

000 001 002 003 004 005 SELF-SUPERVISED LEARNING WITH SPATIALITY PRE- 006 SERVING REPRESENTATION FOR EEG SIGNALS 007 008 009

010 **Anonymous authors**
011 Paper under double-blind review
012
013
014
015
016
017
018
019
020
021
022
023
024
025
026
027
028
029
030
031
032
033
034
035
036
037
038
039
040
041
042
043
044
045
046
047
048
049
050
051
052
053

ABSTRACT

Self-supervised learning (SSL) has revolutionized the field of deep learning with EEG signals, yet current approaches face a critical limitation: the loss of crucial spatial information due to architectures that fail to adequately preserve the one-to-one electrode relationship between the input and representation. To address this, we introduce Spatiality Preserving Representation (SPR) Learning. Unlike existing methods relying on reconstruction or temporal prediction with separate encoders, SPR learns spatial relationships through an innovative coherence pseudolabel prediction task, teaching models to understand the intricate topographical organization of brain signals that conventional approaches overlook. Through comprehensive evaluation, SPR demonstrates superior performance over state-of-the-art methods (4.7%, 9.7%, 1.6%, and 15.6% fine-tuning improvements over different datasets), learning meaningful spatial representations that capture the complex spatial-temporal dynamics inherent in EEG data. Our work opens new avenues for interpreting the relationships of different brain regions by prioritizing spatial awareness, and thus bridge the gap between functional connectivity analysis and self-supervised EEG representation learning.

1 INTRODUCTION

Electroencephalography (EEG) is a noninvasive neuroimaging technique that captures brain activity by recording electrical signals through electrodes placed on the scalp (Teplan et al., 2002). Compared to other neuroimaging modalities such as functional magnetic resonance imaging (fMRI) and magnetic resonance imaging (MRI), EEG offers greater affordability and accessibility for early disease detection (Bosl et al., 2018; Parmar & Paunwala, 2023; Ehteshamzad, 2024) and serves as particularly valuable diagnostic tools (Jeong, 2004; Britton et al., 2016). Consequently, machine learning (ML) approaches applied to EEG analysis have gained significant research attention in recent years (Rafiei et al., 2022; Hosseini et al., 2021).

Self-supervised learning (SSL) is an ML paradigm that leverages unlabeled data to extract inherent features within datasets (Chen et al., 2020; Chen & He, 2021; Zbontar et al., 2021). Given that supervised learning requires substantial amounts of annotated data, and medical data collection and annotation can be prohibitively time-consuming and expensive, SSL presents an attractive alternative. Contemporary SSL methods for EEG signals typically employ pretext tasks with separate encoders to generate learning objectives (Liu et al., 2023; Guetschel et al., 2024; Weng et al., 2025). Invariance-based methods aim to optimize encoders that produce similar embeddings for multiple views of the same EEG time series (Eldele et al., 2021; Mohsenvand et al., 2020), while self-prediction methods involve masking or corrupting parts of the input data and training models through self-prediction or reconstruction (Banville et al., 2021; Cui et al., 2024). Nonetheless, these methods tend to overlook spatial relationships between EEG electrodes, and our framework is motivated to design a channel-aware pretext task based on spatial coherence predictions.

EEG oscillations within specific frequency bands often reflect similar underlying physiological processes. Although many previous models did not explicitly select frequency bands, their architectures implicitly encode preferences for certain frequency ranges (Abello et al., 2021; Molina et al., 2024). This observation motivates our use of both a band-specific approach and a band-mixture design, particularly across higher frequencies. Higher-frequency components exhibit more localized and specific connectivity patterns that quickly attenuate with distance (Łęski et al., 2013; Amo et al.,

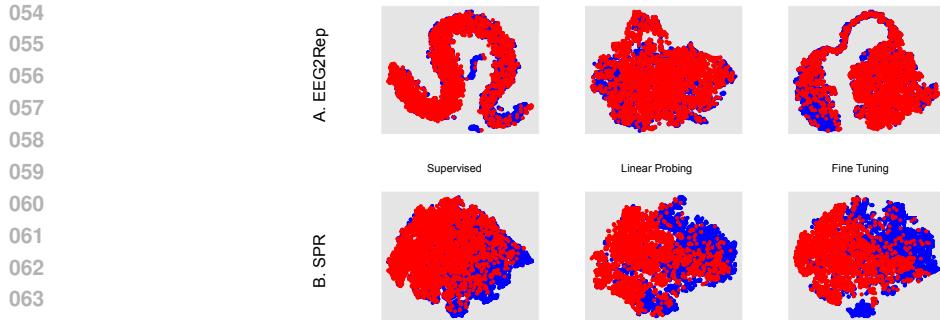


Figure 1: Comparison of t-SNE visualizations between EEG2Rep and SPR (our method) on the STEW dataset (Lim et al., 2018a). Blue and red points represent the two classes: low workload and high workload conditions, respectively. **(A)** EEG2Rep representations (top row). **(B)** SPR representations (bottom row). Our method demonstrates superior class separability with more distinct clustering patterns. Data samples were preprocessed as detailed in Appendix A.1.

2017; Fitzgerald & Watson, 2018; Arnulfo et al., 2020), whereas lower frequencies spread more broadly and are visible in large areas of the scalp. In particular, coherence decays with distance, and higher frequencies show steeper falloffs, indicating that higher-frequency activities are more local while lower-frequency activities support longer-range interactions (Jia et al., 2011; Buzsáki & Schomburg, 2015). In our pre-training, we focus on shorter-range interactions and aim to capture informative local spatial relationships by teaching the model the coherence between channel pairs in higher-frequency bands.

To support our design choices, we present t-distributed stochastic neighbor embedding (t-SNE) visualizations in Fig. 1. The visualization clearly demonstrates that our method generates more visually separable representations compared to the state-of-the-art (SOTA) EEG2Rep (Mohammadi Foumani et al., 2024) approach, indicating the creation of more meaningful and robust representations through our SSL framework. Notably, fine-tuning with our framework also exhibits superior class separability compared to supervised methods, demonstrating the enhanced performance achieved through SSL.

In summary, our contributions are as follows: **(A)** We propose the *Spatiality Preserving Representation (SPR) Learning* framework. EEG signals provide crucial input for early disease detection, and our SPR representation captures the one-to-one electrode relationship useful to functional connectivity analysis and disease localization. **(B)** Our SSL pretext task introduces a novel *band-mixture coherence approach* that captures a meaningful spatial structure of the electrode arrangements of EEG. This pretext task also does not require the keeping of separate encoders, fundamentally differing from existing invariance-based and self-prediction methods. **(C)** Despite employing a compact model architecture, our method consistently achieves *SOTA performance* through straightforward in-domain pre-training (4.7%, 9.7%, 1.6%, and 15.6% fine-tuning improvements over different datasets), demonstrating promising directions for future research.

2 RELATED WORK

SSL methods in EEG analysis offer substantial advantages in addressing the persistent challenges associated with acquiring high-quality labeled EEG datasets, and the field has recently witnessed significant advancement. BENDR learns representations of raw data signals through reconstruction (Kostas et al., 2021). MAEGG learns to reconstruct the masked EEG features using a transformer architecture (Chien et al., 2022). TS-TCC transforms the data into two different views and uses temporal contrasting to learn robust temporal features (Eldele et al., 2021). TF-C embeds a time-based neighborhood of an example close to its frequency-based neighborhood (Zhang et al., 2022). BIOT is a cross-data learning framework that can be used with EEG signals and SSL by tokenizing the signals into unified sentences (Yang et al., 2023). TFM-Tokenizer transforms EEG signals into a sequence of discrete, well-represented tokens in capturing critical time-frequency features from a single channel (Pradeepkumar et al., 2025). EEG2Rep learns to predict masked input in latent

108 representation space (Mohammadi Foumani et al., 2024). Compared to those methods, SPR does
 109 not contrast different views of the input signal or keep separate encoders, but rather adopts novel
 110 spatial coherence targets that can be precomputed efficiently.

111 Foundation models represent large-scale, pre-trained neural network architectures that serve as ver-
 112 satile backbone systems for a wide array of downstream tasks. Motivated by significant advances
 113 in natural language processing (NLP) and computer vision (CV) (Caron et al., 2021; Radford et al.,
 114 2021; Bao et al., 2022; Peng et al., 2022; Wang et al., 2023), foundation models have attracted
 115 considerable interest in recent years. Recently, LaBraM and NeuroLM have used vector-quantized
 116 encoders that establish strong conceptual and methodological connections to their CV counterparts
 117 (Jiang et al., 2024; 2025), demonstrating the transferability of foundation model concepts across
 118 different data modalities. Other notable works include EEGFormer (Chen et al., 2024), EEGM2
 119 (Hong et al., 2025), LEAD (Wang et al., 2025), Neuro-GPT (Cui et al., 2024), EEGPT (Wang et al.,
 120 2024a), CBraMod (Wang et al., 2024b). In particular, our SPR model is not similar to those cross-
 121 domain foundation models, but rather to preserve and capture useful spatial representation through
 122 in-domain pre-training.

123 3 METHOD

124 Our study introduces a novel approach for generating coherence-based pseudolabels from EEG data,
 125 which involves computing band-limited magnitude-squared coherence across all possible channel
 126 pairs. This process allows for the extraction of frequency-specific connectivity patterns within pre-
 127 defined EEG frequency bands, which are subsequently converted into mixed probabilistic pseudola-
 128 bels for downstream machine learning tasks.

131 3.1 BAND-LIMITED SPECTRAL ANALYSIS

132 To focus on physiologically relevant oscillatory interactions, we employ a band-limited spectral
 133 analysis rather than computing coherence across the entire frequency spectrum. This is akin to a
 134 periodogram with averaging but with an important distinction: it performs frequency-bin averaging
 135 within a specific band, unlike the full-spectrum approach. The process begins by applying the Fast
 136 Fourier Transform (FFT) to the input time series \mathbf{x} to obtain its frequency-domain representation:

$$139 \quad \mathbf{X}_{fft} = \text{FFT}(\mathbf{x}) \quad (1)$$

140 Frequency bins corresponding to the target band were selected using a frequency mask:

$$143 \quad M(f) = \begin{cases} 1 & \text{if } f_{low} \leq |f| \leq f_{high} \\ 0 & \text{otherwise} \end{cases} \quad (2)$$

144 Here, f_{low} and f_{high} define the boundaries of the frequency band, and f is the frequency axis,
 145 computed as $f = \frac{k \cdot f_s}{N}$ for $k = 0, 1, \dots, N - 1$, where f_s is the sampling frequency and N is the
 146 number of samples.

147 Our band-limited approach differs significantly from the traditional Welch's method for spectral
 148 density estimation, as summarized in Table 1. Welch's method enhances variance reduction by
 149 dividing the signal into overlapping segments, applying a windowing function (e.g., Hanning) to
 150 each segment, and then averaging the periodograms across these segments. In contrast, our approach
 151 processes the signal as a single time window and later uses a band-mixture approach to trade a small
 152 bias for reduced variance. More details can be found in Appendix C.

