake is transparent. Cats are snakes. Cats are nervous. Each sheep is not angry. Each animal is a cat. Animals are earthy. Sam is an animal. Therefore, Sam is not angry. (Incorrect)

For WinoGrande, which tests commonsense reasoning through contextual coreference resolution, the model must leverage world knowledge and pragmatic inference to identify the correct referent under minimal lexical cues. A pair of examples is:

- a. John moved the couch from the garage to the backyard to create space. The garage is small. (Correct)
- b. John moved the couch from the garage to the backyard to create space. The backyard is small. (Incorrect)

Together, these datasets complement COMPS by spanning property inheritance, deduction, and commonsense inference. Detailed experiments and cross-validation results are provided in Appendix D.

Bad Probing Task Filtering We employ a Bag-of-Words (BoW) baseline and exclude tasks where the BoW model achieves an accuracy above 0.6, resulting in a final set of 29 BLiMP tasks (out of 67). The meaning (COMPS-BASE) and reasoning (COMPS-WUGS-DIST) tasks remain unchanged.

This filtering step is motivated by the fact that BoW readily captures tasks solvable via simple lexical cues, indicating a shallow design. By removing these superficially “easy” tasks, we ensure that the retained tasks demand genuine syntactic or semantic understanding rather than mere lexicon-based heuristics.

Appendix C. Hierarchical Emergence of Linguistic Features

In the main text, we reported probing results for Qwen2.5-14B (Section 3), showing that syntax, meaning, and reasoning features saturate at layers 6, 20, and 30, respectively. To evaluate the robustness of this emergence pattern, we extended our probing pipeline to 17 models across the Qwen family, including Qwen, Qwen1.5, Qwen2, Qwen2.5, and Qwen3 models.

Across nearly all models, we observed the same progression of feature emergence: syntax saturates earliest, followed by meaning, and then reasoning. The only exception was Qwen-1.8B, where syntax and meaning saturated at the same layer. Table 1 reports the saturation layers identified for all models.

Table 1: Saturation layers for syntax, meaning, and reasoning across Qwen models.

Model	Syntax	Meaning	Reasoning
Qwen-1.8B	11	11	14
Qwen-7B	11	13	15
Qwen-14B	9	16	20
Qwen1.5-1.8B	10	11	14
Qwen1.5-7B	9	13	16
Qwen1.5-14B	8	16	20
Qwen2-1.5B	10	14	17
Qwen2-7B	7	14	18
Qwen2.5-1.5B	7	14	18
Qwen2.5-7B	7	15	19
Qwen2.5-14B	6	20	30
Qwen3-1.7B	9	16	19
Qwen3-8B	7	20	23
Qwen3-14B	5	17	27

Appendix D. Consistency Across Reasoning Benchmarks

In the main text, reasoning saturation layers were identified using COMPS-WUGS-DIST. To evaluate robustness across tasks of varying complexity, we additionally applied two further reasoning probes: a 5-hop deductive reasoning task from ProntoQA and the WinoGrande benchmark for commonsense coreference.

We tested all three probes, COMPS-WUGS-DIST, ProntoQA, and WinoGrande, across 17 Qwen models spanning multiple sizes and training modes. Results show strong consistency: the average difference in reasoning saturation layers between COMPS-WUGS-DIST

and ProntoQA was 0.94, and between COMPS-WUGS-DIST and WinoGrande was 0.53. Given overall model depths of 25–49 layers, these differences are negligible and support the robustness of our reasoning probe.

D.1. Consistent Saturation Layers Across Reasoning Benchmarks.

Table 2 reports the reasoning saturation layers identified by each probe. The high degree of agreement across the three tasks indicates that reasoning-specific representations are consistently localized in deep layers, regardless of task format. The agreement across COMPS-WUGS-DIST, ProntoQA, and WinoGrande confirms that our identification of reasoning-specific saturation layers generalizes beyond a single benchmark. This consistency reinforces the conclusion that reasoning features emerge reliably in deep layers across architectures and reasoning formats.

Table 2: Reasoning saturation layers identified by different probes across Qwen models.

Model	ProntoQA	COMPS-WUGS-DIST	WinoGrande
Qwen-1.8B	13	14	14
Qwen-7B	13	15	16
Qwen-14B	18	20	20
Qwen1.5-1.8B	13	14	14
Qwen1.5-7B	14	16	16
Qwen1.5-14B	20	20	22
Qwen2-1.5B	17	17	17
Qwen2-7B	16	18	19
Qwen2.5-1.5B	16	18	17
Qwen2.5-7B	17	19	19
Qwen2.5-14B	29	30	29
Qwen3-1.7B	19	19	19
Qwen3-8B	23	23	23
Qwen3-14B	27	27	26

D.2. Qwen2.5-14B Achieves the Best Reasoning Performance (Small-Scale Models)

We evaluated the Qwen family on three reasoning benchmarks: ProntoQA (multi-hop deductive reasoning), COMPS-WUGS-DIST (property inheritance with distractors), and WinoGrande (commonsense coreference). Table 3 reports accuracy on each benchmark, along with the average across tasks. We observe consistent improvements in reasoning performance with increasing model scale and newer Qwen generations. Among the base models, Qwen2.5-14B achieves the highest average accuracy across benchmarks, reflecting its stronger reasoning capability, which is why we select it as the primary model for our experiments. These results complement our saturation layer analysis by showing that models allocating greater depth to reasoning layers also demonstrate superior task-level reasoning performance.

