Acoustic Degradation Reweights Cortical and ASR Processing: A Brain-Model Alignment Study

Francis Pingfan Chien[*,1], **Chia-Chun Dan Hsu**[*,2], **Po-Jang Hsieh**[3], **Yu Tsao**[2]
[1] Taiwan International Graduate Program in Interdisciplinary Neuroscience,
National Taiwan University and Academia Sinica, Taipei, Taiwan
[2] Research Center for Information Technology Innovation, Academia Sinica, Taipei, Taiwan
[3] Department of Psychology, National Taiwan University, Taipei, Taiwan
f12b49004@ntu.edu.tw, hali2680@gmail.com
hsiehpj@ntu.edu.tw, yu.tsao@citi.sinica.edu.tw

Abstract

We tested whether acoustic degradation changes how a modern ASR represents speech and whether those changes explain human brain and behavior. Twenty-five participants listened to clean and noisy (-3 dB SNR) Mandarin sentences during fMRI while we extracted layer-wise embeddings from Whisper-Tiny. We computed brain scores normalized to each ROI's noise ceiling. Behavioral assessments, including intelligibility, perceived quality, and comprehension, declined under noise. Under clean speech, alignment emphasized frontal predictive processing, with encoder layers 3 and 4 peaking in the right inferior frontal gyrus (IFG) and decoder layer 2 peaking in the left middle frontal gyrus (MFG). Under noisy speech, alignment shifted toward early acoustic and evaluative regions, with encoder layer 1 peaking in the right Heschl's gyrus and encoder layer 4 peaking in the right IFG pars orbitalis (IFGorb), and decoder peaks were weaker and more diffuse. Condition contrasts showed higher alignment for clean speech in the right IFG (encoder layers 3 and 4) and the left MFG (decoder layer 2). These findings demonstrate a processing account, a behavioral link, and a compact layer-to-region map across listening conditions.

1 Introduction

Everyday listening often involves challenges such as background noise, overlapping speech, or degraded signals. When acoustic signals degrade, listening becomes effortful. Comprehension declines, and cortical resources are reallocated toward the frontal regions responsible for attention and control, as demonstrated by previous studies [1, 2]. These findings raise an important question: can modern machine systems, particularly automatic speech recognition (ASR) models, exhibit analogous processing shifts under similar acoustic conditions? Recent work goes beyond surface-level prediction of neural responses and emphasizes brain–model alignment that maps internal model computations to cortical anatomy and function [3, 4]. In language processing, integrative modeling converges on predictive processing as a central organizing principle [5]. Research suggests that alignment quality depends on both the depth and the dynamics of computation, including mechanisms such as recurrence and attention [6, 7]. Additionally, models directly optimized to predict neural data have helped clarify which architectural features support brain-like representations [8]. Studies of auditory and language processing have shown that computational stages in models often correspond to specific brain regions [9–11]. Speech comprehension is a critical domain in which both the human cortex and Transformer-based ASR models exhibit hierarchical processing, transitioning from low-level

* Equal contribution.

39th Conference on Neural Information Processing Systems (NeurIPS 2025) Workshop: CogInterp.