155 3.2 COHERENCE PSEUDOLABELS

156 Cross-power spectral density (CPSD) is a complex-valued measure that quantifies the statistical
 157 relationship between two signals or two channels of a multi-channel signal in the frequency domain.
 158 The magnitude of the CPSD at a given frequency indicates the shared power between the two signals,
 159 and a high magnitude signifies a strong correlation of power at that specific frequency. Meanwhile,
 160 the phase of the CPSD at a given frequency represents the phase difference or time lag between the
 161 two signals' components.

162
163
164
165
166
167
168
169
170
171
172
173
174
175
176
177
178
179
180
181
182
183
184
185
186
187
188
189
190
191
192
193
194
195
196
197
198
199
200
201
202
203
204
205
206
207
208
209
210
211
212
213
214
215

Table 1: Comparison of estimation methods

	Traditional Welch's	Band-limited
Segmentation	Overlapping time windows	Single time window
Averaging	Across time segments	Across frequency bins
Resolution	Time-frequency trade-off	Fixed by band selection
Computation	Higher (multiple segments)	Lower (single FFT)
Specificity	Full spectrum	Band-specific

We compute the CPSD for all possible pairs of channels from the input EEG data. For any two channels, i and j , the CPSD, $S_{ij}(f)$, is defined as:

$$S_{ij}(f) = X_i(f) \cdot X_j^*(f) \quad (3)$$

This resulted in a matrix of dimensions $\mathbf{S} \in \mathbb{C}^{C \times C \times N_{freq}}$, where C is the number of channels. Each element is a complex number representing the CPSD between channel i and channel j .

For our coherence calculation, we take the average of these complex values across the frequency bins within the specified band. Each element in the resulting matrix $\tilde{\mathbf{S}} \in \mathbb{C}^{C \times C}$ entails the average of shared power and phase difference between channel i and j over the entire frequency band.

The diagonal elements of $\tilde{\mathbf{S}}$, where $i = j$, represent the average power spectral density (PSD) for each channel. By extracting the diagonal from $\tilde{\mathbf{S}}$, we obtain the values corresponding to the total average power of a specific channel across the entire frequency band. The magnitude-squared coherence, γ_{ij}^2 , is thus defined by:

$$\gamma_{ij}^2 = \frac{|\tilde{S}_{ij}|^2}{\tilde{S}_{ii} \cdot \tilde{S}_{jj}} \quad (4)$$

Here, the numerator corresponds to the squared magnitude of the average CPSD, and the denominator is the product of the average PSDs of channel i and j , effectively serving as a normalization factor. This approach yields a single coherence value per channel pair for each frequency band, capturing the overall connectivity strength within that physiological frequency range, as illustrated in Fig. 2.

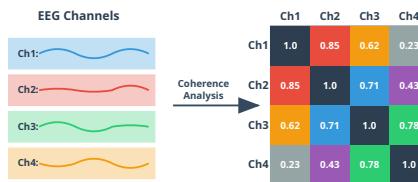


Figure 2: An example of the coherence matrix for a four-channel EEG signal. Each element represents the normalized connectivity between channel i and channel j .

To convert the derived coherence matrices into probabilistic pseudolabels for our downstream learning task, we apply softmax normalization with a temperature parameter τ to control the sharpness of the probability distribution. The pseudolabel for each channel pair is computed as:

$$\mathbf{L}_{ij} = \frac{\exp(\gamma_{ij}^2 / \tau)}{\sum_{k=1}^C \exp(\gamma_{ik}^2 / \tau)} \quad (5)$$

The resulting pseudolabels have dimensions $C \times C$, where each element represents the normalized pseudolabel for a specific channel pair. This provides a channel-wise probability distribution over all possible connections for the frequency band, which itself can be used as a soft target.

216
217

3.3 BAND MIXTURE

218 To capture a comprehensive representation of higher-frequency connectivity, we aggregate the pseudolabels from the α (8-13 Hz), β (13-30 Hz) and γ (30-100 Hz) bands, excluding lower-frequency
 219 δ (0.5-4 Hz) and θ (4-8 Hz) bands. The α , β and γ oscillations are robustly associated with active
 220 cognitive processes, conscious perception, and local or inter-regional communication during wakeful states (Başar et al., 2001; Newson & Thiagarajan, 2019). In particular, α rhythms are linked to
 221 attentional states and sensorimotor integration, β to motor control and higher-level cognitive pro-
 222 cessing, and γ to feature binding and conscious awareness.
 223

225 At the center of our interest are spatial relationships. Higher frequency oscillations typically exhibit
 226 more spatially localized patterns compared to the widespread, diffuse nature of δ and θ rhythms
 227 (Yang et al., 2020). This property makes α , β and γ bands more suitable for distinguishing fine-
 228 grained connectivity patterns between specific channel pairs. The chosen higher-frequency oscil-
 229 lations are also less affected by artifacts, whereas lower frequency bands are highly susceptible to
 230 physiological artifacts, such as eye movements and muscle activity, which can confound connectiv-
 231 ity analysis.
 232

233 For our method, the final coherence pseudolabel \mathbf{L}_{ij} is computed as the average of the pseudolabels
 234 from the α , β , and γ bands:

$$235 \quad \mathbf{L}_{ij} = \frac{1}{N_b} \sum_{k \in \mathcal{B}} \mathbf{L}_{ij,k} \quad (6)$$

236 Here, $\mathcal{B} = \{\alpha, \beta, \gamma\}$, $N_b = 3$, and $\mathbf{L}_{ij,k}$ is the coherence-based pseudolabel for the connection be-
 237 tween channel i and j in the frequency band k . The resulting \mathbf{L}_{ij} represents the integrated coherence
 238 pseudolabel for the three selected physiological bands. Each band is weighted equally to ensure a
 239 balanced contribution to the final representation.
 240

241 This band-mixture coherence approach provides a stable and interpretable measure of connectivity
 242 across functionally relevant frequency ranges, which is well-suited for providing pseudolabels for
 243 the SSL pretext task. It forces the model to learn the underlying inter-channel relationships and
 244 spatial-temporal patterns inherent to EEG data, a task more aligned with downstream cognitive
 245 analysis than low-level signal reconstruction.
 246

247
248 4 MODEL ARCHITECTURE

249 Our proposed model architecture is a two-stage transformer-based framework designed for EEG
 250 signal analysis, comprising both an SSL pre-training stage and a supervised fine-tuning stage. An
 251 illustration of the model architecture is shown in Fig. 3. It maintains a strict one-to-one relationship
 252 between input channels and the channels of learned representations, which is crucial for interpretable
 253 EEG analysis.
 254

255
256 4.1 SELF-SUPERVISED PRE-TRAINING

257 The pre-training stage takes input a multi-channel EEG sequence $\mathbf{X} \in \mathbb{R}^{C \times T}$, where C is the num-
 258 ber of channels and T is the number of time points, the signal is first divided into non-overlapping
 259 temporal patches. Each patch of length P is treated as a token. For an input of $T = 2000$ and
 260 $P = 125$, this yields 16 tokens per channel, resulting in an initial sequence of $16 \times C$ tokens. The
 261 sequence is passed through a linear embedding layer that projects each patch into an embedding
 262 space.
 263

264 Since we are interested in the spatial coherence of whole sequences instead of time segments, a
 265 continuous masking strategy is applied, where a subset of approximately 30% of patches is masked,
 266 followed by the addition of positional encodings to preserve the temporal order. These embeddings
 267 are then processed by a lightweight context encoder, which consists of four transformer blocks, to
 268 generate contextualized representations of the visible signal segments. A cross-attention predictor,
 269 composed of two transformer blocks, is then used to predict the masked representations. Here,
 learnable mask tokens act as queries, while the output features from the context encoder serve as

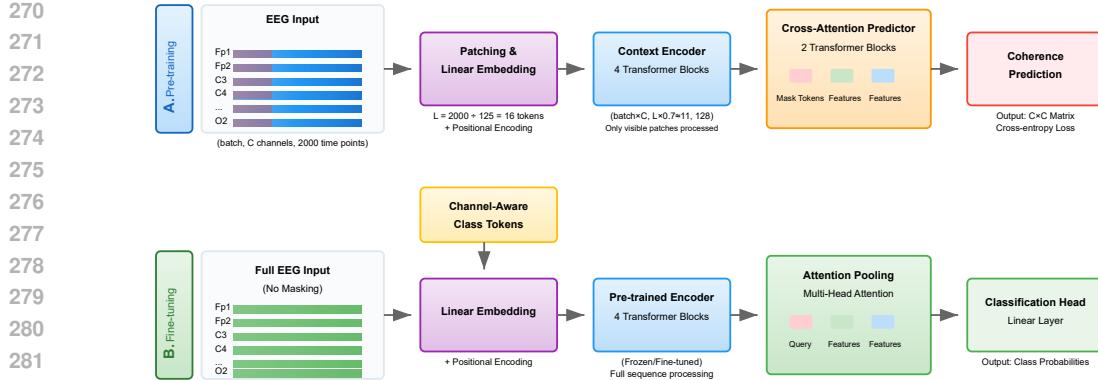


Figure 3: Two-stage SPR architecture with masked pre-training and supervised fine-tuning. **(A)** pre-training stage: A raw C -channel EEG signal (e.g., with channel Fp1, Fp2, ..., O2, and 2000 time points) is temporally patched and continuously masked. The starting indices of masks are the same cross channels but randomly drawn for different batches. Visible patches are processed by a four-block context encoder. A cross-attention predictor then uses mask tokens as queries to predict the masked features, leading to final coherence patterns. **(B)** Fine-tuning stage: Complete EEG signals are processed alongside C channel-aware class tokens. The outputs from the pre-trained encoder are aggregated via attention pooling. Finally, a linear head outputs task-specific predictions.

both keys and values. This design allows the model to predict the content of the masked segments based on the surrounding temporal context.

Crucially, the pre-training objective is not signal reconstruction but rather *coherence prediction*. The model input is partially masked, but targets are computed from the unmasked signal. The model predicts a $C \times C$ coherence matrix that is optimized with a cross-entropy loss.

4.2 SUPERVISED FINE-TUNING

In the fine-tuning stage, the masking strategy is removed, allowing the model to process complete, unmasked EEG sequences. The input undergoes the same patching and linear embedding process as in pre-training, ensuring consistency with the learned representations. The pre-trained context encoder can either be frozen (linear probing) or fine-tuned.

A key feature in this stage is the introduction of *channel-aware class tokens*. The encoder processes the complete sequence of patch tokens and channel-aware class tokens, where the latter are C distinct learnable tokens, each corresponding to a specific EEG channel. An attention-based pooling mechanism is then applied, where the queries are learned and the encoder’s output features serve as keys and values. This attention-based pooling allows the model to selectively aggregate information into channel-weighted representations. Finally, a linear layer maps the channel-weighted representations to the class probability output.