Table 3: Reasoning task performance across Qwen models. The Average column reports the mean accuracy across ProntoQA, COMPS-WUGS-DIST, and WinoGrande.

Model	ProntoQA	COMPS-WUGS-DIST	WinoGrande	Average
Qwen-1.8B	0.797	0.522	0.523	0.614
Qwen-7B	0.886	0.641	0.602	0.710
Qwen-14B	0.880	0.695	0.647	0.741
Qwen1.5-1.8B	0.792	0.513	0.523	0.609
Qwen1.5-7B	0.848	0.667	0.603	0.706
Qwen1.5-14B	0.910	0.670	0.653	0.744
Qwen2-1.5B	0.784	0.586	0.552	0.640
Qwen2-7B	0.851	0.636	0.660	0.716
Qwen2.5-1.5B	0.783	0.605	0.566	0.651
Qwen2.5-7B	0.879	0.673	0.664	0.739
Qwen2.5-14B	0.922	0.691	0.698	0.770
Qwen3-1.7B	0.739	0.500	0.530	0.589
Qwen3-8B	0.876	0.629	0.651	0.719
Qwen3-14B	0.912	0.716	0.670	0.766

Appendix E. Disentanglement Validation

E.1. Matrix-Level Orthogonality of Residual Embeddings

We justify the approximate orthogonality of the feature-specific embeddings E_l, E_s, E_m, E_r at the level of their embedding matrices. Each residual embedding can be represented as an $n \times d$ matrix, where n is the number of tokens in the dataset and d is the embedding dimension. Orthogonality here is considered along the sample axis: the n -dimensional feature columns of one residual embedding matrix are approximately uncorrelated with those of another.

This follows from the progressive emergence of linguistic features across LLM layers (Table 1): syntax peaks early, meaning in mid layers, and reasoning in later layers. Once a feature reaches saturation, its accuracy score remains stable in deeper layers, indicating that later representations retain earlier features. As a result, representations like H_m already embed information from H_l and H_s , making regression from H_m alone comparably as informative as from all three:

$$g_r(H_m) \approx g'_r([H_l, H_s, H_m]) \Rightarrow E_r \approx H_r - g'_r([H_l, H_s, H_m]) =: E'_r$$

Since E'_r is the residual of a linear projection onto $[H_l, H_s, H_m]$, it is orthogonal to each:

$$E'_r \perp H_l, \quad E'_r \perp H_s, \quad E'_r \perp H_m$$

Each residual embedding $E_i \in \{E_l, E_s, E_m\}$ is a linear combination of earlier hidden states (e.g., $E_m = H_m - g_m(H_s) = H_m - W_m H_s$). By the bilinearity of covariance, we have:

$$\text{Cov}(E_i, E'_r) = \text{Cov}(H_i - W_i H_j, E'_r) = \text{Cov}(H_i, E'_r) - W_i \text{Cov}(H_j, E'_r) = 0$$

whenever $E'_r \perp H_i$ and H_j , for appropriate $i, j \in \{l, s, m\}$. This implies:

$$\langle E'_r, E_i \rangle = 0 \quad \forall i \in \{l, s, m\}$$

Applying this proof across all residual stages, we conclude approximate mutual orthogonality of residual embedding matrices:

$$\langle E_i, E_j \rangle \approx 0 \quad \text{for all } i \neq j.$$

This sample-axis orthogonality ensures that residual embedding matrices contribute non-overlapping signals across the dataset, which is precisely the property required for the subsequent brain encoding analysis.

E.2. Token-Level Independence via Cosine Similarity

While Section E.1 establishes matrix-level orthogonality across the dataset, we also test independence at the level of individual tokens. For each token $i \in \{1, \dots, N\}$, let $E_l^i, E_s^i, E_m^i, E_r^i$ denote the four feature-specific residual vectors. We compute a 4×4 cosine similarity matrix C_i , take its absolute value $|C_i|$, and average across samples:

$$[C_i]_{j,k} = \frac{\langle E_j^i, E_k^i \rangle}{\|E_j^i\| \cdot \|E_k^i\|}, \quad \bar{C} = \frac{1}{N} \sum_{i=1}^N |C_i|, \quad j, k \in \{l, s, m, r\}$$

In the ideal case of perfect disentanglement, off-diagonal entries of \bar{C} would be zero, indicating orthogonality between different embeddings. Figure 2a compares the pairwise mean absolute cosine similarity among raw hidden states at the four saturation layers (top) and among the corresponding residual embeddings (bottom). While the hidden states show substantial overlap, especially between meaning and reasoning layers ($\bar{C} = 0.751$), the residual embeddings exhibit near-zero off-diagonal similarity (all ≤ 0.045) across all pairs.

Thus, Section E.1 and Section E.2 together provide complementary evidence: residuals are linearly novel across the sample axis, and they also align with near-orthogonal directions at the token level. These findings confirm that residualization yields disentangled representations of lexical, syntactic, semantic, and reasoning information.