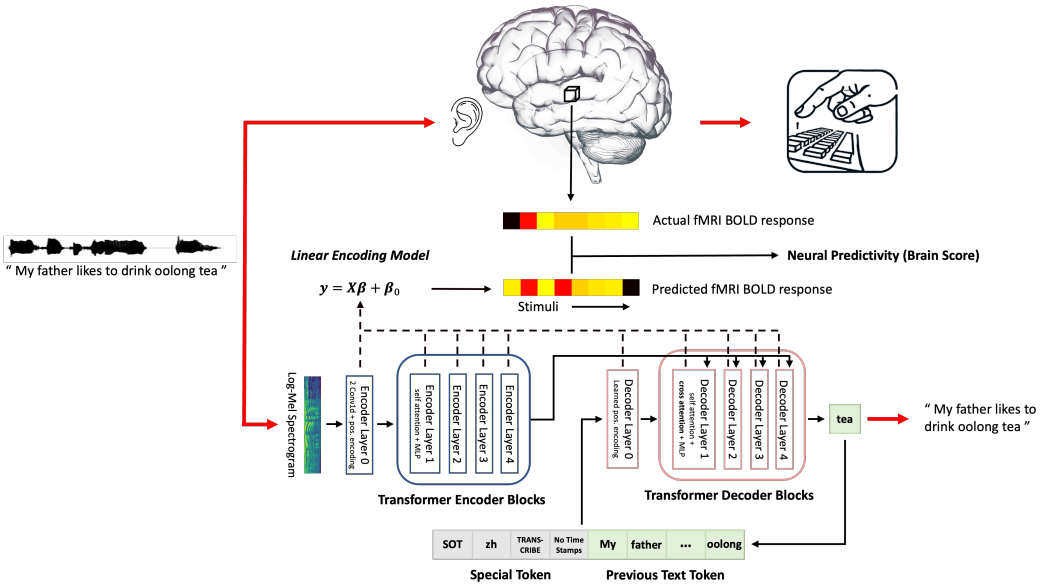

Figure 1: **Brain–model alignment pipeline from stimulus to model-brain to behavior.** Clean or noisy sentences are presented to Whisper Tiny (encoder–decoder Transformer) and to participants during fMRI. Layer-wise embeddings are used to predict BOLD in predefined ROIs—primary auditory cortex (Heschl's gyrus) and bilateral language-network regions: anterior temporal lobe (AntTemp), posterior temporal lobe (PostTemp), angular gyrus (AngG), inferior frontal gyrus (IFG), orbital IFG (IFGorb), and middle frontal gyrus (MFG)—using ridge encoding. Prediction accuracy, normalized to each ROI's noise ceiling (brain score), indexes alignment. After each trial, participants report intelligibility, perceived quality, and comprehension.

acoustic to higher-level lexical and semantic representations [12, 13]. Prior evidence suggests that when sensory input is unreliable, the brain downweights predictive mechanisms and relies more on bottom-up acoustic features. If ASR models reflect similar internal shifts, alignment with brain activity may reveal where and how such reweighting occurs. Emerging research also links linguistic embedding spaces with brain signals during natural conversation, underscoring the value of alignment for cognitive understanding [14, 15].

Here, we extend this approach by contrasting clean and noisy speech, a defining feature of everyday listening. We examine whether background noise alters model–brain correspondences. Specifically, we test whether noise alters alignment toward earlier encoder layers associated with low-level auditory processing, whereas clean input strengthens coupling with later predictive representations in frontal regions. Figure 1 outlines the full pipeline, from auditory stimuli to model representations, to fMRI responses and behavioral outcomes. We address three questions:

(1) Do intelligibility, quality, and comprehension differ across conditions, and do alignment patterns reflect these differences?

(2) Where in the brain, and at what model depth, does alignment change with degradation?

(3) Can condition-specific alignment explain how neural processing mechanisms lead to behavioral outcomes?

2 Methods

2.1 Participants & Stimuli & Experiment

We recruited 25 healthy native Mandarin speakers with normal hearing. Participants listened to 24 clean and 24 noisy Mandarin sentences (10 words; ∼3-second each) presented via MRI-compatible headphones. Noisy trials were created by mixing each sentence with stationary speech-shaped noise

(SSN) at -3 dB SNR [16]. After each sentence, while still in the scanner, they (i) rated intelligibility (how well did you comprehend this speech?), (ii) rated perceived speech quality (how would you rate the quality of this speech?), and (iii) answered a two-option comprehension question about the sentence. Intelligibility and quality were reported on a 5-point mean opinion score (MOS) scale (1 = bad, 2 = poor, 3 = fair, 4 = good, 5 = excellent) [17]. Comprehension was scored as correct or incorrect.

2.2 fMRI Data & ROI Definition

MRI scanning was performed on a Siemens Magnetom Skyra 3T scanner. Anatomical images were acquired using a T1-weighted multi-echo MP-RAGE sequence (1 mm^3 voxels). Functional images were collected with a gradient-echo EPI sequence (TR = 2000 ms; TE = 24 ms; FOV = 220 $\times$ 220 mm; 38 slices; voxel = 3.4 $\times$ 3.4 $\times$ 4.0 mm ; flip angle = 90°). MRI preprocessing was performed in SPM12 (`http://www.fil.ion.ucl.ac.uk/spm/software/spm12`), including slice-timing correction, motion correction, co-registration of functional to anatomical images, normalization to MNI space, and smoothing with an 8 mm FWHM Gaussian kernel. First-level GLMs were specified with boxcar regressors for clean and noisy conditions, convolved with the Hemodynamic Response Function (HRF), plus six motion parameters as nuisance covariates. We employed a fixed set of subject-specific Fedorenko language fROIs: bilateral anterior temporal cortex (AntTemp), posterior temporal cortex (PostTemp), angular gyrus (Angular), inferior frontal gyrus (triangular and orbital subdivisions), and middle frontal gyrus (MFG) [18, 19]. Two additional auditory ROIs (bilateral Heschl's gyri) were defined anatomically using the Automated Anatomical Labeling (AAL) atlas, extracted with the MarsBaR toolbox [20].

