

## Appendices

### A. Detailed Subject-Specific Findings and Discussion

#### A.1 Spatial neural patterns

We first provide a region-level interpretability analysis of BrainStack using a clustered heatmap of normalized expert attribution scores (Figure 6). Each row corresponds to a subject and each column to a predefined cortical region, with hierarchical clustering applied along both dimensions to reveal latent structure in regional reliance.

The results show that the left and right temporal experts consistently receive higher attribution across subjects, underscoring their dominant role in silent speech decoding. This pattern aligns with well-established evidence that the temporal cortex supports phonological and auditory-semantic processing, even in the absence of overt articulation. Parietal and central regions exhibit generally lower attribution, indicating limited involvement in this task. Notably, subjects with stronger reliance on temporal and pre-frontal experts tend to cluster together, suggesting shared internal speech strategies or more robust covert speech representations.

The clustering patterns also reveal clear inter-subject heterogeneity in regional contributions, emphasizing the need for individualized or adaptive decoding strategies. Our findings provide neurophysiologically plausible explanations for BrainStack’s expert routing behavior and further motivate region-aware modeling in silent-speech BCI systems.

#### A.2 Word-level confusion analysis

Figure 7 presents the confusion matrices for all 10 subjects across the 24 word classes. High-performing subjects (e.g., S01, S07, S09) exhibit strong diagonal dominance with minimal off-diagonal errors, indicating confident and well-separated predictions. The sparsity of misclassifications suggests that errors, when present, are isolated rather than systematic.

In contrast, low-performing subjects such as S05 and S06 show highly dispersed confusion patterns, with predictions spread across many incorrect labels. This reflects weak discriminative signal and substantial trial-level variability, consistent with their low accuracy and F1-scores reported earlier.

Some subjects (e.g., S03, S10) show intermediate patterns, with partially formed diagonals and localized confusion clusters. These clusters may correspond to semantically or phonologically related word categories, suggesting class-dependent decoding difficulty. Such observations point toward future directions including semantic-aware decoding architectures or adaptive class grouping strategies.

### B. Additional Experiments: Cross-Subject and Public Dataset

While the main paper focuses on within-subject evaluation on SS-EEG, we additionally assess cross-subject and cross-dataset generalization. Specifically, we report results for: (i) cross-subject decoding using a leave-one-subject-out (LOSO) protocol on SS-EEG, and (ii) evaluation on the public *Thinking Out Loud* dataset. Here we summarize the overall comparative performance; detailed analyses such as ablations, regional contributions, and error breakdowns are provided only for the within-subject SS-EEG setting.

**Cross-subject evaluation.** Table 4 reports results under the leave-one-subject-out (LOSO) protocol on SS-EEG. As expected, accuracies drop substantially compared to the within-subject setting—a well-documented challenge in EEG decoding. This degradation stems from strong inter-individual variability and non-stationarity in EEG, including differences in channel characteristics, noise profiles, and individual-specific linguistic and cognitive patterns, all of which create a pronounced domain shift between training and test subjects.

Despite the low absolute accuracy, BrainStack achieves the best overall performance with an average accuracy of 8.84%, outperforming CNN-based models (EEGNet: 5.88%, TCNet: 6.15%) and Transformer-based models (EEGConformer: 7.19%). These results suggest that the proposed heterogeneous Neuro-MoE architecture with adaptive expert routing can extract a limited but meaningful set of cross-subject invariant features, maintaining a consistent performance advantage even in this challenging regime.

While these findings underscore the difficulty of cross-subject EEG-to-text decoding, they also highlight promising research directions. Reducing the distributional gap between individuals may require domain adaptation, subject-aligned representation learning, or explicit modeling of subject-specific characteristics. Parameter-efficient transfer methods—such as adapter layers—could further enable rapid personalization for unseen subjects. Additionally, multimodal fusion (e.g., EMG, subtle speech-motor signals) may inject complementary priors that help stabilize decoding across participants.