E.3. Feature Probing on Residual Embeddings

While mutual independence ensures that residual embeddings are separated from one another, each residual should also preserve information relevant to its intended linguistic feature. We evaluate this by reapplying the same probing tasks used to define feature emergence. Since the lexical embedding E_l is directly extracted rather than residualized, we focus on the syntactic (E_s), meaning (E_m), and reasoning (E_r) residuals.

For each case, we train a logistic regression classifier on the corresponding residual embedding and report accuracy scores:

$$\hat{y}_i = \sigma(w_j^\top E_j^i + b_j), \quad j \in \{s, m, r\},$$

where $\sigma(\cdot)$ is the sigmoid function and $E_j^i \in \mathbb{R}^d$ is the residual embedding for token i .

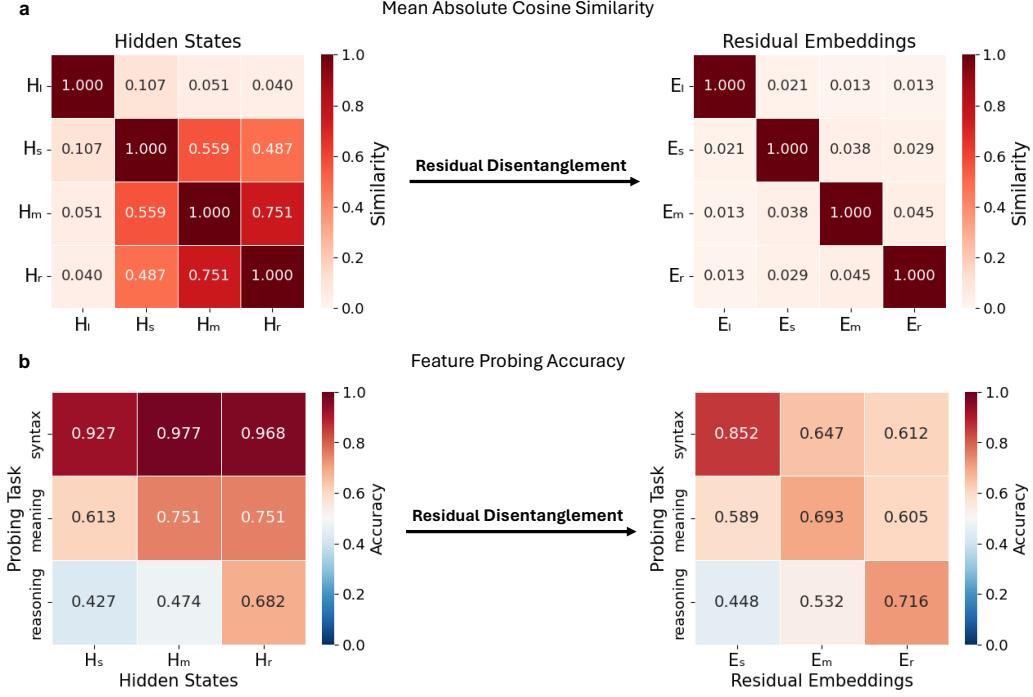


Figure 2: **a)** Pairwise cosine similarity among representations before (left) and after (right) residual disentanglement. The hidden states at feature saturation layers (H_l, H_s, H_m, H_r) exhibit substantial overlap. In contrast, the residual embeddings (E_l, E_s, E_m, E_r) show near-zero off-diagonal similarity. **b)** Feature probing results on hidden states (left) and residual embeddings (right), shown as accuracy. Each residual embedding achieves the highest performance on its corresponding task, while performing worse on unrelated tasks.

Results are shown in Figure 2b. Each residual achieves the highest performance on its own task (e.g., E_s on syntax probing, E_m on meaning, E_r on reasoning) while performing worse on unrelated tasks. The diagonal entries retain accuracy similar to the hidden-state baseline, whereas the off-diagonal entries drop substantially, with decreases of 33.8% (syntax task evaluated on the meaning residual), 36.8% (syntax task evaluated on the reasoning residual), and 19.4% (meaning task evaluated on the reasoning residual). These findings demonstrate that the residual embeddings are meaningfully disentangled and capture feature-specific information rather than reflecting general task complexity. A more detailed analysis of cross-task accuracies, including Bag-of-Word baseline, is provided in Appendix F.

Summary of Disentanglement Validation Across three complementary tests, layerwise saturation analysis, matrix-level orthogonality proof, and token-wise independence probing, residualization consistently produces four nearly orthogonal, interpretable embed-

dings corresponding to distinct linguistic levels. These embeddings form the basis for our subsequent neural encoding experiments.

Appendix F. Clarification on Residual Embedding Probing Results

In Section E.3, we reported that each residual embedding performs best on its corresponding task while performing worse on unrelated tasks. Some cross-task accuracies remain non-trivial, particularly on the meaning and reasoning probes. We provide additional clarification here.

Lexical-level baseline. A control experiment with a bag-of-words (BoW) model, which ignores syntax and word order, achieved 0.665 accuracy on COMPS-BASE. This shows that lexical cues alone provide a strong baseline for the meaning task, well above random chance (0.5). In this context, the raw scores of the syntax and reasoning residuals on COMPS-BASE (0.589 and 0.605, respectively) are in fact *below* the BoW baseline, indicating that residualization successfully removes spurious lexical and structural information from the meaning embedding. This control confirms that the observed cross-task scores are largely explained by lexical baselines artifacts. The syntax, meaning, and reasoning residuals each retain information most relevant to their target feature, supporting the interpretation that residual embeddings are meaningfully disentangled.