2.3 ASR Model Representations

In this study, we evaluate the alignment relationship between brain and ASR model using layer-wise representations from Whisper-tiny [21]. In addition to Whisper-tiny's four transformer blocks in both encoder and decoder, we include the zeroth layer representation before each module. For each encoder layer ℓ, Whisper-tiny generates 1500 hidden states for a 30-second input (each state representing $\sim$20ms), with dimensionality 384. Since our stimuli last approximately 3 seconds, we retain only the first 200 hidden states (corresponding to 4 seconds) to capture the entire stimuli duration, obtaining $\mathbf{E}^{(\ell)} \in \mathbb{R}^{200\times384}$. We then flatten to $\text{vec}\big(\mathbf{E}^{(\ell)}\big) \in \mathbb{R}^{76800}$ and use this as the encoder layer-ℓ representation. The decoder in Whisper is an auto-regressive transformer that predicts the next token at each step, conditioned on the encoder's final representation and previously generated tokens. This objective encourages the pre-terminal hidden state to summarize the entire sentence information. Hence, we use the hidden state at the last non-special token position as the decoder layer-ℓ representation, $\mathbf{D}^{(\ell)} \in \mathbb{R}^{384}$.

2.4 Linear Encoding

We employ ridge regression as our linear encoding model to map ASR layer representations to BOLD responses across different ROIs and stimulus conditions. Ridge regression imposes constraints on the magnitude and covariance of regression weights to enhance weight estimation when dealing with limited and noisy data [22]. For implementation, we first reduce the encoder layer representation dimensionality from 76,800 to 2,796 using sparse random projection [23, 24]. This target dimensionality is determined by the Johnson-Lindenstrauss lemma with 24 samples and a distortion parameter $\epsilon = 0.1$ [25]. We apply scikit-learn's RidgeCV [26], selecting optimal regularization hyperparameters α for each voxel via nested leave-one-out cross-validation from 60 logarithmically spaced values between 10^{-6} and 10^{6}. To evaluate the alignment relationship between ROIs and ASR model representations, we adopt the brain score metric [5]. We compute Pearson correlations between predicted and actual BOLD responses on a test set within the 4-fold cross-validation scheme. For each participant, we calculate the median predictivity across all voxels averaged over the 4 splits as the ROI raw score. Finally, we normalize this raw score by the noise ceiling, representing the theoretical upper bound of model performance given cross-subject consistency [5], yielding a brain score in the range [0, 1].

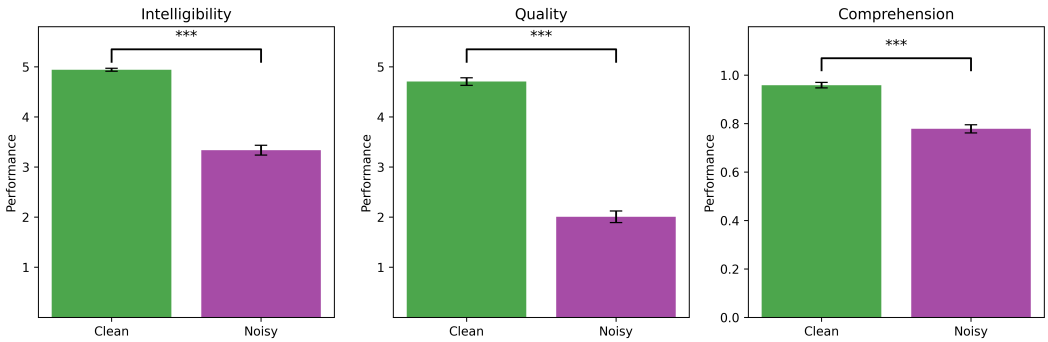

Figure 2: **Behavior under clean and noisy speech: Intelligibility, Quality, and Comprehension.** Group means for intelligibility and perceived quality (MOS 1–5) and comprehension accuracy (0–1) while participants ($n = 25$) listened to 24 clean and 24 noisy (-3 dB SNR) Mandarin sentences during fMRI. Asterisks indicate significant clean–noisy differences after Benjamini–Hochberg FDR correction.