5 DATASETS

We evaluate our approach on four publicly available EEG datasets spanning different applications and acquisition settings. The **Simultaneous Task EEG Workload (STEW)** dataset contains 14-channel recordings from 48 participants performing multitasking workload experiments, with binary classification labels for high/low mental workload (Lim et al., 2018b). From the Temple University Hospital (TUH) EEG Corpus, we utilize two subsets: **TUH Abnormal EEG Corpus (TUAB)** for binary normal/abnormal EEG classification, and **TUH EEG Events (TUEV)** for event detection across six categories (Obeid & Picone, 2016). To compare the performances with foundation models, we also include experimental results for seizure detection on the **Children’s Hospital Boston-MIT (CHB-MIT)** database (Shoeb, 2009; Goldberger et al., 2000).

324 These datasets exhibit diverse characteristics in terms of channel configuration (14-16 channels),
 325 sampling rates (128-256 Hz), and temporal segments (2-10 seconds), providing a comprehensive
 326 evaluation framework. Dataset statistics are summarized in Table 2, with detailed preprocessing and
 327 acquisition protocols described in Appendix A.

329
330 Table 2: Dataset overview

331 Dataset	332 Rate	333 #Channels	334 Length	335 #Samples	336 Task
332 STEW	333 128Hz	334 14	335 2s	336 28,512	337 Mental workload
332 TUAB	333 200Hz	334 16	335 10s	336 409,455	337 Abnormal classification
332 TUEV	333 250Hz	334 16	335 5s	336 113,353	337 Event detection
332 CHB-MIT	333 256Hz	334 16	335 10s	336 326,991	337 Seizure detection

338

6 RESULTS

340 We present our main results in Table 3 and 4, which demonstrate the performance of our proposed
 341 SPR method under linear probing and fine-tuning evaluation protocols, respectively. The reported
 342 baseline metrics are consistent with those in EEG2Rep (Mohammadi Foumani et al., 2024). Each
 343 entry represents the mean value and standard deviation computed over five independent random
 344 seeds to ensure statistical reliability.

345 Our SPR method consistently achieves SOTA performance across all datasets and evaluation metrics
 346 and demonstrates substantial improvements compared to the second-best method under the linear
 347 probing protocol. We observe gains of 4.9% in accuracy and 12.7% in AUROC for STEW. There
 348 are improvements of 11.4% in balanced accuracy and 1.7% in weighted F1 score on TUEV, and
 349 improvements of 2.5% in accuracy and 4.7% in AUROC on TUAB.

350 The fine-tuning results further validate the effectiveness of our approach, with notable improvements
 351 compared to the second-best method across all datasets. We observe gains of 4.7% in accuracy and
 352 11.9% in AUROC on STEW for our pre-trained model. There are also improvements of 9.7% in
 353 balanced accuracy and 3% in weighted F1 score on TUEV, and improvements of 1.6% in accuracy
 354 and 1% in AUROC on TUAB.

355 The consistent performance gains across both evaluation protocols indicate that our learned repre-
 356 sentations are both semantically meaningful and amenable to task-specific adaptation. The fact that
 357 our randomly initialized model serves as a strong baseline indicates that our contributions extend
 358 beyond simply learning better representations to encompassing a more effective overall approach to
 359 EEG analysis. Meanwhile, our pre-trained SPR model establishes new SOTA results by substan-
 360 tial margins. The consistent performance gains across diverse datasets with different characteristics
 361 further underscore the generalizability and robustness of our method.

362
363 Table 3: Linear Probing Performance (%) of SPR in comparison to other methods

364 Method	365 STEW		366 TUEV		367 TUAB	
	368 Accuracy	369 AUROC	370 B-Accuracy	371 W-F1	372 Accuracy	373 AUROC
374 BIOT	375 67.5(2.1)	376 67.7(3.6)	377 40(1.9)	378 66(2)	379 75.1(2.8)	380 82.9(2)
374 BENDR	375 63(1.1)	376 63(1.1)	377 37.4(3.1)	378 61.3(3.2)	379 72.8(4.2)	380 79.9(4.1)
374 MAEEG	375 68(1.9)	376 68.6(1.9)	377 37.2(3)	378 61.4(3.1)	379 72.6(4)	380 79.8(4.1)
374 TS-TCC	375 64.5(1.6)	376 64.6(1.7)	377 36(2.9)	378 60.7(3)	379 74.4(3.1)	380 81(3)
374 TF-C	375 58.8(2.4)	376 58.7(2.4)	377 30.1(4.1)	378 56.2(4.1)	379 69.3(5.8)	380 75.8(3.8)
374 EEG2Rep	375 69(1)	376 69.1(1.2)	377 43.3(3.1)	378 70(3.2)	379 76.6(3.3)	380 83.2(3.3)
374 SPR	375 73.9(1.5)	376 81.8(2.1)	377 54.7(4.3)	378 71.7(3.5)	379 79.1(0.5)	380 87.9(0.5)

375

6.1 ABLATIONS

376 An important observation from our experiments is that loss function selection significantly impacts
 377 performance on highly imbalanced datasets such as TUEV. Given this class imbalance, we employed

378

379

Table 4: Fine-tuning performance (%) of SPR in comparison to other methods

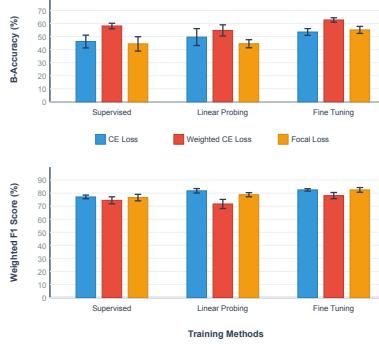
380

381

Method	STEW		TUEV		TUAB	
	Accuracy	AUROC	B-Accuracy	W-F1	Accuracy	AUROC
BIOT	69.9(2.2)	70.1(2.6)	46(1.7)	70(2)	79.2(2.2)	87.4(2)
BENDR	69.7(2.1)	69.8(2)	41.2(2.9)	67.3(3)	77(4)	84(3.4)
MAEEG	72.5(3.7)	72.5(3.2)	41.2(3.7)	67.4(3.7)	77.6(3.6)	86.6(3.3)
TS-TCC	71(3)	71(3)	41(2.6)	68.7(2.9)	79.7(3)	87(2.7)
TF-C	68.7(1.1)	68.7(1.8)	40.1(3.7)	66.2(3.9)	72.3(5.6)	78.5(3.9)
EEG2Rep	73.6(1.5)	74.4(1.5)	53(1.6)	75.1(1.2)	80.5(2.2)	88.4(3.1)
SPR(Random)	75.3(0.9)	83.0(0.8)	58.1(2.1)	75.5(2.7)	79.2(0.3)	86.9(0.3)
SPR (Pre-train)	78.3(1.9)	86.3(2.5)	62.7(1.7)	78.1(2.4)	82.1(0.2)	89.4(0.3)

388

389



390

391

392

393

394

395

396

397

398

399

400

401

402

403

Figure 4: Evaluation of different loss functions on the TUEV dataset. The top subplot shows balanced accuracy performance, and the bottom subplot shows weighted F1 score performance.

404

405

406

weighted cross-entropy loss for supervised learning on TUEV, but it is crucial to examine how different loss function choices affect our model’s performance across various metrics. Fig. 4 reveals distinct patterns in loss function effectiveness. Weighted cross-entropy loss achieves the highest balanced accuracy while maintaining competitive performance on weighted F1 score, though the improvement in F1 is modest. Interestingly, when focusing specifically on fine-tuning results, focal loss demonstrates an optimal balance between balanced accuracy and weighted F1 score, suggesting its effectiveness for task-specific adaptation on imbalanced data.

407

408

409

410

411

412

413

These findings align with our main results and highlight the trade-off in loss function selection. While weighted cross-entropy loss yields significantly higher balanced accuracy, alternative loss functions can achieve substantial improvements in weighted F1 score, albeit with some compromise in balanced accuracy. Importantly, our SPR method demonstrates consistent improvements across all loss function choices, indicating the robustness of our approach.

414

415

416

417

418

419

420

421

422

423

424

425

426

427

428

To provide a comprehensive assessment of each key component’s contribution, we present the averaged accuracy across linear probing and fine-tuning protocols on the STEW dataset in Table 5, which is obtained over five random seeds. In particular, the most substantial performance degradation (4.2% accuracy drop) occurs when positional encoding is removed, and the removal of channel-aware class tokens also results in a 1.7% accuracy reduction. These significant declines underscore the importance of maintaining spatial and temporal relationships within the learned representations, as positional encoding and channel-aware class tokens enable the model to distinguish between identical signal patterns occurring at different time points and electrode locations, providing strong empirical support for our hypothesis that the framework effectively captures spatial-temporal dynamics in EEG signals.

429

430

431

We observe that excessively large τ values significantly degrade performance. Performance deterioration with overly high masking ratios indicates that excessive signal occlusion prevents the model from extracting meaningful coherence features. This suggests an optimal balance where sufficient signal information remains visible to infer spatial-temporal relationships while still providing

432

433

Table 5: Key component analysis on average performance (%)

434

435

436

437

438

439

440

441

442

443

444

445

446

447

448

449

450

451

452

453

454

455

adequate reconstruction challenge for effective learning. Meanwhile, our spatial coherence computation demonstrates superior performance even when diagonal elements are retained. As detailed in Appendix B, the model successfully learns spatial relationships despite strong diagonal coherence values.

460

461

462

6.2 COMPARISON TO FOUNDATION MODELS

463

464

465

466

467

468

469

470

471

472

473

We provide an additional comparison in terms of fine-tuning performance to recently developed models in Table 6, in which we compare our SPR model with both non-foundational models including EEGNet (Lawhern et al., 2018), EEGConformer (Song et al., 2022), SPaRCNet (Jing et al., 2023), CNN-Transformer (Peh et al., 2022), ST-Transformer (Song et al., 2021), and foundation models including BIOT (Yang et al., 2023), CBraMod (Wang et al., 2024b), LaBraM (Jiang et al., 2024). We can observe that the balanced accuracy on TUAB of our method is significantly higher than all other non-foundational models, and is comparable to LaBraM, a much larger foundation model. For the seizure detection task on CHB-MIT, we achieve the highest balanced accuracy and AUROC by a large margin (15.6% and 5.2% higher than the second best model CBraMod), possibly indicating a more useful technique to pre-train the model for the seizure detection task based on local-connectivity patterns.

474

475

476

477

478

479

480

481

482

To explain the CHB-MIT results, we examine the neuroanatomy more closely via visualizations (Fig. 7D and 8 in Appendix B) and observe structured activation patterns that highlight functionally connected brain regions that can help identify focal epilepsy (Watanabe, 1989; Plummer et al., 2008; Foldvary et al., 2001). In Fig. 8, there is more pronounced local activation over **frontal** and **prefrontal** areas, as well as the **temporal** lobe in the β and γ bands, and these regions exhibit correspondingly elevated local connectivity in Fig. 7D. In contrast, central and parietal regions show relatively weak activation and limited local connectivity in the learned heat map. The learned heat map also reveals dominant local connectivity rather than long-range interactions, consistent with our design choice to emphasize local connectivity through higher-frequency coherence.