Appendix G. Encoding Results

G.1. Shallow features dominate neural prediction, but lexical activations are sparse.

Across time-aligned word events, encoding models based on shallower linguistic features, such as lexical identity and syntactic structure, showed the strongest neural correlations (Figure 3). Electrodes responsive to these features exhibited markedly higher peak correlation values than those driven by meaning or reasoning embeddings (Welch’s t-test, $p < 0.001$), confirming that lower-level linguistic information accounts for a disproportionate share of brain–model alignment under linear encoding. However, despite its high correlation strength, the lexicon embedding activated the fewest electrodes overall, suggesting that its predictive precision arises from a smaller, more specialized subset of cortical sites. In contrast, deeper features such as meaning and reasoning recruited broader but less strongly tuned populations. This dissociation between activation magnitude and spatial extent indicates that shallow representations, though limited in coverage, achieve tighter alignment within localized neural circuits, whereas deeper representations engage distributed cortical networks with weaker linear correspondence. Moreover, without disentanglement, the temporal dynamics in Figure 4 reveal that the full embedding is dominated by lexical and syntactic signals, whose early and sharp peaks mask the slower, high-level temporal signatures of meaning and reasoning. This highlights the necessity of residual separation for isolating deeper cognitive processes from low-level linguistic confounds.

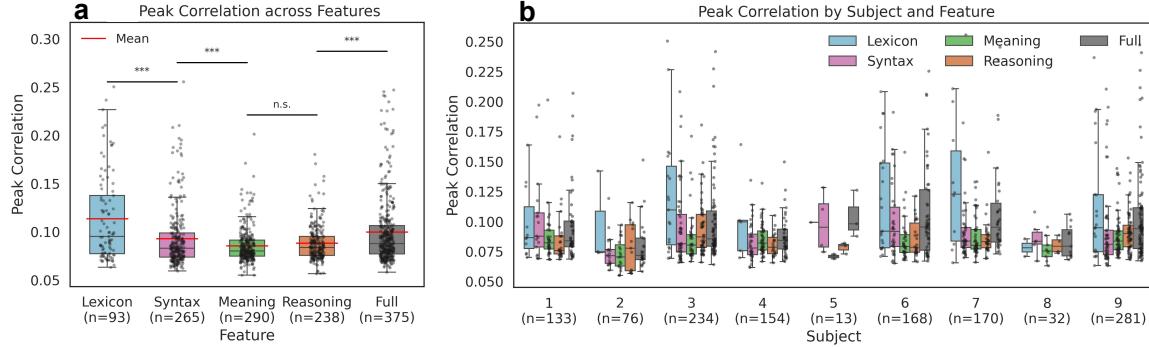


Figure 3: **a)** Peak correlation across linguistic feature models. As described in the Methods section, electrodes with $z \geq 3.95$ (one-tailed $\alpha = .05$, Bonferroni-corrected across $N = 1268$ electrodes) were deemed responsive, corresponding to values exceeding 3.95 standard deviations above the shuffle mean. Boxplots indicate the median and interquartile range across activated electrodes with red lines indicating means additionally. Asterisks mark significant differences between adjacent features (Welch's t-test; $p < 0.001$), showing that lexical and syntax features yield significantly higher correlations than high-level linguistic representations. **b)** Peak correlation by subject and feature. The number of responsive electrodes for each subject is shown below the x-axis. Lexical features generally exhibit higher correlations, indicating stronger neural alignment for lower-level linguistic representations.

G.2. Reasoning embedding shows different temporal pattern compared to Lexicon, Syntax and Meaning.

The temporal correlation profiles of Figure 4 illustrate a clear processing hierarchy among linguistic representations. Syntactic signals exhibit an early rise that begins before word onset and reach their maximum slightly pre-onset, suggesting that a large portion of predictive activity for upcoming words might be syntactically driven. Lexical features show a sharp and temporally confined peak immediately after onset, consistent with rapid low-level mapping between acoustic input and lexical identity. Meaning representations start to increase even before onset but peak afterward, indicating that contextual semantics contribute both to anticipatory prediction and to continued integration once the word is heard. In contrast, reasoning signals begin to rise only after onset and reach their peak at around 362 ms, reflecting a delayed, high-level computation that likely involves integrative or inferential processes beyond immediate linguistic parsing. Finally, the full embedding closely tracks the envelope of the four disentangled features, mirroring their composite dynamics and reinforcing the reliability of our residual separation procedure. Together, these temporal patterns reveal a cascading progression from predictive syntactic structure to contextual semantics, and finally to post-lexical reasoning operations in the human cortex.