3 Results

Behavior While undergoing fMRI, 25 participants rated the intelligibility and perceived quality of each sentence and answered a comprehension question. All three measures were significantly lower for noisy than for clean sentences (FDR–BH, $p < 0.001$; Figure 2).

Clean speech Encoder predictivity increased with depth, peaking at layer 3 in the right IFG (BS = 0.76). Decoder alignment peaked at layer 2 in the left MFG (BS = 0.80). Consistent with Table 1, the right IFG was the best-performing ROI for encoder layers 1–4, whereas the left MFG and left Heschl's gyrus showed the strongest decoder peaks. (Figure 4a, Table 1)

Noisy speech The layer profile flattened, and encoder maxima shifted toward primary sensory and evaluative regions, with peaks at layer 1 in the right Heschl's gyrus (BS = 0.76) and layer 4 in the right IFG pars orbitalis (IFGorb; BS = 0.80). Decoder peaks were weaker and more dispersed, with highest values in the MFG and anterior temporal cortex (typical peak $\sim$0.71).(Figure 4b, Table 1)

Condition effects To quantify degradation effects, we compared BS across conditions using paired tests with Benjamini–Hochberg correction (Figure 3, 5). Brain scores were significantly higher for clean than for noisy speech in the right IFG at encoder layers 3 and 4 and in the left MFG at decoder layer 2 ($q < 0.05$). In Figure3, clean is plotted in green and noisy in purple. The gray line marks the noise ceiling, and error bars indicate ± 1 SEM (n = 25).

4 Discussion

Our results show that clean speech preferentially engages frontal language areas, consistent with predictive, hierarchical processing. Specifically, encoder layers 3 and 4 aligned with the right inferior frontal gyrus (IFG), and decoder layer 2 aligned with the left middle frontal gyrus (MFG), echoing prior links between IFG and MFG and speech integration and prediction in natural language processing [14, 27–30]. Under noisy speech, alignment shifted toward earlier encoder layers and sensory or orbitofrontal regions, with peaks at encoder layer 1 in the right Heschl's gyrus and encoder layer 4 in the right IFG pars orbitalis (IFGorb), while decoder alignment was weaker and more diffuse. This redistribution supports a precision-weighted account, in which lower input reliability shifts processing toward bottom-up acoustic analysis over top-down prediction, consistent with multisensory stabilization mechanisms [31–34]. Taken together, these effects produce a layer–ROI map that reflects cortical hierarchy. Deeper model stages align with frontal regions under clean input, whereas shallower stages align with auditory regions under noise. Behaviorally, lower intelligibility, perceived quality, and comprehension in the noisy condition mirror these neural shifts, reinforcing a functional link between brain–model alignment and perceptual performance.

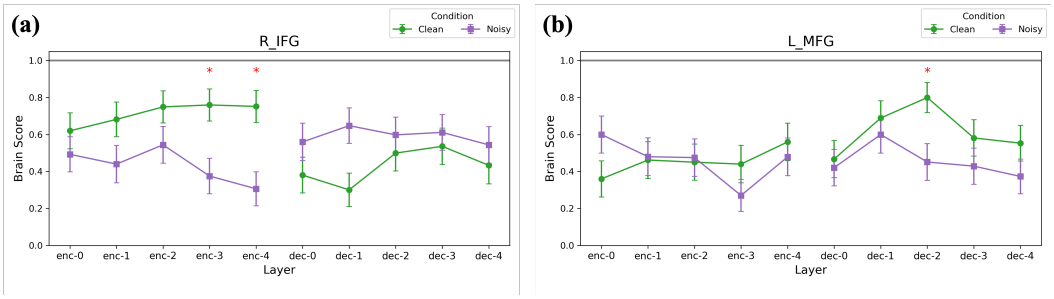

Figure 3: **Significant differences in model–brain alignment between clean and noisy speech conditions after FDR-BH correction.** (a) Right inferior frontal gyrus (R_IFG): Brain scores for clean (green) and noisy (purple) speech across encoder and decoder layers. Clean speech showed significantly higher alignment in encoder layers 3 and 4 (FDR-corrected $p < 0.05$, red asterisks). (b) Left middle frontal gyrus (L_MFG): Clean speech exhibited significantly higher alignment in decoder layer 2 (FDR-corrected $p < 0.05$, red asterisk). The horizontal grey line denotes the noise ceiling (normalized to 1). Error bars indicate ± 1 SEM across 25 participants.