**Thinking Out Loud dataset.** Table 5 reports the within-subject results on the public *Thinking Out Loud* dataset. Overall accuracies are lower than those on SS-EEG, largely due to the dataset’s limited scale and task design: the vocabulary is restricted to four Spanish words, repetitions per class are few, and the total recording duration is short. These factors reduce the model’s ability to learn stable discriminative patterns and increase sensitivity to noise. Cross-linguistic differences—such as the phonological structure

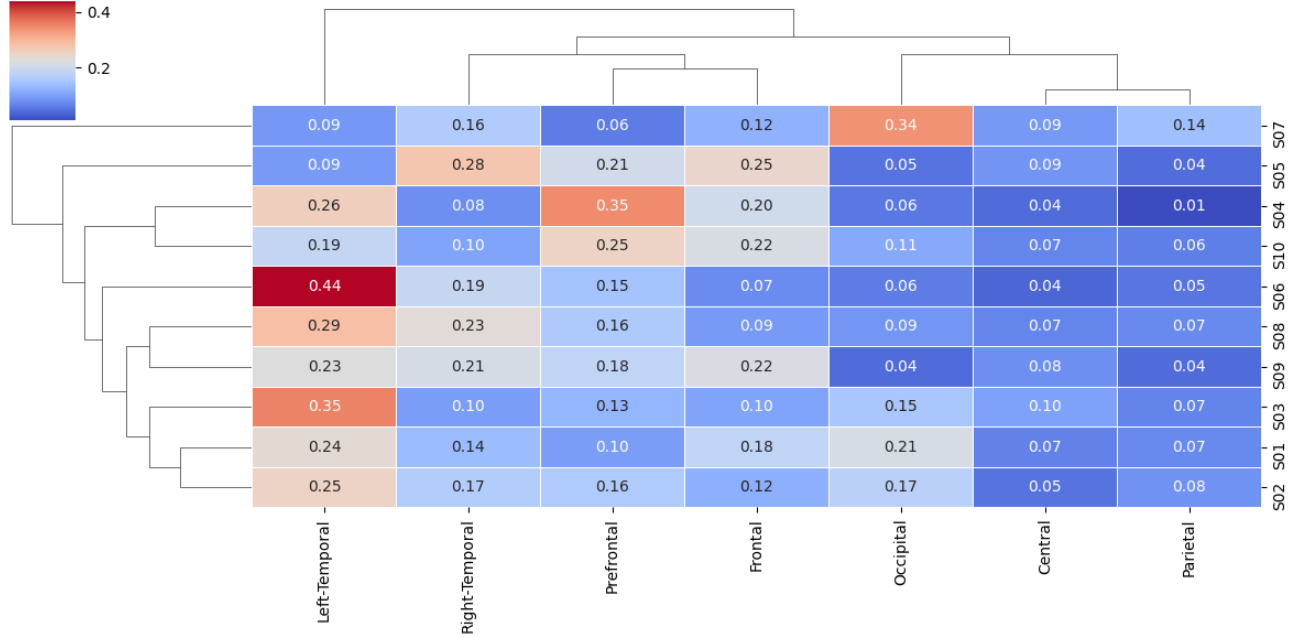


Figure 6. Clustered heatmap of normalized regional contributions across subjects. Both the subject dimension (rows) and the cortical region dimension (columns) are hierarchically clustered to reveal latent structure in expert reliance.