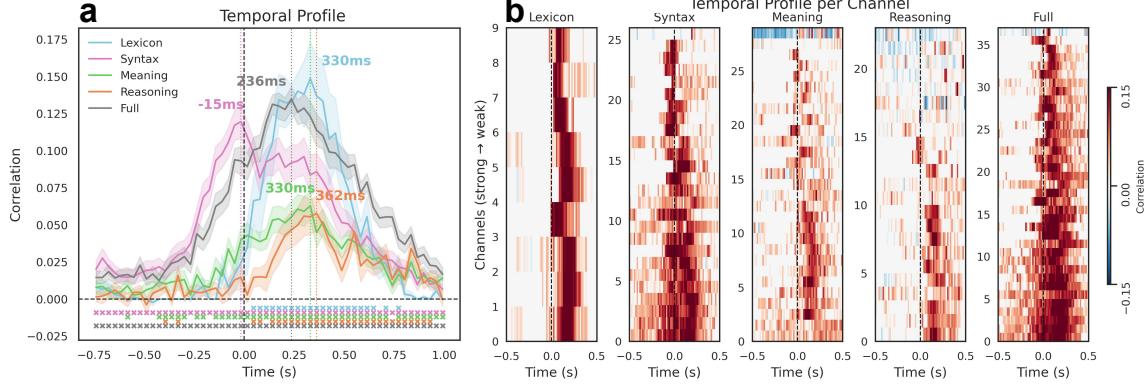


Figure 4: **a)** Average correlation time courses for the top 10% most responsive electrodes reveal distinct temporal dynamics across linguistic features. Responsive electrodes were identified following the strategy described in the Methods section. Syntactic signals rise before word onset, followed by meaning, lexicon and reasoning, which peak latest at ~ 362 ms. Significance markers (\times) indicate time points where correlations are significantly greater than zero (one-tailed t -test, FDR-corrected $q < 0.05$). **b)** Performance-based temporal receptive field (TRF) across time for electrodes selective to one of the five features: Lexicon, Syntax, Meaning, Reasoning, and Full. Electrodes were selected the same way as in panel a.

G.3. Reasoning recruits more than language area compared to low-level aspects.

As shown in Figure 5, shallow linguistic features such as lexicon exhibit high peak correlations but remain confined to classical language regions, including the inferior frontal gyrus (IFG) and superior temporal gyrus (STG). The strong yet spatially restricted activations indicate that lexical representations are predominantly localized within the traditional perisylvian language network, reflecting low-level linguistic encoding. In contrast, syntax, meaning, and particularly reasoning embeddings engage a broader set of cortical regions. While all three maintain robust responses in IFG and STG, reasoning-selective electrodes extend anteriorly into the superior frontal gyrus (SFG) and posteriorly into the visual cortex, showing significant activation compared with other features (Welch's t -test; $p < 0.001$; Figure 6). This expanded recruitment beyond canonical language areas suggests that reasoning might involve higher-order cognitive operations that integrate linguistic input with abstract inference, visual imagery, and executive control. Such cross-modal engagement underscores that reasoning-related brain activity cannot be fully explained by linguistic processing alone, but rather reflects the neural substrates of generalizable, multimodal cognition (Kosslyn et al., 1997; Ganis et al., 2004).

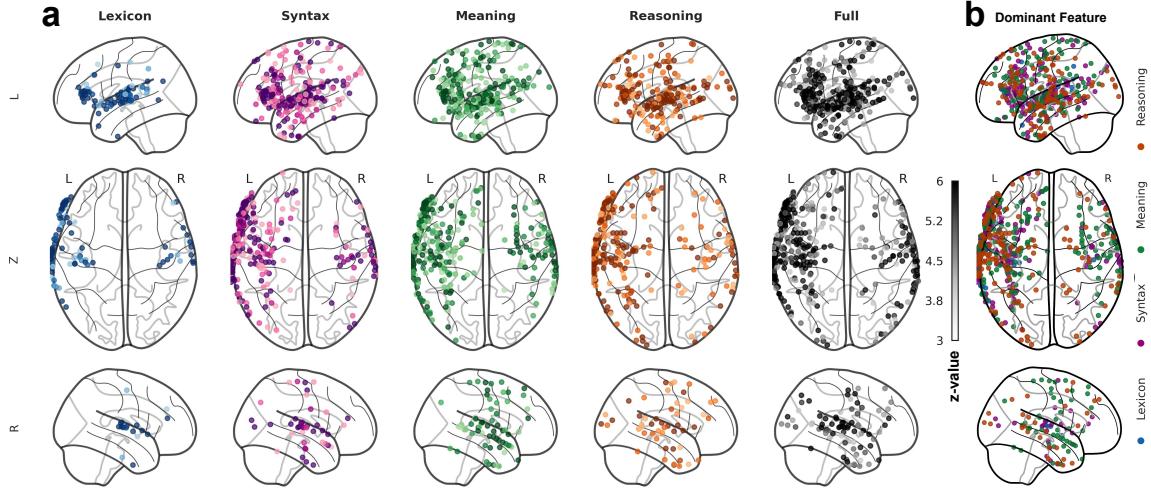


Figure 5: **a)** Spatial distribution of feature-selective responsive electrodes (MNI space). Displayed electrodes are those that met the responsiveness criterion defined in the Methods section. For visualization, z-scores were clipped to 3–6 to prevent extreme values from dominating the color scale and obscuring spatial patterns. **b)** For each electrode, the feature with the highest peak z-score among Lexicon, Syntax, Meaning, and Reasoning models is shown as its dominant feature. Colored dots indicate the dominant model. Across all electrodes, the counts of dominant features were: Syntax = 166, Meaning = 161, Reasoning = 128, and Lexicon = 42.