Table 1: Best-performing ROI for each Whisper-tiny layer under clean and noisy speech conditions

	Clean Speech		Noisy Speech	
Layer	ROI	Brain Score	ROI	Brain Score
Encoder layer 0	Right IFGorb	0.63	Right IFGorb	0.66
Encoder layer 1	Right IFG	0.68	Right Heschl	0.76
Encoder layer 2	Right IFG	0.75	Right IFGorb	0.68
Encoder layer 3	Right IFG	0.76	Right IFGorb	0.65
Encoder layer 4	Right IFG	0.75	Right IFGorb	0.80
Decoder layer 0	Right Heschl	0.67	Right MFG	0.60
Decoder layer 1	Left MFG	0.69	Right IFG	0.65
Decoder layer 2	Left MFG	0.80	Right AntTemp	0.70
Decoder layer 3	Left Heschl	0.66	Left AntTemp	0.64
Decoder layer 4	Left Heschl	0.67	Right MFG	0.71

Contributions and Novelty:

(1) A processing account showing condition-sensitive reweighting in both the model and the brain, with the clean speech favoring decoder-centric predictive alignment and the noisy speech favoring encoder-centric acoustic alignment.

(2) A behavioral account that ties this reweighting to measurable declines in intelligibility, quality, and comprehension.

(3) An algorithm to a brain mapping framework that locates strategy changes across model depth and cortical regions, for example, the early encoder with Heschl's gyrus, the middle to late encoder with the IFG or IFGorb, and the middle decoder with MFG, providing a compact process-level explanation of how the brain adapts to acoustic challenges.

Future studies should explore causal manipulations, such as constraining decoder context or modifying early encoder features, and examine whether individual differences in alignment shifts can predict behavioral variability in speech comprehension.

References

[1] Jonathan E Peelle. Listening effort: How the cognitive consequences of acoustic challenge are reflected in brain and behavior. *Ear and hearing*, 39(2):204–214, 2018.

[2] Conor J Wild, Afiqah Yusuf, Daryl E Wilson, Jonathan E Peelle, Matthew H Davis, and Ingrid S Johnsrude. Effortful listening: the processing of degraded speech depends critically on attention. *Journal of Neuroscience*, 32(40):14010–14021, 2012.

[3] Daniel LK Yamins and James J DiCarlo. Using goal-driven deep learning models to understand sensory cortex. *Nature neuroscience*, 19(3):356–365, 2016.

[4] Martin Schrimpf, Jonas Kubilius, Michael J Lee, N Apurva Ratan Murty, Robert Ajemian, and James J DiCarlo. Integrative benchmarking to advance neurally mechanistic models of human intelligence. *Neuron*, 108(3):413–423, 2020.

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

[6] Tim C Kietzmann, Courtney J Spoerer, Lynn KA Sörensen, Radoslaw M Cichy, Olaf Hauk, and Nikolaus Kriegeskorte. Recurrence is required to capture the representational dynamics of the human visual system. *Proceedings of the National Academy of Sciences*, 116(43):21854–21863, 2019.

[7] Mariya Toneva and Leila Wehbe. Interpreting and improving natural-language processing (in machines) with natural language-processing (in the brain). *Advances in neural information processing systems*, 32, 2019.

[8] Ghislain St-Yves, Emily J Allen, Yihan Wu, Kendrick Kay, and Thomas Naselaris. Brain-optimized deep neural network models of human visual areas learn non-hierarchical representations. *Nature communications*, 14(1):3329, 2023.

[9] Greta Tuckute, Jenelle Feather, Dana Boebinger, and Josh H McDermott. Many but not all deep neural network audio models capture brain responses and exhibit correspondence between model stages and brain regions. *Plos Biology*, 21(12):e3002366, 2023.

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

[11] Eghbal A Hosseini, Martin Schrimpf, Yian Zhang, Samuel Bowman, Noga Zaslavsky, and Evelina Fedorenko. 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, 2024.

[12] Juliette Millet, Charlotte Caucheteux, Yves Boubenec, Alexandre Gramfort, Ewan Dunbar, Christophe Pallier, Jean-Remi King, et al. Toward a realistic model of speech processing in the brain with self-supervised learning. *Advances in Neural Information Processing Systems*, 35: 33428–33443, 2022.