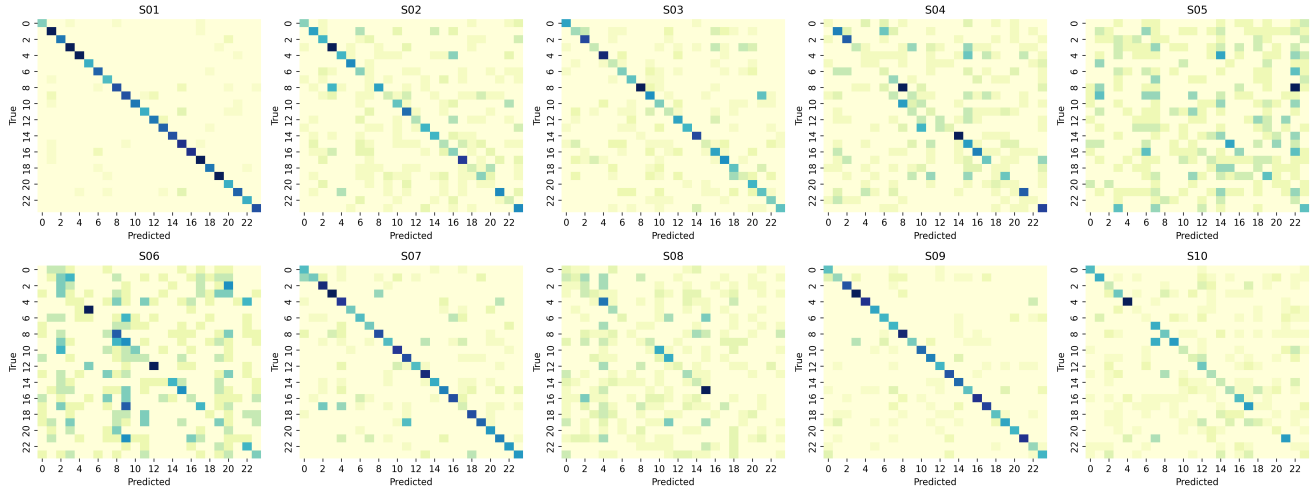


Figure 7. Confusion matrices for all 10 subjects in the SS-EEG dataset. Each matrix reports classification performance across 24 target words, with rows representing ground-truth labels and columns representing model predictions.

874 of Spanish words—may further contribute to inter-subject  
875 variability.

876 Under these constraints, BrainStack achieves the highest  
877 average accuracy of 34.83%, slightly outperforming EEG-  
878 Net (33.11%) and EEGConformer (32.03%), and clearly  
879 surpassing TCNet (30.06%). Although the absolute gains  
880 are modest, the consistent advantage across subjects sug-  
881 gests that the proposed heterogeneous Neuro-MoE with

adaptive expert routing can effectively exploit complemen-  
882 tary local and global features even in small-scale datasets.  
883

884 Taken together, the cross-subject results on SS-EEG and  
885 the within-subject results on *Thinking Out Loud* demon-  
886 strate that BrainStack is not tailored to a single dataset or  
887 evaluation protocol. In the highly challenging cross-subject  
888 setting, where accuracies drop due to substantial inter-  
889 individual variability, BrainStack still outperforms repre-

Table 4. Performance comparison on the SS-EEG dataset under the cross-subject setting (leave-one-subject-out).

Model	S01	S02	S03	S04	S05	S06	S07	S08	S09	S10	Avg. Acc.
EEGNet	10.83%	7.52%	5.05%	5.00%	4.78%	5.20%	4.52%	4.82%	5.75%	5.30%	5.88%
TCNet	5.97%	6.50%	6.07%	4.89%	5.62%	7.13%	6.25%	4.55%	7.35%	7.18%	6.15%
EEGConformer	6.91%	7.55%	8.27%	5.10%	5.53%	7.11%	10.28%	4.73%	9.05%	7.33%	7.19%
<b>BrainStack</b>	<b>13.42%</b>	<b>8.27%</b>	<b>11.25%</b>	<b>5.13%</b>	<b>6.60%</b>	<b>8.05%</b>	<b>11.97%</b>	<b>5.10%</b>	<b>10.18%</b>	<b>8.41%</b>	<b>8.84%</b>

Table 5. Performance comparison on the Thinking Out Loud dataset.

Model	S01	S02	S03	S04	S05	S06	S07	S08	S09	S10	Avg. Acc.
EEGNet	33.33%	30.56%	36.67%	<b>36.11%</b>	30.56%	33.33%	26.67%	40.00%	33.33%	30.56%	33.11%
TCNet	26.67%	<b>36.11%</b>	36.67%	27.78%	27.78%	27.27%	27.78%	26.67%	30.56%	33.33%	30.06%
EEGConformer	40.00%	27.78%	30.00%	30.56%	27.78%	30.30%	27.78%	36.67%	36.11%	33.33%	32.03%
<b>BrainStack</b>	<b>40.00%</b>	30.56%	<b>36.67%</b>	27.78%	<b>33.33%</b>	<b>40.00%</b>	<b>30.56%</b>	<b>40.00%</b>	<b>36.11%</b>	<b>33.33%</b>	<b>34.83%</b>

sentative CNN- and Transformer-based baselines, indicating its ability to capture partially transferable neural representations. On the public dataset—despite its limited vocabulary, short duration, and linguistic differences from SS-EEG—BrainStack again achieves the best overall performance. The convergence of these supporting results demonstrates that the proposed functionally guided NeuroMoE architecture exhibits cross-domain efficacy beyond SS-EEG, thereby establishing it as a versatile methodological framework for advancing broader EEG-to-text decoding paradigms.