Appendix H. Supplemented Neuroscience Analysis

H.1. Activated-electrode Overlap Across Linguistic Feature Models

Figure 7. Activated-electrode overlap across feature models. Each cell shows the number of electrodes that were identified as active in both the feature model indicated by the row and that indicated by the column. Diagonal values represent the total number of electrodes active within each model, while off-diagonal values quantify cross-feature overlaps. Notably, high overlaps between Syntax and Meaning (143) and between Full and all subcomponents (around 180) suggest that electrodes responsive to linguistic processing are largely shared across multiple feature spaces. This pattern indicates that while distinct feature models capture partially unique neural activations, substantial commonality exists among them, particularly for integrative language representations.

H.2. Lateralization Analysis

To quantify hemispheric lateralization, we computed the proportion of active electrodes relative to total electrodes in the left and right hemispheres for each linguistic feature. Although the total number of electrodes is highly asymmetric across hemispheres (1057 on the left vs. 210 on the right), the normalized activation proportions still reveal clear trends.

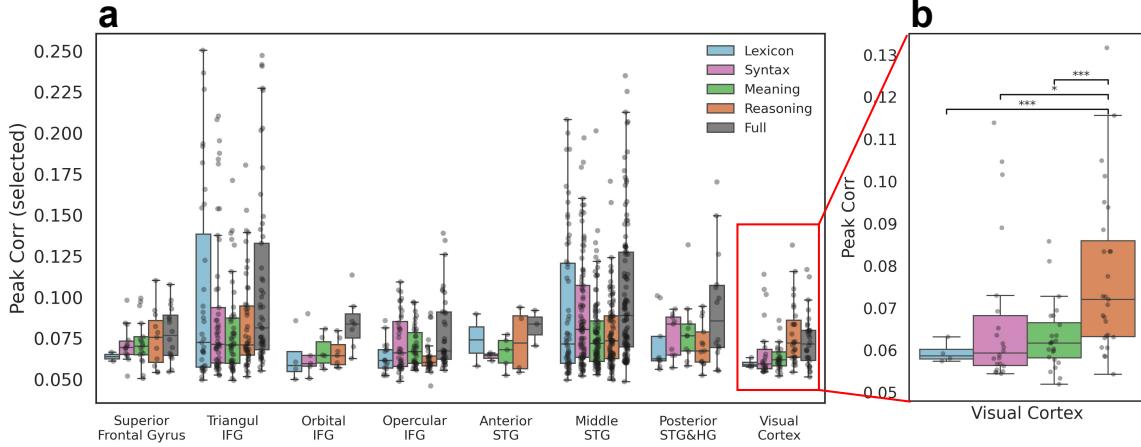


Figure 6: **a)** Peak correlation across cortical regions and feature models. Responsive electrodes were selected based on bilateral significance above the null shuffle distribution, $\alpha < 0.05$. The analysis focuses on language-related cortical regions, including the inferior frontal gyrus (IFG), superior temporal gyrus (STG), and surrounding Heschl's Gyrus (HG), as well as the superior frontal gyrus and visual cortex. The visual cortex group includes electrodes located in the superior occipital sulcus, transverse occipital sulcus, calcarine sulcus, occipital pole, superior occipital gyrus, middle occipital gyrus, and inferior occipital gyrus and sulcus. **b)** Welch's t-test (one-tailed) revealed that electrodes in visual cortex activated by the Reasoning feature showed significantly higher correlations than those activated by all other features ($p < 0.001$ for '***', $p < 0.05$ for '*').

Table 4: **Hemispheric lateralization ratios.** Each proportion represents the fraction of active electrodes relative to total electrodes within each hemisphere (1057 left vs. 211 right). The ratio indicates $\frac{\text{Left proportion}}{\text{Right proportion}}$.

Feature	#L	#R	Prop.L	Prop.R	L/R
Lexicon	76	17	0.07	0.08	0.89
Syntax	225	41	0.21	0.20	1.10
Meaning	227	63	0.21	0.30	0.72
Reasoning	193	51	0.18	0.24	0.76
Full	300	79	0.28	0.37	0.76

Specifically, Syntax-related encoding shows a higher proportion of left-hemisphere activations (0.21 vs. 0.20), resulting in the highest left/right ratio (1.09) among all features. In contrast, Meaning and Reasoning features exhibit relatively stronger right-hemisphere activations (ratios < 1). Given the uneven electrode coverage across hemispheres, these findings should not be overinterpreted as true right-lateralization for semantic or reasoning process-

		Activated-electrode Overlap				
		Lexicon	77	62	63	86
Active set from Reasoning	Lexicon	93	77	62	63	86
	Syntax	77	265	143	118	185
	Meaning	62	143	290	128	181
	Reasoning	63	118	128	238	175
	Full	86	185	181	175	375

Also active in

Figure 7: **Activated-electrode overlap across feature models.** Each cell shows the number of electrodes active in both feature sets, revealing substantial overlap among higher-level linguistic representations.

ing; however, they do suggest a stronger left-lateralization effect for shallow representations (lexicon and syntax).

H.3. Activated Channels in All Brain Area

Please refer to Figure 8.