483

484

485

In Table 7, we further demonstrate the results evaluated on TUAB with cross-domain pre-training. With adding less than 15% pre-train samples, the standard deviation decreases considerably, indicating better generalization capability, and the balanced accuracy increases by 0.4%. This improved performance indicates the robustness and generalization capability of our model across datasets.

486

487

488

489

490

491

492

493

494

495

496

497

498

499

500

501

502

503

504

505

506

507

508

509

510

511

512

513

514

515

516

517

518

519

520

521

522

523

524

525

526

527

528

529

530

531

532

533

534

535

536

537

538

539

Table 6: Fine-tuning performance (%) comparisons with foundation models

Method	Params	TUAB		CHB-MIT	
		B-Accuracy	AUROC	B-Accuracy	AUROC
EEGNet	0.003M	76.4(0.4)	84.1(0.3)	56.6(1.1)	80.5(1.4)
EEGConformer	0.55M	77.6(0.5)	84.5(0.4)	59.8(1.4)	82.3(1.7)
SPaRCNet	0.79M	79(0.2)	86.8(0.1)	58.8(1.9)	81.4(1.5)
CNN-Transformer	3.2M	77.8(0.2)	84.6(0.1)	63.9(0.7)	86.6(0.8)
ST-Transformer	3.5M	79.7(0.2)	87.1(0.2)	59.2(2)	82.4(4.9)
BIOT	3.2M	79.6(0.6)	88.2(0.4)	70.7(4.6)	87.6(2.8)
CBraMod	4M	82.9(0.2)	92.3(0.1)	74(2.8)	88.9(1.5)
LaBraM	5.8M	81.4(0.2)	90.2(0.1)	70.8(3.6)	86.8(2)
SPR (Random)	0.8M	78.7(0.3)	86.9(0.3)	54.7(1.2)	76.3(4.1)
SPR (In-domain)	0.8M	81.3(0.6)	89.4(0.3)	89.6(0.8)	94.1(1.4)

Table 7: Fine-tuning performance (%) of SPR on TUAB with cross-domain pre-training

Dataset	#Samples	B-Accuracy	AUROC
TUAB (Pre-train)	372,510	81.3(0.6)	89.4(0.3)
+ TUEV	425,873	81.7(0.2)	89.4(0.3)

7 DISCUSSION

The results presented in Table 3 and 4 provide evidence for the effectiveness of our SPR framework, demonstrating that our approach not only successfully addresses fundamental challenges in EEG representation learning, but also establishes new performance benchmarks across multiple evaluation protocols and datasets. Our approach introduces a novel time-frequency self-prediction paradigm that combines spatial coherence prediction with implicit reconstruction through masking. The ablation studies further validate our architectural choices and substantiate our hypothesis that effective EEG analysis requires explicit modeling of channel-specific characteristics, as different electrode locations capture distinct neural activities that must be processed with spatial awareness.

The computational complexity of coherence calculation is overall $\mathcal{O}(C^2 N_{freq})$ as the FFT computation $\mathcal{O}(CT \log T)$ is typically smaller for segmented EEG signals ($T \log T \ll CN_{freq}$), where N_{freq} is the number of frequency bins. While the quadratic scaling with the number of channels might seem concerning for high-density EEG systems, the computational burden remains manageable since coherence pseudolabels can be precomputed only once per sample. For extremely high-density channel systems, channel subgrouping strategies can be attractive.

Our band-mixture coherence approach provides a single connectivity value per channel pair within specified physiological frequency bands, effectively summarizing spectral relationships into lower-dimensional representations suitable for deep learning applications, and reducing variance while capturing essential connectivity patterns. However, we acknowledge that averaging information across frequency ranges may occasionally obscure fine-grained frequency-specific relationships. The trade-off between frequency resolution and representation dimensionality represents a conscious design decision that prioritizes robust learning over spectral detail.

Although our primary motivation is to design a spatiality preserving EEG representation learning framework rather than developing a general foundation model, comprehensive comparisons with existing foundation models demonstrate the competitive advantages of our SPR framework. Future research directions include exploring cross-domain pre-training capabilities of the framework, investigating adaptive frequency band selection mechanisms, and extending the approach to other neurophysiological signal modalities. The demonstrated effectiveness of our spatiality preserving paradigm also opens new avenues for developing more sophisticated neural signal representation learning methods that explicitly model the complex spatial-temporal dynamics inherent in brain activities.

540 REFERENCES
541

542 Antonio A. Abello, Roberto Hirata, and Zhangyang Wang. Dissecting the high-frequency bias
543 in convolutional neural networks. In *2021 IEEE/CVF Conference on Computer Vision and*
544 *Pattern Recognition Workshops (CVPRW)*, pp. 863–871, 2021. doi: 10.1109/CVPRW53098.
545 2021.00096.

546 Carlos Amo, Luis De Santiago, Rafael Barea, Almudena López-Dorado, and Luciano Boquete.
547 Analysis of gamma-band activity from human eeg using empirical mode decomposition. *Sensors*,
548 17(5):989, 2017.

549 Gabriele Arnulfo, Sheng H Wang, Vladislav Myrov, Benedetta Toselli, Jonni Hirvonen, MM Fato,
550 L Nobili, F Cardinale, Annalisa Rubino, Alexander Zhigalov, et al. Long-range phase synchroniza-
551 tion of high-frequency oscillations in human cortex. *Nature communications*, 11(1):5363,
552 2020.

553 Hubert Banville, Omar Chehab, Aapo Hyvärinen, Denis-Alexander Engemann, and Alexandre
554 Gramfort. Uncovering the structure of clinical eeg signals with self-supervised learning. *Journal*
555 *of Neural Engineering*, 18(4):046020, 2021.

556 Hangbo Bao, Li Dong, Songhao Piao, and Furu Wei. BEiT: BERT pre-training of image trans-
557 formers. In *International Conference on Learning Representations*, 2022. URL <https://openreview.net/forum?id=p-BhZSz59o4>.

558 Adrien Bardes, Jean Ponce, and Yann LeCun. Vicreg: Variance-invariance-covariance regulariza-
559 tion for self-supervised learning. *arXiv preprint arXiv:2105.04906*, 2021.

560 Erol Başar, Canan Başar-Eroglu, Sirel Karakaş, and Martin Schürmann. Gamma, alpha, delta, and
561 theta oscillations govern cognitive processes. *International journal of psychophysiology*, 39(2-3):
562 241–248, 2001.

563 Walter G Besio, Iris E Martínez-Juárez, Oleksandr Makeyev, John N Gaitanis, Andrew S Blum,
564 Robert S Fisher, and Andrei V Medvedev. High-frequency oscillations recorded on the scalp
565 of patients with epilepsy using tripolar concentric ring electrodes. *IEEE journal of translational*
566 *engineering in health and medicine*, 2:2000111, 2014.

567 William J Bosl, Helen Tager-Flusberg, and Charles A Nelson. Eeg analytics for early detection of
568 autism spectrum disorder: a data-driven approach. *Scientific reports*, 8(1):6828, 2018.

569 Jeffrey W Britton, Lauren C Frey, Jennifer L Hopp, Pearce Korb, Mohamad Z Koubeissi, William E
570 Lievens, Elia M Pestana-Knight, and Erk K St Louis. *Electroencephalography (eeg): An intro-
571 ductory text and atlas of normal and abnormal findings in adults, children, and infants*. 2016.

572 György Buzsáki and Erik W Schomburg. What does gamma coherence tell us about inter-regional
573 neural communication? *Nature neuroscience*, 18(4):484–489, 2015.

574 Mathilde Caron, Hugo Touvron, Ishan Misra, Hervé Jégou, Julien Mairal, Piotr Bojanowski, and
575 Armand Joulin. Emerging properties in self-supervised vision transformers. In *Proceedings of*
576 *the International Conference on Computer Vision (ICCV)*, 2021.

577 Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. A simple framework for
578 contrastive learning of visual representations. In *International conference on machine learning*,
579 pp. 1597–1607. PMLR, 2020.

580 Xinlei Chen and Kaiming He. Exploring simple siamese representation learning. In *Proceedings of*
581 *the IEEE/CVF conference on computer vision and pattern recognition*, pp. 15750–15758, 2021.

582 Yuqi Chen, Kan Ren, Kaitao Song, Yansen Wang, Yifan Wang, Dongsheng Li, and Lili Qiu. Eeg-
583 former: Towards transferable and interpretable large-scale eeg foundation model. *arXiv preprint*
584 *arXiv:2401.10278*, 2024.

585 Zhuying Chen, David B Grayden, Anthony N Burkitt, Udaya Seneviratne, Wendy J D’Souza, Chris
586 French, Philippa J Karoly, Katrina Dell, Kent Leyde, Mark J Cook, et al. Spatiotemporal patterns
587 of high-frequency activity (80–170 hz) in long-term intracranial eeg. *Neurology*, 96(7):e1070–
588 e1081, 2021.

589

594 Hsiang-Yun Sherry Chien, Hanlin Goh, Christopher M Sandino, and Joseph Y Cheng. Maeeg:
 595 Masked auto-encoder for eeg representation learning. [arXiv preprint arXiv:2211.02625](https://arxiv.org/abs/2211.02625), 2022.
 596

597 Wenhui Cui, Woojae Jeong, Philipp Thölke, Takfarinas Medani, Karim Jerbi, Anand A Joshi, and
 598 Richard M Leahy. Neuro-gpt: Towards a foundation model for eeg. In [2024 IEEE International
 599 Symposium on Biomedical Imaging \(ISBI\)](#), pp. 1–5. IEEE, 2024.

600 Silvia López de Diego. [Automated interpretation of abnormal adult electroencephalograms](#). Temple
 601 University, 2017.
 602

603 Sharareh Ehteshamzad. Assessing the potential of eeg in early detection of alzheimer’s disease: A
 604 systematic comprehensive review (2000–2023). [Journal of Alzheimer’s Disease Reports](#), 8(1):
 605 1153–1169, 2024.

606 Emadeldeen Eldele, Mohamed Ragab, Zhenghua Chen, Min Wu, Chee Keong Kwoh, Xiaoli Li, and
 607 Cuntai Guan. Time-series representation learning via temporal and contextual contrasting. [arXiv
 608 preprint arXiv:2106.14112](https://arxiv.org/abs/2106.14112), 2021.
 609

610 Paul J Fitzgerald and Brendon O Watson. Gamma oscillations as a biomarker for major depression:
 611 an emerging topic. [Translational psychiatry](#), 8(1):177, 2018.

612 Nancy Foldvary, G Klem, J Hammel, W Bingaman, I Najm, and H Luders. The localizing value of
 613 ictal eeg in focal epilepsy. [Neurology](#), 57(11):2022–2028, 2001.
 614

615 Ary L Goldberger, Luis AN Amaral, Leon Glass, Jeffrey M Hausdorff, Plamen Ch Ivanov, Roger G
 616 Mark, Joseph E Mietus, George B Moody, Chung-Kang Peng, and H Eugene Stanley. Physiobank,
 617 physiotoolkit, and physionet: components of a new research resource for complex physiologic
 618 signals. [circulation](#), 101(23):e215–e220, 2000.