### C. The SS-EEG Dataset

The SilentSpeech-EEG (SS-EEG) dataset was developed to support non-invasive silent speech decoding research based on electroencephalography (EEG). The dataset comprises over 120 hours of EEG recordings from 12 right-handed native English speakers.

#### C.1 Participants

- **Total Subjects:** 12 right-handed native English speakers (6 male, 6 female)
- **Age:** Mean = 28.0 years, SD = 0.85 years
- **Health Status:** No reported visual, speech, or neurological disorders
- **Recruitment and Consent:** All participants were recruited in Australia and provided written informed consent
- **Ethics Approval:** Approved by the IRB ethics committee

#### C.2 Experimental environment and task design

- **Setup:** Participants sat in a quiet lab, facing a monitor at a comfortable distance

- **Instructions:** Subjects kept their eyes open and minimized body movement
- **Task:** Silent articulation (no sound or subvocalization) of prompted words
- **Vocabulary:** 24 words from 6 semantic categories: *motion, emotion, location, person, object, number*
- **Trial Design:** Each word repeated 5 times per session; randomized order
- **Implementation:** Task presented via PsychoPy (Python)

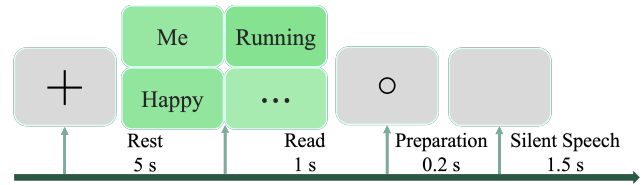


Figure 8. Schematic of a single trial in the silent speech EEG (SS-EEG) experiment. Each trial consists of four phases: a rest period (5 s) with fixation, word display (1 s), an auditory cue (0.2 s), and the silent articulation phase (1.5 s).

#### C.3 Experimental pipeline

- **Total Duration:** Over 120 hours of EEG data
- **Sessions:** Each subject completed 16 sessions, totaling 6,000 trials per participant
- **Session Length:** Each session includes 375 trials and lasts approximately 40 minutes (including breaks)
- **Trial Phases:** Each trial follows a four-stage protocol:
  - *Rest (5 s)* — includes a 1.5 s fixation cross
  - *Read (1 s)* — target word displayed on screen
  - *Preparation (0.2 s)* — auditory cue signaling onset
  - *Silent Speech (1.5 s)* — subject imagines articulating the word

- **Fatigue Mitigation:** 5-minute break every 20 trials
- **Vocabulary:** 24 English words across 6 semantic groups (e.g., “go”, “happy”, “python”, “running”)
- **Final Dataset:** Due to severe signal contamination, 2 subjects were excluded. Final release includes data from 10 subjects (Subject-01 to Subject-10)

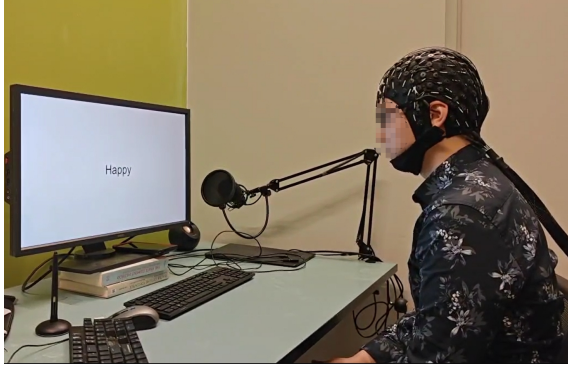


Figure 9. Experimental setup for silent speech EEG recording. A participant silently articulates the displayed word while wearing a 128-channel Neuroscan EEG cap. Recordings were performed in a controlled, sound-attenuated lab environment.