Appendix I. Discussion

By isolating distinct types of linguistic information using residual embeddings, we disentangled the contributions of lexicon, syntax, meaning, and reasoning to the neural encoding of natural speech. This revealed new insights into the spatial and temporal organization of linguistic features that cannot be captured by full LLM embeddings used in prior studies (Jain and Huth, 2018; Toneva and Wehbe, 2019; Goldstein et al., 2022; Oota et al., 2022; Heilbron et al., 2022; Chen et al., 2023; Li et al., 2023; Antonello et al., 2023; Aw and Toneva, 2023; Chen et al., 2024a,b; Mischler et al., 2024; Caucheteux and King, 2022). Previous work has tried hierarchical approaches to decompose low-level linguistic and acoustic features in the brain (Keshishian et al. (2023, 2025)). However, our study is the first to apply hierarchical disentanglement framework to large language models by first localizing feature emergence layers through probing and then performing residual decoupling to construct layer-wise feature representations, thereby demonstrating its utility for isolating higher-order reasoning representations.

Our findings show that lower-level features, such as syntax and lexicon, elicit stronger neural activations than higher-level features like meaning or reasoning. This reflects both the stronger linear accessibility of shallow features and an inherent bias in full LLM embeddings, which primarily capture syntactic regularities while obscuring deeper computations