619 Pierre Guetschel, Sara Ahmadi, and Michael Tangermann. Review of deep representation learning
 620 techniques for brain–computer interfaces. [Journal of Neural Engineering](#), 21(6):061002, 2024.
 621

622 Jiazen Hong, Geoffrey Mackellar, and Soheila Ghane. Eegm2: An efficient mamba-2-based self-
 623 supervised framework for long-sequence eeg modeling. [arXiv preprint arXiv:2502.17873](https://arxiv.org/abs/2502.17873), 2025.
 624

625 Mohammad-Parsa Hosseini, Amin Hosseini, and Kiarash Ahi. A review on machine learning for
 626 eeg signal processing in bioengineering. [IEEE Reviews in Biomedical Engineering](#), 14:204–218,
 627 2021. doi: 10.1109/RBME.2020.2969915.

628 Jaeseung Jeong. Eeg dynamics in patients with alzheimer’s disease. [Clinical neurophysiology](#), 115
 629 (7):1490–1505, 2004.

630 Xiaoxuan Jia, Matthew A Smith, and Adam Kohn. Stimulus selectivity and spatial coherence of
 631 gamma components of the local field potential. [Journal of Neuroscience](#), 31(25):9390–9403,
 632 2011.
 633

634 Wei-Bang Jiang, Li-Ming Zhao, and Bao-Liang Lu. Large brain model for learning generic represen-
 635 tations with tremendous EEG data in BCI. In [The Twelfth International Conference on Learning
 636 Representations](#), 2024. URL <https://openreview.net/forum?id=QzTpTRVtrP>.

637 Wei-Bang Jiang, Yansen Wang, Bao-Liang Lu, and Dongsheng Li. NeuroLM: A universal multi-
 638 task foundation model for bridging the gap between language and EEG signals. In [The Thirteenth
 639 International Conference on Learning Representations](#), 2025. URL <https://openreview.net/forum?id=Io9yFt7XH7>.
 640

641 Jin Jing, Wendong Ge, Shenda Hong, Marta Bento Fernandes, Zhen Lin, Chaoqi Yang, Sungtae An,
 642 Aaron F Struck, Aline Herlopian, Ioannis Karakis, et al. Development of expert-level classifi-
 643 cation of seizures and rhythmic and periodic patterns during eeg interpretation. [Neurology](#), 100
 644 (17):e1750–e1762, 2023.
 645

646 Demetres Kostas, Stephane Aroca-Ouellette, and Frank Rudzicz. Bendr: Using transformers and a
 647 contrastive self-supervised learning task to learn from massive amounts of eeg data. [Frontiers in
 Human Neuroscience](#), 15:653659, 2021.

648 Vernon J Lawhern, Amelia J Solon, Nicholas R Waytowich, Stephen M Gordon, Chou P Hung, and
 649 Brent J Lance. Eegnet: a compact convolutional neural network for eeg-based brain–computer
 650 interfaces. *Journal of neural engineering*, 15(5):056013, 2018.

651

652 Szymon Łęski, Henrik Lindén, Tom Tetzlaff, Klas H Pettersen, and Gaute T Einevoll. Frequency
 653 dependence of signal power and spatial reach of the local field potential. *PLoS computational
 654 biology*, 9(7):e1003137, 2013.

655 W. L. Lim, O. Sourina, and L. P. Wang. Stew: Simultaneous task eeg workload dataset, 2018a. URL
 656 <https://dx.doi.org/10.21227/44r8-ya50>.

657

658 W. L. Lim, O. Sourina, and L. P. Wang. Stew: Simultaneous task eeg workload data set. *IEEE
 659 Transactions on Neural Systems and Rehabilitation Engineering*, 26(11):2106–2114, 2018b. doi:
 660 10.1109/TNSRE.2018.2872924.

661 Ziyu Liu, Azadeh Alavi, Minyi Li, and Xiang Zhang. Self-supervised contrastive learning for med-
 662 ical time series: A systematic review. *Sensors*, 23(9):4221, 2023.

663

664 Navid Mohammadi Foumani, Geoffrey Mackellar, Soheila Ghane, Saad Irtza, Nam Nguyen, and
 665 Mahsa Salehi. Eeg2rep: enhancing self-supervised eeg representation through informative
 666 masked inputs. In *Proceedings of the 30th ACM SIGKDD Conference on Knowledge Discovery
 667 and Data Mining*, pp. 5544–5555, 2024.

668 Mostafa Neo Mohsenvand, Mohammad Rasool Izadi, and Pattie Maes. Contrastive representation
 669 learning for electroencephalogram classification. In *Machine learning for health*, pp. 238–253.
 670 PMLR, 2020.

671 Juan Molina, Mircea Petrache, Francisco Sahli Costabal, and Matías Courdurier. Understanding the
 672 dynamics of the frequency bias in neural networks. *arXiv preprint arXiv:2405.14957*, 2024.

673

674 Jennifer J Newson and Tara C Thiagarajan. Eeg frequency bands in psychiatric disorders: a review
 675 of resting state studies. *Frontiers in human neuroscience*, 12:521, 2019.

676

677 Iyad Obeid and Joseph Picone. The temple university hospital eeg data corpus. *Frontiers in
 678 neuroscience*, 10:196, 2016.

679

680 Shankar Parmar and Chirag Paunwala. Early detection of dyslexia based on eeg with novel predictor
 681 extraction and selection. *Discover Artificial Intelligence*, 3(1):33, 2023.

682

683 Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor
 684 Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Kopf, Edward
 685 Yang, Zachary DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner,
 686 Lu Fang, Junjie Bai, and Soumith Chintala. Pytorch: An imperative style, high-performance
 687 deep learning library. In *Advances in Neural Information Processing Systems* 32, pp. 8024–8035.
 Curran Associates, Inc., 2019.

688

689 Wei Yan Peh, Yuanyuan Yao, and Justin Dauwels. Transformer convolutional neural networks for
 690 automated artifact detection in scalp eeg. In *2022 44th Annual International Conference of the
 691 IEEE Engineering in Medicine & Biology Society (EMBC)*, pp. 3599–3602. IEEE, 2022.

692

693 Zhiliang Peng, Li Dong, Hangbo Bao, Qixiang Ye, and Furu Wei. Beit v2: Masked image modeling
 694 with vector-quantized visual tokenizers. *arXiv preprint arXiv:2208.06366*, 2022.

695

696 Chris Plummer, A Simon Harvey, and Mark Cook. Eeg source localization in focal epilepsy: where
 697 are we now? *Epilepsia*, 49(2):201–218, 2008.

698

699 Jathurshan Pradeepkumar, Xihao Piao, Zheng Chen, and Jimeng Sun. Single-channel eeg tokeniza-
 700 tion through time-frequency modeling. *arXiv preprint arXiv:2502.16060*, 2025.

701

Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal,
 Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual
 models from natural language supervision. In *International conference on machine learning*, pp.
 8748–8763. PMLR, 2021.

702 Mohammad H Rafiei, Lynne V Gauthier, Hojjat Adeli, and Daniel Takabi. Self-supervised learning
 703 for electroencephalography. *IEEE Transactions on Neural Networks and Learning Systems*, 35
 704 (2):1457–1471, 2022.

705

706 Ali Hossam Shoeb. Application of machine learning to epileptic seizure onset detection and
 707 treatment. PhD thesis, Massachusetts Institute of Technology, 2009.

708

709 Yonghao Song, Xueyu Jia, Lie Yang, and Longhan Xie. Transformer-based spatial-temporal feature
 710 learning for eeg decoding. *arXiv preprint arXiv:2106.11170*, 2021.

711

712 Yonghao Song, Qingqing Zheng, Bingchuan Liu, and Xiaorong Gao. Eeg conformer: Convolutional
 713 transformer for eeg decoding and visualization. *IEEE Transactions on Neural Systems and
 Rehabilitation Engineering*, 31:710–719, 2022.

714

715 Michal Teplan et al. Fundamentals of eeg measurement. *Measurement science review*, 2(2):1–11,
 716 2002.

717

718 S.P van den Broek, F Reinders, M Donderwinkel, and M.J Peters. Volume conduction effects in
 719 eeg and meg. *Electroencephalography and Clinical Neurophysiology*, 106(6):522–534, 1998.
 720 ISSN 0013-4694. doi: [https://doi.org/10.1016/S0013-4694\(97\)00147-8](https://doi.org/10.1016/S0013-4694(97)00147-8). URL <https://www.sciencedirect.com/science/article/pii/S0013469497001478>.

721

722 Guangyu Wang, Wenchao Liu, Yuhong He, Cong Xu, Lin Ma, and Haifeng Li. Eegpt: Pre-
 723 trained transformer for universal and reliable representation of eeg signals. *Advances in Neural
 724 Information Processing Systems*, 37:39249–39280, 2024a.

725

726 Jiquan Wang, Sha Zhao, Zhiling Luo, Yangxuan Zhou, Haiteng Jiang, Shijian Li, Tao Li, and
 727 Gang Pan. Cbramod: A criss-cross brain foundation model for eeg decoding. *arXiv preprint
 728 arXiv:2412.07236*, 2024b.

729

730 Wenhui Wang, Hangbo Bao, Li Dong, Johan Bjorck, Zhiliang Peng, Qiang Liu, Kriti Aggarwal,
 731 Owais Khan Mohammed, Saksham Singhal, Subhajit Som, and Furu Wei. Image as a foreign
 732 language: BEiT pretraining for vision and vision-language tasks. In *Proceedings of the IEEE/CVF
 Conference on Computer Vision and Pattern Recognition*, 2023.

733

734 Yihe Wang, Nan Huang, Nadia Mammone, Marco Cecchi, and Xiang Zhang. Lead: Large founda-
 735 tion model for eeg-based alzheimer’s disease detection. *arXiv preprint arXiv:2502.01678*, 2025.

736

737 Kazuyoshi Watanabe. The localization-related epilepsies: Some problems with subclassification.
 738 *Psychiatry and Clinical Neurosciences*, 43(3):471–475, 1989.

739

740 Weining Weng, Yang Gu, Shuai Guo, Yuan Ma, Zhaohua Yang, Yuchen Liu, and Yiqiang Chen. Self-
 741 supervised learning for electroencephalogram: A systematic survey. *ACM Computing Surveys*,
 742 57(12):1–38, 2025.

743

744 Chaoqi Yang, M Brandon Westover, and Jimeng Sun. Biot: Biosignal transformer for cross-data
 745 learning in the wild. In *Thirty-seventh Conference on Neural Information Processing Systems*,
 2023. URL <https://openreview.net/forum?id=c2LZyTyddi>.

746

747 Kai Yang, Li Tong, Jun Shu, Ning Zhuang, Bin Yan, and Ying Zeng. High gamma band eeg closely
 748 related to emotion: evidence from functional network. *Frontiers in human neuroscience*, 14:89,
 2020.