[13] Charlotte Caucheteux, Alexandre Gramfort, and Jean-Rémi King. Evidence of a predictive coding hierarchy in the human brain listening to speech. *Nature human behaviour*, 7(3):430–441, 2023.

[14] Ariel Goldstein, Haocheng Wang, Leonard Niekerken, Mariano Schain, Zaid Zada, Bobbi Aubrey, Tom Sheffer, Samuel A Nastase, Harshvardhan Gazula, Aditi Singh, et al. A unified acoustic-to-speech-to-language embedding space captures the neural basis of natural language processing in everyday conversations. *Nature human behaviour*, pages 1–15, 2025.

[15] Subba Reddy Oota, Manish Gupta, and Mariya Toneva. Joint processing of linguistic properties in brains and language models. *Advances in Neural Information Processing Systems*, 36: 18001–18014, 2023.

[16] Paul Glad Mihai, Nadja Tschentscher, and Katharina von Kriegstein. Modulation of the primary auditory thalamus when recognizing speech with background noise. *Journal of Neuroscience*, 41(33):7136–7147, 2021.

[17] Hassan Akbari, Bahar Khalighinejad, Jose L Herrero, Ashesh D Mehta, and Nima Mesgarani. Towards reconstructing intelligible speech from the human auditory cortex. *Scientific reports*, 9 (1):874, 2019.

[18] Evelina Fedorenko, Po-Jang Hsieh, Alfonso Nieto-Castañón, Susan Whitfield-Gabrieli, and Nancy Kanwisher. New method for fmri investigations of language: defining rois functionally in individual subjects. *Journal of neurophysiology*, 104(2):1177–1194, 2010.

[19] Evelina Fedorenko, Anna A Ivanova, and Tamar I Regev. The language network as a natural kind within the broader landscape of the human brain. *Nature Reviews Neuroscience*, 25(5): 289–312, 2024.

[20] Matthew Brett, Jean-Luc Anton, Romain Valabregue, Jean-Baptiste Poline, et al. Region of interest analysis using the marsbar toolbox for spm 99. *Neuroimage*, 16(2):S497, 2002.

[21] Alec Radford, Jong Wook Kim, Tao Xu, Greg Brockman, Christine McLeavey, and Ilya Sutskever. Robust speech recognition via large-scale weak supervision. In *International conference on machine learning*, pages 28492–28518. PMLR, 2023.

[22] Arthur E Hoerl and Robert W Kennard. Ridge regression: Biased estimation for nonorthogonal problems. *Technometrics*, 12(1):55–67, 1970.

[23] Colin Conwell, David Mayo, Andrei Barbu, Michael Buice, George Alvarez, and Boris Katz. Neural regression, representational similarity, model zoology & neural taskonomy at scale in rodent visual cortex. *Advances in Neural Information Processing Systems*, 34:5590–5607, 2021.

[24] Colin Conwell, Jacob S Prince, Kendrick N Kay, George A Alvarez, and Talia Konkle. A large-scale examination of inductive biases shaping high-level visual representation in brains and machines. *Nature communications*, 15(1):9383, 2024.

[25] Dimitris Achlioptas. Database-friendly random projections. In *Proceedings of the twentieth ACM SIGMOD-SIGACT-SIGART symposium on Principles of database systems*, pages 274–281, 2001.

[26] Fabian Pedregosa, Gaël Varoquaux, Alexandre Gramfort, Vincent Michel, Bertrand Thirion, Olivier Grisel, Mathieu Blondel, Peter Prettenhofer, Ron Weiss, Vincent Dubourg, et al. Scikit-learn: Machine learning in python. *the Journal of machine Learning research*, 12:2825–2830, 2011.

[27] Angela D Friederici. Hierarchy processing in human neurobiology: how specific is it? *Philosophical Transactions of the Royal Society B*, 375(1789):20180391, 2020.

[28] Adeen Flinker, Anna Korzeniewska, Avgusta Y Shestyuk, Piotr J Franaszczuk, Nina F Dronkers, Robert T Knight, and Nathan E Crone. Redefining the role of broca's area in speech. *Proceedings of the National Academy of Sciences*, 112(9):2871–2875, 2015.

[29] Ariel Goldstein, Avigail Grinstein-Dabush, Mariano Schain, Haocheng Wang, Zhuoqiao Hong, Bobbi Aubrey, Samuel A Nastase, Zaid Zada, Eric Ham, Amir Feder, et al. Alignment of brain embeddings and artificial contextual embeddings in natural language points to common geometric patterns. *Nature communications*, 15(1):2768, 2024.