#### C.4 Recording setup

- **EEG Cap:** 128-channel Neuroscan Quik-Cap®
- **Amplifier:** SynAmps RT 128-channel amplifier
- **Software:** CURRY 9.0 for acquisition
- **Montage:** 133 channels total:
  - 122 EEG channels (brain activity)
  - 4 EOG electrodes (eye movement)
  - 6 reference electrodes (environment/muscle noise)
  - 1 trigger channel (condition labeling)
- **Electrode Placement:** 10–10 international system
- **Impedance:** Maintained below 10 k $\Omega$  and balanced across channels

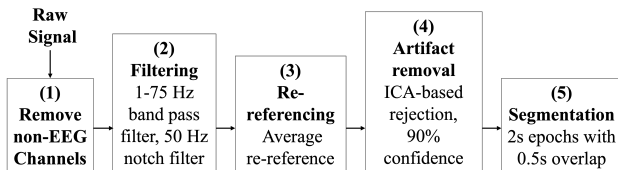


Figure 10. The pipeline of raw EEG data preprocessing.

#### C.5 Preprocessing

- **Filtering:** 1–75 Hz bandpass (FIR), 50 Hz notch filter for line noise
- **Referencing:** Common average across 122 EEG channels
- **Artifact Removal:** ICA-based rejection of eye, muscle, and cardiac noise; 90% confidence threshold

- **Segmentation:** Epochs of 2 s length with 0.5 s overlap
- **Software:** EEGLAB toolbox (MATLAB)

#### (a) Read v.s. Silent Speech

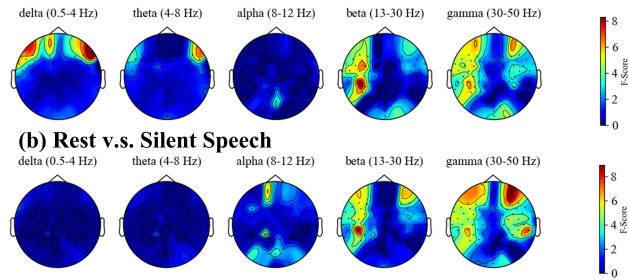


Figure 11. Topographic visualization of F-scores from repeated-measures ANOVA on five canonical EEG frequency bands. (a) Read vs. Silent Speech; (b) Rest vs. Silent Speech. Strong discriminative patterns emerge in the alpha, beta, and gamma bands, particularly in left-lateralized regions.

#### C.6 Frequency band analysis

To investigate which neural frequency bands are most discriminative for silent speech decoding, we performed a repeated-measures analysis of variance (rm-ANOVA) on the power spectral density (PSD) of EEG signals across five canonical frequency bands: delta (1–4 Hz), theta (4–8 Hz), alpha (8–13 Hz), beta (13–30 Hz), and gamma (30–50 Hz). The analysis compared three condition pairs: rest versus reading, rest versus silent speech, and reading versus silent speech.

F-scores from the ANOVA were used to quantify inter-condition differences, with higher values indicating greater spectral variation. In the delta band, elevated frontal activation was observed, likely reflecting ocular artifacts. The alpha band exhibited moderate but consistent differences in occipital regions, particularly near the POOZ electrode, suggesting changes in visual attention. Notably, the beta and gamma bands showed the most significant activation differences, especially in the left hemisphere. These variations are potentially associated with motor planning and cognitive control mechanisms specific to internal speech generation.

Further comparison between rest and silent speech confirmed that occipital activity diminished, while fronto-temporal beta/gamma activation increased, reinforcing the hypothesis of reduced visual processing and elevated cognitive engagement during silent articulation. These findings highlight alpha, beta, and gamma bands as the most informative for discriminating silent speech states, motivating their selection in our downstream modeling and multi-band fusion analyses.

Table 6. Summary of SS-EEG dataset configuration and vocabulary stimuli with syllable counts.