749

750 Jure Zbontar, Li Jing, Ishan Misra, Yann LeCun, and Stéphane Deny. Barlow twins: Self-supervised
 751 learning via redundancy reduction. In *International Conference on Machine Learning*, pp. 12310–
 752 12320. PMLR, 2021.

753

754 Xiang Zhang, Ziyuan Zhao, Theodoros Tsiligkaridis, and Marinka Zitnik. Self-supervised
 755 contrastive pre-training for time series via time-frequency consistency. *Advances in neural
 information processing systems*, 35:3988–4003, 2022.

756 A DETAILED DATASET DESCRIPTION AND PREPROCESSING
757758 A.1 SIMULTANEOUS TASK EEG WORKLOAD (STEW)
759760 The STEW dataset comprises raw EEG recordings from 48 participants performing multitasking
761 workload experiments using the SIMKAP multitasking test. Each experimental session began with
762 baseline resting-state recordings to establish individual neural activity baselines prior to task engage-
763 ment. EEG signals were acquired using a 14-channel Emotiv EPOC headset at 128 Hz sampling
764 frequency, providing 2.5 minutes of continuous recording per participant. Following each experi-
765 mental stage, participants provided subjective mental workload assessments using a 9-point scale (1
766 = very low workload, 9 = very high workload).767 **Preprocessing:** EEG signals were segmented into 2-second non-overlapping epochs, yielding a total
768 of 28,512 samples. The dataset followed a standardized train-validation-test partition: 17,820 training
769 samples, 3,564 validation samples, and 7,128 test samples. For self-supervised pre-training, we
770 combined the training and validation sets (21,384 samples in total) to maximize the effectiveness of
771 representation learning. Binary classification labels were derived by thresholding subjective ratings,
772 with scores > 4 categorized as “high workload” and scores ≤ 4 as “low workload”. We followed the
773 preprocessing method in EEG2Rep (Mohammadi Foumani et al., 2024) and applied an additional
774 14 bipolar montage to the dataset: “AF3-F7”, “F7-T7”, “T7-P7”, “P7-O1”, “AF4-F8”, “F8-T8”,
775 “T8-P8”, “P8-O2”, “AF3-F3”, “F3-FC5”, “AF4-F4”, “F4-FC6”, “F3-F4”, and “FC5-FC6”.
776777 **Access:** STEW is publicly accessible through IEEE DataPort at <https://dx.doi.org/10.21227/44r8-ya50> (Lim et al., 2018a).778 A.2 TEMPLE UNIVERSITY HOSPITAL EEG CORPUS
779780 A.2.1 TUH ABNORMAL EEG CORPUS (TUAB)
781782 The Temple University Abnormal EEG Corpus (TUAB) represents a curated subset of the larger
783 Temple University Hospital EEG dataset, specifically comprising clinical recordings with abnormal
784 neurological findings. This clinical dataset contains 2339 recordings and has a slightly unbalanced
785 class distribution (42% abnormal, 58% normal) (de Diego, 2017). TUAB serves as a critical benchmark
786 for evaluating automated neurological disorder detection systems.787 **Preprocessing:** Following the preprocessing protocol established by BIOT (Yang et al., 2023), we
788 employed a standardized 16-channel bipolar montage configuration adhering to the international
789 10-20 electrode placement system. The bipolar electrode pairs were organized into four chains:
790 “FP1-F7”, “F7-T3”, “T3-T5”, “T5-O1”, “FP2-F8”, “F8-T4”, “T4-T6”, “T6-O2”, “FP1-F3”, “F3-
791 “C3-C4”, “C3-P3”, “P3-O1”, “FP2-F4”, “F4-C4”, “C4-P4”, and “P4-O2”. Each channel underwent
792 amplitude normalization using the 95th percentile of absolute amplitude values as the scaling factor,
793 providing robust normalization that minimizes the influence of extreme artifacts while preserving
794 physiologically relevant signal dynamics. EEG signals were subsequently segmented into 10-second
795 non-overlapping epochs, yielding a total of 409,455 samples for binary normal/abnormal classifica-
796 tion tasks. Those samples were partitioned to 299,656 training samples (73.2%), 72,854 validation
797 samples (17.8%), and 36,945 test samples (9.0%). For self-supervised pre-training, we combined
798 the training and validation sets to create an expanded pre-training corpus of 372,510 samples.799 A.2.2 TUH EEG EVENTS (TUEV)
800801 The Temple University EEG Events (TUEV) corpus comprises clinically annotated EEG segments
802 with expert-labeled neurological events across six distinct categories: (1) spike and sharp wave, (2)
803 generalized periodic epileptiform discharges, (3) periodic lateralized epileptiform discharges, (4)
804 eye movement, (5) artifacts, and (6) background activity. This comprehensive annotation scheme
805 encompasses both pathological epileptiform patterns and common non-pathological events, mak-
806 ing TUEV particularly valuable for developing robust automated EEG event detection systems in
807 clinical settings.808 **Preprocessing:** TUEV follows an identical preprocessing pipeline to TUAB, employing the same
809 standardized 16-channel bipolar montage configuration based on the international 10-20 system.
Each channel underwent robust amplitude normalization using the 95th percentile of absolute am-

810 **plitude** values. Given the shorter duration requirements for event detection tasks, EEG signals were
 811 segmented into 5-second non-overlapping epochs, producing a total of 113,353 labeled samples for
 812 multi-class event classification. The dataset exhibits inherent class imbalance typical of clinical EEG
 813 event detection scenarios, where pathological events are naturally less frequent than background ac-
 814 tivity. The dataset was partitioned into 75,859 training samples (66.9%), 8,073 validation samples
 815 (7.1%), and 29,421 test samples (26.0%). For self-supervised pre-training, we combined training
 816 and validation sets to create an expanded corpus of 83,932 samples.

817 **Access:** Both TUAB and TUEV datasets are accessible through the Temple University EEG Re-
 818 sources repository upon request at [https://isip.piconepress.com/projects/nedc/](https://isip.piconepress.com/projects/nedc/html/tuh_eeg/index.shtml)
 819 [html/tuh_eeg/index.shtml](https://isip.piconepress.com/projects/nedc/html/tuh_eeg/index.shtml).

823 A.3 THE CHILDREN’S HOSPITAL BOSTON-MIT (CHB-MIT)

824
 825 The Children’s Hospital Boston-MIT (CHB-MIT) database (Shoeb, 2009; Goldberger et al., 2000)
 826 represents a publicly available dataset comprising long-term EEG recordings from 24 pediatric pa-
 827 tients (ages 3-22 years) with medically intractable epilepsy. The recordings were acquired during
 828 presurgical evaluation periods following controlled withdrawal of anti-seizure medications, enabling
 829 comprehensive characterization of spontaneous seizure patterns and assessment of surgical candi-
 830 dacy. This clinical protocol provides authentic seizure events in their natural temporal context,
 831 making CHB-MIT a good standard for evaluating automated seizure detection algorithms. EEG
 832 signals were recorded using scalp electrodes positioned according to the international 10-20 system,
 833 with data acquisition at 256 Hz sampling frequency to capture high-frequency seizure components.
 834 The dataset encompasses multi-day continuous recordings, providing both ictal (seizure) and exten-
 835 sive interictal (non-seizure) periods. Each recording session contains expert neurologist annotations
 836 marking precise seizure onset and offset times, enabling binary classification into seizure and non-
 837 seizure states with clinical-grade ground truth labels.

838 **Preprocessing:** Following the preprocessing methodology established by BIOT, we employed
 839 a standardized 16-channel bipolar montage configuration to ensure consistency with established
 840 benchmarks. The selected electrode pairs form four chains: “FP1-F7”, “F7-T7”, “T7-P7”, “P7-O1”,
 841 “FP2-F8”, “F8-T8”, “T8-P8”, “P8-O2”, “FP1-F3”, “F3-C3”, “C3-P3”, “P3-O1”, “FP2-F4”, “F4-
 842 “C4”, “C4-P4”, “P4-O2”. Each channel underwent robust amplitude normalization using the 95th
 843 percentile of absolute amplitude values to maintain signal integrity while reducing inter-subject vari-
 844 ability. The continuous recordings were segmented into 10-second non-overlapping epochs, yielding
 845 326,991 labeled samples suitable for seizure detection tasks. The dataset was divided into 286,964
 846 training samples (87.8%), 23,065 validation samples (7.1%), and 16,962 test samples (5.2%). For
 847 self-supervised pre-training, we combined training and validation sets to create an expanded corpus
 848 of 310,029 samples.

849 **Access:** The CHB-MIT dataset is accessible through PhysioNet at [https://doi.org/10.](https://doi.org/10.13026/C2K01R)
 850 [13026/C2K01R](https://doi.org/10.13026/C2K01R).

851 852 A.4 DATA PARTITIONING

853
 854 For the STEW dataset, we implemented a subject-wise train/validation/test split to evaluate the
 855 model’s cross-subject generalization capabilities and account for inter-subject variability. For TUAB
 856 and TUEV, we adhered to the inherent training/test splits provided with the datasets, which are con-
 857 sistent with established benchmarking protocols. Within the designated pre-training sets, we created
 858 subject-wise validation splits. For TUAB, the pre-training set was split into 80% training and 20%
 859 validation subsets. For TUEV, we divided pre-training sets into 90% training and 10% validation
 860 subsets. The CHB-MIT dataset was partitioned by subject to evaluate performance on unseen indi-
 861 viduals. Subjects 1-20 were designated for the training set, subjects 21-22 for the validation set, and
 862 subjects 23-24 for the test set.