[30] 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.

[31] Evelina Fedorenko, John Duncan, and Nancy Kanwisher. Broad domain generality in focal regions of frontal and parietal cortex. *Proceedings of the National Academy of Sciences*, 110 (41):16616–16621, 2013.

[32] Roel M Willems, Stefan L Frank, Annabel D Nijhof, Peter Hagoort, and Antal Van den Bosch. Prediction during natural language comprehension. *Cerebral cortex*, 26(6):2506–2516, 2016.

[33] Christopher W Bishop and Lee M Miller. A multisensory cortical network for understanding speech in noise. *Journal of cognitive neuroscience*, 21(9):1790–1804, 2009.

[34] 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.

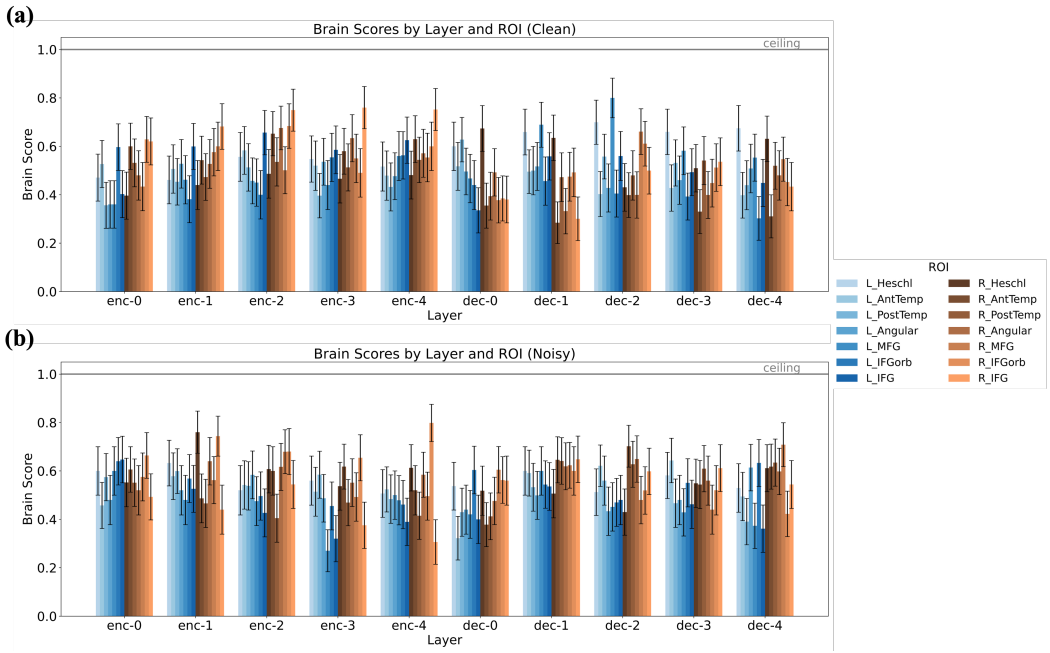

Figure 4: **Layer-wise alignment between Whisper Tiny representations and fMRI responses in 12 language- and auditory-related ROIs, quantified via normalized brain scores (predictivity) under clean and noisy speech. (a)** Clean speech: Alignment patterns across 5 encoder layers (enc-0 to enc-4) and 5 decoder layers (dec-0 to dec-4) for each ROI. **(b)** Noisy speech: Corresponding alignment patterns for the noisy (-3 dB SNR) condition. The horizontal grey line indicates the estimated noise ceiling (normalized to 1), representing the upper bound of model–brain predictivity given measurement noise. Error bars indicate ± 1 SEM across 25 participants.