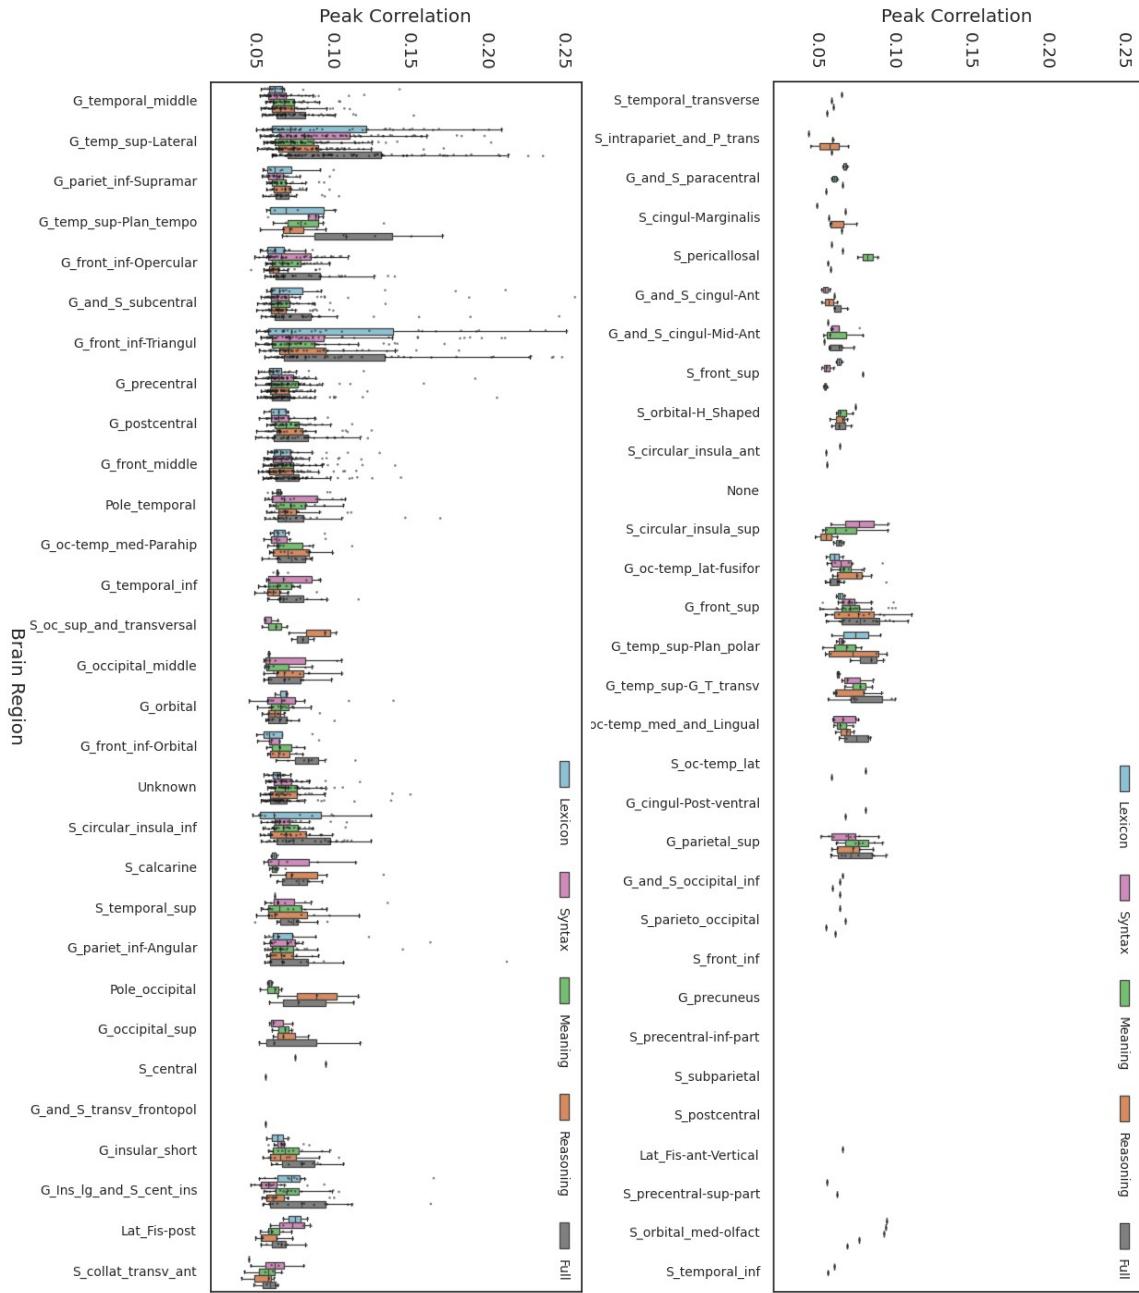


Figure 8: Peak correlation values of encoding performance across all cortical regions for five feature types. Responsive electrodes were selected based on bilateral significance above the null shuffle distribution, $\alpha < 0.05$.

(Lampinen et al., 2024). Residual disentanglement reveals clear spatiotemporal hierarchies: shallow features emerge earlier and localize to classical language regions (IFG, STG),

whereas reasoning signals arise later and extend beyond the perisylvian network into SFG and visual cortex. These results indicate that reasoning engages broader cortical systems that integrate linguistic processing with higher-level cognitive and multimodal functions. Together, our framework provides a principled approach for separating linguistic and cognitive components in LLMs and mapping them onto the human brain, offering an alignment-based perspective on how language and reasoning may correspond between artificial and neural systems.

Appendix J. Limitations and Future Directions

Despite these advances, several limitations remain. Our reliance on ECoG data, although temporally precise, provides limited spatial coverage, with insufficient sampling in frontal regions that are crucial for reasoning. Future work should integrate complementary modalities such as fMRI to improve spatial resolution. Also, while Qwen2.5-14B balances scale and accessibility, larger models optimized for multi-step reasoning may reveal stronger or more specialized reasoning signals. While reasoning representations saturate around layer 30 in Qwen2.5-14B, the role of subsequent layers remains unclear. We hypothesize that deeper layers may encode more abstract or generative representations, potentially related to discourse organization, narrative planning, or world modeling, beyond explicit reasoning. Exploring how these higher-level representations might align with cortical hierarchies or large-scale cognitive organization in the brain presents an especially interesting direction for future work.

Appendix K. Compute Resources and Execution Time

We report here the compute infrastructure and approximate runtime for each stage of our pipeline to support reproducibility.

- **Hidden State Extraction:** Performed using 4× NVIDIA L40 GPUs. Extracting hidden states for around 164,000 tokens from the Qwen2.5-14B model. The extraction took approximately 40 minutes and required around 4×30 GB of GPU memory.
- **Layer-wise Probing:** Conducted using 1× NVIDIA L40 GPUs. Each BLiMP task paradigm took 2.2 minutes on average, totaling around 57.2 minutes for the selected 26 paradigms focusing on syntax. COMPS-BASE consisting of 49,340 English sentence pairs took 1 hour and 47 minutes on average. COMPS-WUGS-DIST consisting of 27,792 pairs took 1 hour on average.
- **Residual Embedding Construction:** Ridge regression training was done using Scikit-learn on CPU with 30 GB of memory. Each regression model took less than 10 minutes to converge. NVIDIA cuML and GPU training was not adopted due to lack of support for multi-output ridge regression training.
- **Brain Encoding (ECoG):** Conducted using CPU with self-developed package. Aligning model activations and training encoding models over all channels and time lags took approximately 5 hours for all experiments done.

Overall, the full pipeline can be reproduced on a modern workstation or cloud instance equipped with 1 NVIDIA L40 GPU + 30 GB RAM.

Appendix L. Licenses and Attribution of External Assets

This work uses several publicly available datasets and a language model, all of which are properly cited and used in accordance with their respective licenses. No proprietary or scraped data was used.

- **BLiMP Dataset** (Warstadt et al., 2020): A suite of syntactic probing tasks for language models.
 - URL: <https://github.com/alexwarstadt/blimp>
 - License: CC-BY 4.0
- **COMPS Datasets** (Misra et al., 2023): COMPS-BASE and COMPS-WUGS-DIST are used to assess semantic and reasoning representations.
 - URL: <https://github.com/kanishkamisra/comps/>
 - License: Apache License 2.0
- **The Podcast ECoG Dataset** (Zada et al., 2025): High-resolution electrocorticographic recordings from participants listening to a natural podcast.
 - URL: <https://openneuro.org/datasets/ds005574/versions/1.0.2>
 - License: CC0
- **Expanded Podcast Transcripts** (Mao et al., 2020): Text for additional podcast episodes used to extend ECoG analysis.
 - URL: <https://github.com/calclavia/tal-asrd>
 - License: No explicit license is specified on the repository. However, the author provides access to the dataset and indicates that it can be downloaded for research purposes. We used the dataset solely for non-commercial academic research and did not redistribute or modify it.
- **Qwen2.5-14B** (Team, 2025a): Language model for LLM representations extraction.
 - URL: <https://huggingface.co/Qwen/Qwen2.5-14B>
 - License: Apache license 2.0

All datasets and model were used solely for research purposes. We did not modify or redistribute any datasets beyond standard preprocessing steps for experimental use. All license terms were respected.