864 **B VISUALIZATION**
865866 The visualizations on TUAB and CHB-MIT are shown in Fig. 5, 6 and 7, 8, respectively. The
867 channels matrices contain the axis ordered by both the bipolar chains and brain regions. The * sign
868 signals the same channel name but followed by a different referencing path.869 A key observation from the coherence matrices is that despite the dominance of diagonal elements
870 across all target labels, our model successfully learns to focus on meaningful off-diagonal spatial
871 relationships rather than being misled by these trivially high self-coherence values. This suggests that
872 our model effectively captures the underlying neural dynamics while mitigating potential artifacts
873 from self-coherence terms.874 The model predictions exhibit clear chain-wise and region-wise activation patterns in both Fig. 5D
875 and 7D, providing evidence that our framework effectively captures genuine spatial relationships
876 between electrode locations. These structured activation patterns indicate that the model learns to
877 identify functionally connected brain regions rather than simply memorizing statistical regularities
878 in the pre-training dataset.879 When examining predictions across different clinical conditions, we observe distinct spatial
880 relationship patterns. Normal and non-seizure subjects exhibit sparser spatial connectivity predictions,
881 while abnormal and seizure subjects demonstrate visually denser interconnection patterns. This
882 class-dependent variation in learned spatial representations creates discriminative features that
883 naturally separate different underlying pathological states, thereby enhancing the utility of learned rep-
884 resentations for downstream tasks.885 Meanwhile, both ground truth labels and model predictions consistently demonstrate the well-
886 established adjacency effect, where spatially neighboring electrodes exhibit high coherence regard-
887 less of their anatomical brain region assignments. For instance, electrode pairs T5-P3 and P4-T6
888 show strong coherence in the normal subject of Fig. 5A, and Fp2-F4 in the non-seizure state of
889 Fig. 7C as well, reflecting the physical proximity of these recording sites. This neurophysiologically
890 plausible pattern provides strong validation for using coherence as a prediction target, as it confirms
891 that our model learns spatial relationships that align with known principles of neural signal propa-
892 gation. However, we need to point out that this adjacency effect is susceptible to volume conduction
893 effects, which can create spurious correlations between channels due to signal propagation through
894 conductive head tissues (van den Broek et al., 1998).895 The topographic maps presented in Fig. 6 and 8 reveal a clear frequency-dependent pattern in spatial
896 specificity, providing strong empirical support for our frequency band selection strategy. Lower
897 frequency bands (δ and θ) exhibit diffuse, spatially non-specific activation patterns with limited
898 discriminative power between electrode locations. In contrast, higher frequency bands (α , β , and
899 γ) demonstrate increasingly focal and spatially differentiated patterns. In particular, the γ band
900 here exhibits the most pronounced spatial specificity and demonstrates visually distinct topographic
901 patterns among normal versus abnormal subjects, as well as seizure versus non-seizure states.902 The enhanced spatial discriminability in the α , β , and γ ranges suggests that these frequencies
903 carry more informative spatial-temporal signatures for representation learning. By prioritizing these
904 frequency ranges, our model can take advantage of their inherent spatial specificity to learn more
905 meaningful and diagnostically relevant representations.906
907 **C MATHEMATICAL FRAMEWORK**
908909 For discrete-time signals $x[n]$ of length N , the Discrete Fourier Transform is computed as:

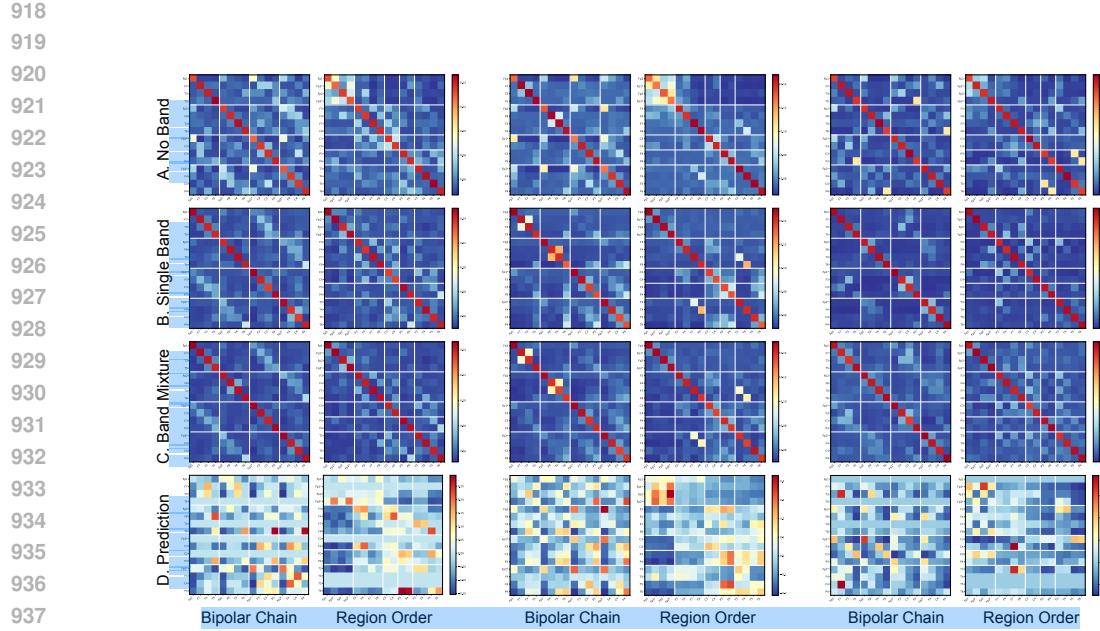
910
$$X[k] = \sum_{n=0}^{N-1} x[n]e^{-j2\pi kn/N}, \quad k = 0, 1, \dots, N-1 \quad (7)$$

911
912
913

914 The frequency bins correspond to $f[k] = \frac{k f_s}{N}$, where f_s is the sampling frequency.915 Given a frequency band defined by $[f_{low}, f_{high}]$, we define the band mask as:

916
$$M[k] = \begin{cases} 1 & \text{if } f_{low} \leq |f[k]| \leq f_{high} \\ 0 & \text{otherwise} \end{cases} \quad (8)$$

917



939
940
941
942
943
944
945
946
947
948
949
950
951
952
953
954
955
956
957
958
959
960
961
962
963
964
965
966
967
968
969
970
971

Figure 5: Channel Matrices for three different samples (abnormal, abnormal, normal) from TUAB. **(A)** Coherence labels without band-limits. **(B)** Coherence labels with single band (8–100 Hz). **(C)** Coherence labels with band mixture (8–13, 13–30, 30–100 Hz). **(D)** Coherence labels predicted by the model.

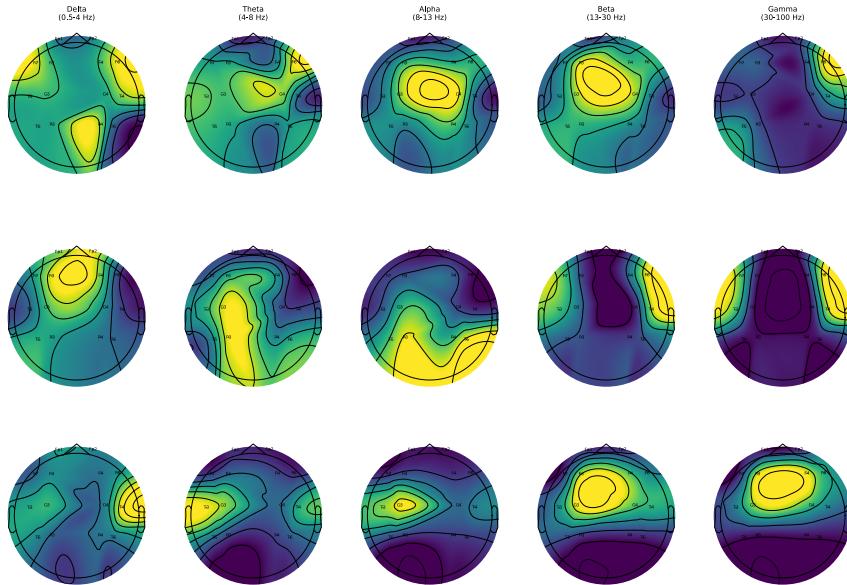
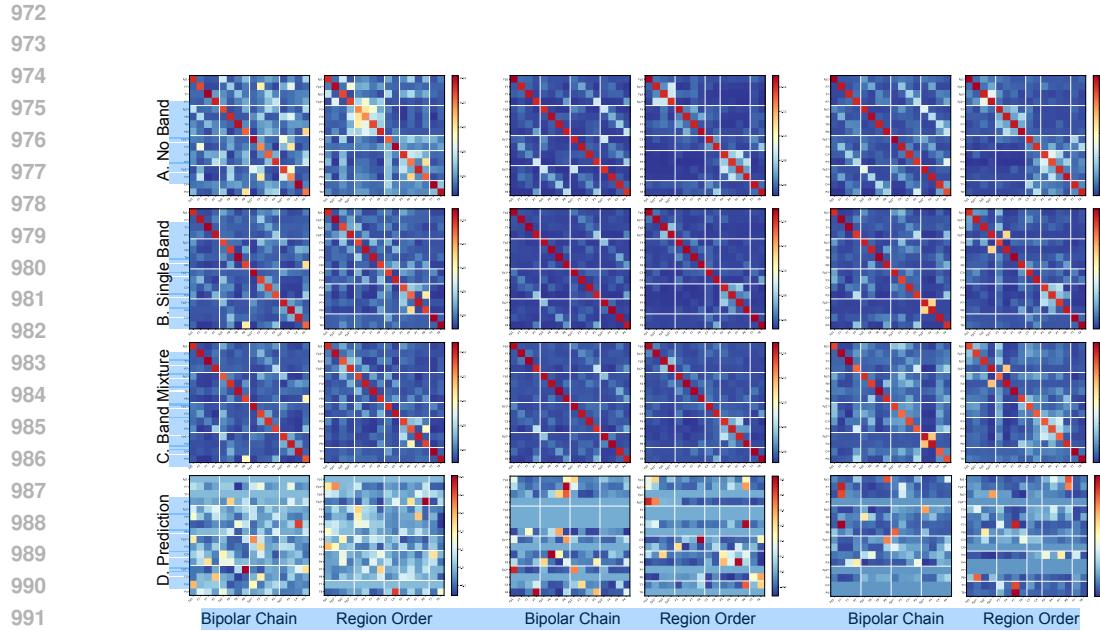


Figure 6: The topographic maps in power (dB) for three different samples (abnormal, abnormal, normal) from TUAB. The values are averaged for channels at the same position. Each sample corresponds to the sample in Fig. 5.



993 Figure 7: Channel Matrices for three different samples (seizure, seizure, non-seizure) from CHB-
994 MIT. (A) Coherence labels without band-limits. (B) Coherence labels with single band (8–100 Hz).
995 (C) Coherence labels with band mixture (8–13, 13–30, 30–100 Hz). (D) Coherence labels predicted
996 by the model.

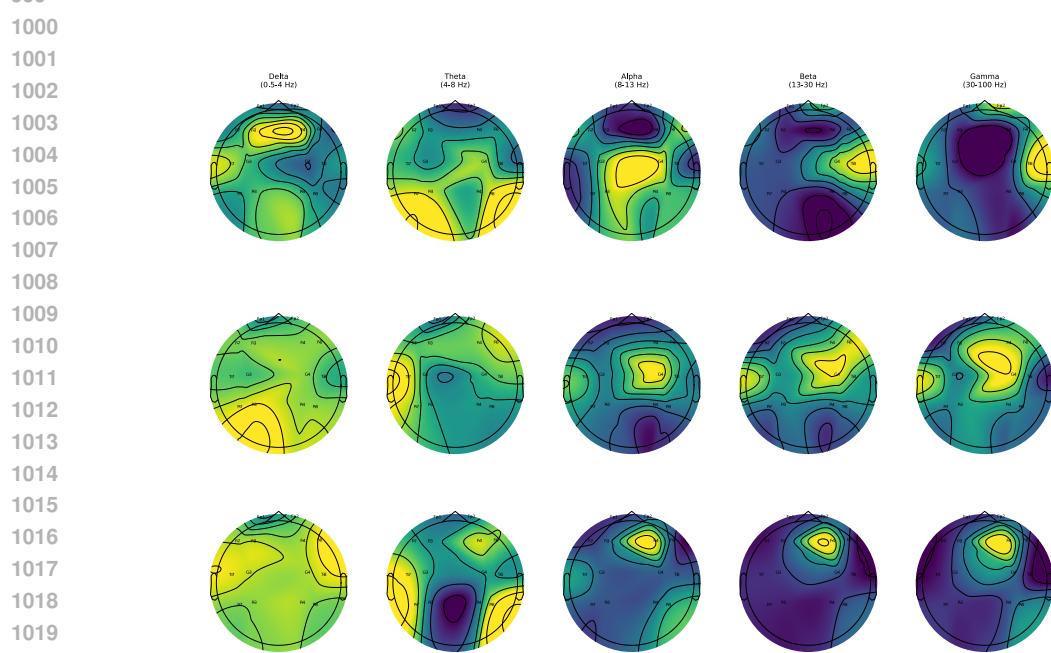


Figure 8: The topographic maps in power (dB) for three different samples (seizure, seizure, non-seizure) from CHB-MIT. The values are averaged for channels at the same position. Each sample corresponds to the sample in Fig. 7.