Data Collection Settings	Stimuli and Semantic Groups
Subjects: 12	<b>Motion:</b> Jumping (2), Running (2), Swimming (2), Going (2)
Modality: EEG	<b>Emotion:</b> Happy (2), Sad (1), Fun (1), Horrible (3), College (3)
EEG Channels: 122	<b>Location:</b> Home (2), Battlefield (3), Here (2)
Sessions / Subject: 16	<b>People:</b> Mother (2), Cowboy (2), Professor (3), Me (1)
Trials / Subject: 6000	<b>Number:</b> One (1), Three (1), Eleven (3), Million (2)
Vocabulary Size: 24	<b>Object:</b> Spoon (1), Alfa (2), Python (2), Telephone (3)
Total Time / Subject: 10 hrs	

## D. Cortical Parcellation

### D.1 Anatomical definitions of the 7 RoIs

To incorporate neuroscientific priors into our EEG modeling framework, we partition the scalp-level EEG channels into anatomically informed functional regions of interest (ROIs). These partitions are based on the standard 10–10 international system and follow spatially adjacent topographies consistent with prior literature on cortical function and EEG spatial layout.

We define seven anatomically informed regions of interest (ROIs) based on standard EEG electrode locations. The detailed channel assignments for each region are listed in Table 7.

### D.2 Simplified 5-region variant (RoI.5)

For comparison purposes, we also introduce a coarse-grained 5-region mapping by merging Prefrontal and Frontal, and Left/Right Temporal into one Temporal group:

- **Frontal (merged):** Frontal\_front + Frontal\_middle\_and\_rear (combined 23 channels)
- **Central, Parietal, Occipital:** Same as RoI.7
- **Temporal:** Combined left and right temporal regions (14 channels)

This version, referred to as BRAINSTACK\_ROI5, is used in ablation studies to evaluate the importance of fine-grained spatial decomposition.

## E. Model Architecture and Training Configuration

BrainStack integrates global and region-specific neural representations through a heterogeneous Neuro-Mixture-of-Experts architecture. The model comprises: (a) seven regional experts based on CNet, each operating on a distinct anatomically defined subset of EEG channels; (b) one global expert implemented with a CTNet encoder over the full channel montage; and (c) an adaptive expert routing gate, which combines expert outputs through a learnable attention mechanism.

The input is a segment of EEG signals of shape `(batch.size, 122, 1000)`.

### E.1 Global and regional architectures

**CTNet (Global Expert).** The global expert operates on the full EEG montage and consists of:

- Patch Embedding: Conv2D + AvgPool  $\rightarrow$  40-d embeddings.
- Transformer Encoder: 2 layers, 4 heads, dropout=0.5.
- Fully Connected: Flatten  $\rightarrow$  32-d hidden vector.
- Final Classifier: Linear(32  $\rightarrow$  24).

**CNet (Regional Experts).** Each regional expert uses a lightweight convolutional architecture tailored for localized EEG processing. The full configuration is provided in Table 8.

### E.2 Adaptive Expert Routing

We implement an adaptive expert routing mechanism to combine global and regional experts:

- Each expert output (32-d) is first projected into a shared 224-d routing space.
- A learnable scoring vector is applied to all expert embeddings and normalized via a softmax with temperature  $T = 1.5$  to obtain routing weights.
- DropBranch regularization (drop ratio = 0.25) is applied during training to improve robustness and prevent expert co-adaptation.
- The routed feature is then processed by a lightweight MLP:  
MLP: Linear(224  $\rightarrow$  512)  $\rightarrow$  ReLU  $\rightarrow$  Linear(512  $\rightarrow$  24)

### E.3 Training configuration

See table 9.

## F. Limitations and Future Work

While our results show that functionally guided region parcellation and heterogeneous expert modeling substantially improve EEG-based silent speech decoding, several limitations remain.

Table 7. EEG channel assignments for 7-region cortical division (RoI-7).