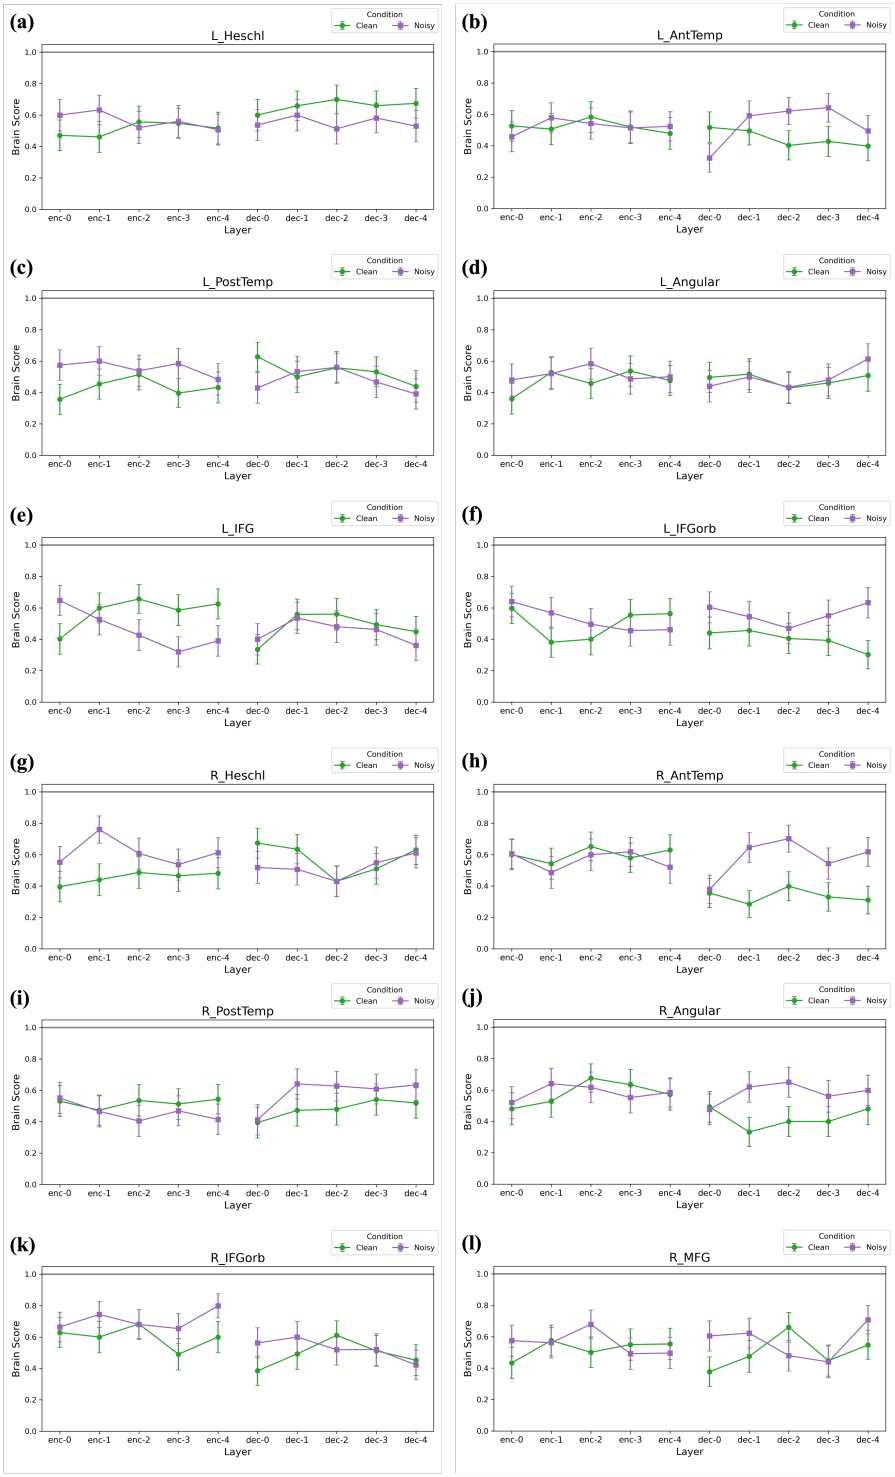

Figure 5: **Model–brain alignment between clean and noisy speech conditions.** Normalized brain scores (noise ceiling = 1.0) are plotted across Whisper Tiny layers for clean (green) and noisy (purple) conditions. The horizontal grey line indicates the noise ceiling. Error bars show ± 1 SEM across 25 participants. Panels (a–l) correspond to: **(a)** Left Heschl, **(b)** Left AntTemp, **(c)** Left PostTemp, **(d)** Left Angular, **(e)** Left IFG, **(f)** Left IFGorb, **(g)** Right Heschl, **(h)** Right AntTemp, **(i)** Right PostTemp, **(j)** Right Angular, **(k)** Right IFGorb, **(l)** Right MFG.