RoI	Region Name	#Channel	Assigned EEG Channels
RoI.1	Frontal_front	8	AF3, FP1, FPz, AFz, AF7, AF4, Fp2, AF8
RoI.2	Frontal_middle_and_rear	15	FFT7h, F5, F7, AFF5h, F3, FFC1h, F1, Fz, FFC2h, F2, AFF6h, F4, F6, F8, FFT8h
RoI.3	Central	39	FC5, FFC5h, FFC3h, FC1, FCz, FCC5h, FC3, FCC3h, FCC1h, FCCz, FC2, FFC4h, FFC6h, FC6, FCC2h, FCC4h, FC4, FCC6h, C5, CCP5h, C3, C1, CCP1h, Cz, CCP2h, C2, C4, CCP6h, C6, TTP7h, TPP5h, CP3, CCP3h, CP1, CP2, CCP4h, CP4, TPP8h, TTP8h
RoI.4	Left-Temporal	7	T9, FT9, FTT9H, T7, TP7, FTT7H, FT7
RoI.5	Right-Temporal	7	FTT8H, FT8, T10, FT10, FTT10H, T8, TP8
RoI.6	Parietal	18	P7, CPP5H, CPP3H, CPP1H, CPPZ, P9, P5, P3, PPO1, PPOZ, CPP2H, CPP4H, CPP6H, P8, PPO2, P4, P6, P10
RoI.7	Occipital	28	P11, PO11, PO9, PPO7, PO3, POO7, POO9H, POO11H, I1, OI1, POO3, PO1, POZ, POOZ, OZ, IZ, P12, PO12, PO10, PPO8, PO4, POO8, POO10H, POO12H, I2, OI2, POO4, PO2

Table 8. Configuration of each CNet expert branch.

Module	Configuration
Temporal Conv	$1 \times 32$ kernel, $f_1 = 16$ , padding=same
Spatial Conv	Depthwise, $d = 2$ , grouped by $f_1$
Separable Conv	$1 \times 16$ kernel, $f_2 = 32$
Pooling	AvgPool ( $1 \times 8$ ), ( $1 \times 16$ ), Dropout(0.5)
Classifier	Linear( $\cdot \rightarrow 24$ ), max-norm

Table 9. Training configuration

Setting	Value
Learning Rate	5e-3
Weight Decay	1e-4
Max Epochs	100
Early Stop Patience	5 Epochs
Optimizer	SGD
Scheduler	ReduceLROnPlateau
Batch Size	16
Num Classes	24
Warm-up Epoch $T_{\text{warmup}}$	5
Transition Epoch $T_{\text{transition}}$	5
$\lambda$	$0.2 \sim 1$
$\alpha$	$0.8 \sim 0$
$\beta_{\text{max}}$	0.5
$\gamma_{\text{max}}$	1
Validation Split	Random Selection
Total Parameters	1.06M
Device	NVIDIA A40 (48GB)

First, our use of a predefined 7-region cortical mapping—derived from standard EEG montages and prior neuroscience—may not represent the optimal granularity for language decoding. More flexible strategies such as data-driven parcellation, task-dependent functional clustering, or subject-adaptive region selection may yield expert decompositions that better align with underlying neural organization.

Second, although BrainStack improves robustness, significant inter-subject variability persists, especially for subjects with low signal-to-noise ratios. This indicates that static anatomical partitioning cannot fully capture individual neural differences. Future extensions may incorporate personalized expert assignments, dynamic region re-weighting, or meta-learning-based adaptation to further enhance cross-participant generalization.

Finally, our current design primarily leverages anatomical priors, while other forms of domain knowledge—such as the temporal hierarchy of language processing, hemi-

spheric asymmetry, and neurosemantic pathways—remain underutilized. Embedding these priors through structured inductive biases, temporal-semantic expert modules, or neuro-informed routing strategies represents a promising

Table 10. Parameter breakdown of the BrainStack Neuro–MoE components.

Component	# Parameters	Notes
Each CNet (×7)	8.4K–13.3K	Varies by regional input size (total ~84K)
Global CTNet	~871K	Includes Patch Embedding, 2-layer Transformer, FC, Output
Feature Projector	7.4K	Linear(32 → 224) projection
Expert Routing Network	127.5K	Linear(224→512→24)
<b>Total</b>	<b>1.06M</b>	Trainable parameters in full model

1094 direction for advancing EEG-to-text decoding.

1095 **G. Ethics**

1096 This study was conducted under strict adherence to ethical  
1097 standards for research involving human participants. All  
1098 procedures were approved prior to data collection, and mea-  
1099 sures were taken to ensure participant safety, privacy, and  
1100 informed participation.

1101 **1.Ethics Approval and Participant Consent** The ex-  
1102 perimental protocol was reviewed and approved by the In-  
1103 stitutional Human Research Ethics Committee at our insti-  
1104 tution. All participants were recruited voluntarily and pro-  
1105 vided written informed consent. Prior to participation, sub-  
1106 jects were fully briefed on the nature of the silent speech  
1107 EEG experiment, including the tasks, session duration, and  
1108 the intended research goals. Participants were explicitly in-  
1109 formed of their right to withdraw at any time without con-  
1110 sequence.

1111 **2.Data Privacy and Anonymization** No personally  
1112 identifiable information (PII) was collected. All EEG data  
1113 were anonymized upon recording and stored on encrypted,  
1114 access-controlled servers. Data usage was strictly limited  
1115 to academic research purposes. The anonymization process  
1116 ensures that individual identities cannot be recovered, pro-  
1117 tecting the privacy of participants throughout all stages of  
1118 data handling and analysis.

1119 **3.Physical Safety and Intended Use** The experimen-  
1120 tal procedure involved only non-invasive EEG measure-  
1121 ments and silent articulation tasks, posing minimal phys-  
1122 ical or psychological risk. Participants were given suffi-  
1123 cient breaks to avoid fatigue. This work is intended to  
1124 advance assistive communication and neurotechnology re-  
1125 search. The dataset and models are not designed for in-  
1126 ferring internal cognitive states without consent. Any fu-  
1127 ture application of this technology must adhere to estab-  
1128 lished ethical norms regarding transparency, privacy, and  
1129 non-coercive use.

**H. Broader Impacts** 1130

**H.1 Positive Impacts** 1131

**Assistive Communication for Speech-Impaired Users** 1132  
Our EEG-based silent speech decoding framework offers  
1133 a non-invasive pathway for individuals with severe speech  
1134 impairments to express themselves. This technology can  
1135 serve as a foundation for future communication aids, en-  
1136 abling users to silently articulate intended words and im-  
1137 proving quality of life, autonomy, and social participation. 1138

**Neurotechnological Advancement** By linking local-  
1139 ized cortical activity to linguistic representations, our work  
1140 contributes to the broader understanding of neural language  
1141 processing. It provides empirical evidence for the role of  
1142 anatomically defined brain regions in internal speech de-  
1143 coding, which may inform clinical diagnostics and cogni-  
1144 tive neuroscience research. 1145

**Scalable and Accessible BCI Design** BrainStack is  
1146 lightweight and compatible with widely available EEG  
1147 hardware, lowering the barrier for practical deployment. Its  
1148 modular Neuro–MoE architecture further enables future ex-  
1149 tensions—such as adaptive personalization or multilingual  
1150 decoding—broadening accessibility to diverse user popula-  
1151 tions and application domains. 1152

**H.2 Negative Impacts** 1153

**Privacy, Consent, and Misuse Risk** Decoding neural  
1154 signals—especially internal speech—raises significant con-  
1155 cerns around user privacy, informed consent, and the po-  
1156 tential for surveillance. Improper collection, storage, or in-  
1157 terpretation of EEG data could unintentionally reveal sensi-  
1158 tive cognitive states or user identities. Moreover, BCI sys-  
1159 tems developed for assistive purposes could be repurposed  
1160 for covert behavioral inference or unauthorized monitoring.  
1161 While our current system is limited to word-level decoding  
1162 under controlled conditions, future models with higher res-  
1163 olution and broader vocabulary coverage may amplify these  
1164 risks. Ethical deployment requires robust data governance,  
1165 transparency, and strict usage constraints to prevent misuse. 1166

**Bias and Representational Limitations** Our dataset  
1167 is derived from 12 healthy, right-handed English speakers, 1168

1169 limiting its demographic and linguistic diversity. This raises  
1170 concerns about potential performance disparities when ap-  
1171 plied to broader populations, including users with speech  
1172 or neurological impairments, non-native speakers, or those  
1173 with atypical cortical patterns. Without inclusive data col-  
1174 lection and fairness-aware training objectives, BCI systems  
1175 risk encoding existing biases or failing to generalize across  
1176 user groups. Future work should prioritize equitable repre-  
1177 sentation, demographic balancing, and bias auditing in both  
1178 data and modeling.