1026 The band-limited cross-spectral density between channel i and j is:
 1027

$$1028 \quad \tilde{S}_{ij} = \frac{1}{N_{freq}} \sum_{k: M[k]=1} X_i[k] \cdot X_j^*[k] \quad (9)$$

1031 Here, N_{freq} is the number of frequency bins in the band.

1032 The coherence function is formally defined as the ratio of the CPSDs to the product of the individual
 1033 PSDs:

$$1034 \quad \gamma_{xy}^2(f) = \frac{|S_{xy}(f)|^2}{S_{xx}(f)S_{yy}(f)} \quad (10)$$

1037 Here, $S_{xx}(f)$ and $S_{yy}(f)$ are the PSDs of $x[n]$ and $y[n]$, respectively, and $S_{xy}(f)$ is the CPSD.

1038 These spectral densities are the Fourier transforms of the corresponding auto-correlation and cross-
 1039 correlation functions. The non-negativity and boundedness of this function are a direct consequence
 1040 of the Cauchy-Schwarz inequality applied to the spectral domain. Meanwhile, the diagonal elements
 1041 of the coherence matrix are always unity:

$$1042 \quad \gamma_{ii} = \frac{|S_{ii}|^2}{S_{ii} \cdot S_{ii}} = 1 \quad (11)$$

1045 **Theorem 1:** Softmax of band-averaged coherence is **NOT** equivalent to averaging individual soft-
 1046 max coherence values:

$$1048 \quad \text{softmax}(\bar{\gamma}^2) \neq \frac{1}{N_b} \sum_{k \in \mathcal{B}} \text{softmax}(\gamma_{xy}^2(f_k)) \quad (12)$$

1050 *Proof.* For two frequency bins with coherences γ_1^2 and γ_2^2 , we compare Softmax of band-averaged
 1051 coherence and average of individual softmax values:

$$1053 \quad \text{softmax}(\bar{\gamma}^2) = \frac{e^{\bar{\gamma}^2}}{\sum_j e^{\bar{\gamma}_j^2}} \quad (13)$$

$$1056 \quad \frac{\text{softmax}(\gamma_1^2) + \text{softmax}(\gamma_2^2)}{2} = \frac{1}{2} \left(\frac{e^{\gamma_1^2}}{\sum_j e^{\gamma_j^2}} + \frac{e^{\gamma_2^2}}{\sum_j e^{\gamma_j^2}} \right) \quad (14)$$

1059 Here,

$$1061 \quad \bar{\gamma}^2 = \frac{4|S_{xy,1} + S_{xy,2}|^2}{(S_{xx,1} + S_{xx,2})(S_{yy,1} + S_{yy,2})} \quad (15)$$

1064 The inequality first arises from coherence nonlinearity: The band-averaged coherence $\bar{\gamma}^2$ is generally
 1065 not equal to the arithmetic mean $\frac{\gamma_1^2 + \gamma_2^2}{2}$ due to:

$$1067 \quad \frac{4|S_{xy,1} + S_{xy,2}|^2}{(S_{xx,1} + S_{xx,2})(S_{yy,1} + S_{yy,2})} \neq \frac{1}{2} \left(\frac{|S_{xy,1}|^2}{S_{xx,1}S_{yy,1}} + \frac{|S_{xy,2}|^2}{S_{xx,2}S_{yy,2}} \right) \quad (16)$$

1070 Even if $\bar{\gamma}^2 = \frac{\gamma_1^2 + \gamma_2^2}{2}$, Jensen's inequality for the convex function e^x would still create inequality in
 1071 most cases. Since $f(x) = e^x$ is strictly convex, Jensen's inequality states:

$$1073 \quad f\left(\frac{x_1 + x_2}{2}\right) < \frac{f(x_1) + f(x_2)}{2} \quad (17)$$

1076 Equality occurs only when $\bar{\gamma}^2 = \frac{\gamma_1^2 + \gamma_2^2}{2}$ and $\gamma_1^2 = \gamma_2^2$ are both satisfied. For EEG signals, this
 1077 requires:

1079

- Identical coherences at both frequency bins.

- Specific phase and amplitude relationships in cross-spectral densities.
- Equal power spectral densities: $S_{xx,1} = S_{xx,2}$ and $S_{yy,1} = S_{yy,2}$.

Thus, the combination of coherence calculation nonlinearity and Jensen's inequality for the exponential function ensures:

$$\text{softmax}(\bar{\gamma}_{xy}^2) \neq \frac{1}{N_b} \sum_{k \in \mathcal{B}} \text{softmax}(\gamma_{xy}^2(f_k))$$

□

Theorem 2: Under the assumption of independent frequency bins, band-averaging acts as a form of spectral smoothing that trades bias for variance reduction.

Proof. The variance of the arithmetic mean is:

$$\begin{aligned} \text{Var}[\bar{\gamma}_{xy}^2] &= \text{Var} \left[\frac{1}{N_b} \sum_i \bar{\gamma}_{xy,i}^2 \right] \\ &= \frac{1}{N_b^2} \text{Var} \left[\sum_i \bar{\gamma}_{xy,i}^2 \right] \\ &= \frac{1}{N_b^2} \left[\sum_i \text{Var}[\bar{\gamma}_{xy,i}^2] + 2 \sum_{i < j} \text{Cov}[\bar{\gamma}_{xy,i}^2, \bar{\gamma}_{xy,j}^2] \right] \end{aligned} \quad (18)$$

If band-limited coherences are independent:

$$\text{Var}[\bar{\gamma}_{xy}^2] = \frac{1}{N_b^2} \sum_i \text{Var}[\bar{\gamma}_{xy,i}^2] = \frac{1}{N_b} \overline{\text{Var}[\bar{\gamma}_{xy}^2]} \quad (19)$$

This shows a variance reduction by factor of $\frac{1}{N_b}$ compared to individual band estimates.

□

Theorem 3: The softmax transformation $\mathbf{L} = \text{softmax}(\bar{\gamma}_{xy}^2)$ preserves ordinal relationships:

$$\bar{\gamma}_{xy,i}^2 > \bar{\gamma}_{xy,j}^2 \Rightarrow y_i > y_j \quad (20)$$

Proof. If $\bar{\gamma}_{xy,i}^2 > \bar{\gamma}_{xy,j}^2$, then $e^{\bar{\gamma}_{xy,i}^2} > e^{\bar{\gamma}_{xy,j}^2}$. Since both have the same denominator:

$$\mathbf{L}_i = \frac{e^{\bar{\gamma}_{xy,i}^2}}{\sum_k e^{\bar{\gamma}_{xy,k}^2}} > \frac{e^{\bar{\gamma}_{xy,j}^2}}{\sum_k e^{\bar{\gamma}_{xy,k}^2}} = \mathbf{L}_j \quad (21)$$

□

Theorem 4: The arithmetic mean of pseudolabels \mathbf{L} satisfies standard averaging bounds:

$$\min_{i \in \{\alpha, \beta, \gamma\}} \mathbf{L}_i \leq \mathbf{L} \leq \max_{i \in \{\alpha, \beta, \gamma\}} \mathbf{L}_i \quad (22)$$

Proof. Let $m = \min_i(\mathbf{L}_i)$ and $M = \max_i(\mathbf{L}_i)$. Then:

$$\begin{aligned} \mathbf{L} &= \frac{1}{N_b} \sum_i \mathbf{L}_i \\ &\geq \frac{1}{N_b} \sum_i m = \frac{1}{N_b} \cdot N_b \cdot m = m \\ &\leq \frac{1}{N_b} \sum_i M = \frac{1}{N_b} \cdot N_b \cdot M = M \end{aligned} \quad (23)$$

1134 Since each $\mathbf{L}_i \in [0, 1]$, we also have:

$$0 \leq \mathbf{L} \leq 1 \quad (24)$$

□

D EXPERIMENTAL DETAILS AND HYPERPARAMETERS

All experiments were conducted with PyTorch 2.6 (Paszke et al., 2019) on a single Nvidia GeForce RTX 3090 GPU and 16-core CPU. LLMs were used as tools to aid or polish coding and writing.

The hyperparameter choices that are the same in all data sets are shown in Table 8. The hyperparameter choices that are dataset-specific are shown in Table 9. Since the numbers of embeddings vary here, the resulting numbers of parameters also vary across datasets. The cross entropy loss function is used for STEW, the binary cross entropy loss function is used for TUAB, the weighted cross entropy loss function is used for TUEV, and the focal loss function is used for CHB-MIT. All loss function choices are consistent with previous works except for TUEV. To demonstrate the effect of the loss function choice on TUEV, a detailed examination of the choice of loss functions for TUEV is included in Fig. 4

Table 8: Default hyperparameter settings

Parameter	pre-train	Supervised
Batch size	128	
Learning schedule	Cosine	
Optimizer	AdamW	
Weight decay	$1e^{-4}$	
Masking ratio	0.3	
Head	8	
Learning rate	$2e^{-3}$	$5e^{-4}$
Dropout	0.1	0.3

Table 9: Dataset-specific settings

Parameter	STEW	TUEV	TUAB	CHB-MIT
Epochs	100	60	60	60
Patience	20	5	5	5
Temperature (τ)	1.2	1.0	1.0	1.0
Sampling rate	128	250	200	256
Patch size	32	25	125	128
Embedding size	32	32	128	128
Learning rate (fine tune)	$5e^{-5}$	$5e^{-4}$	$5e^{-5}$	$5e^{-5}$
Params	0.05M	0.05M	0.8M	0.8M

E EVALUATION METRICS

To evaluate the performance of our method, we employed a set of different metrics, particularly in the context of potentially imbalanced datasets. The following metrics were used for our evaluation:

- **Accuracy** is defined as the ratio of correctly predicted samples to the total number of samples, which is a straightforward and intuitive metric. However, it can be misleading in scenarios with class imbalance, where a model might achieve high accuracy by simply predicting the majority class.
- **Balanced Accuracy (B-Accuracy)** is an alternative to the standard accuracy, which can be particularly useful for imbalanced datasets. It is defined as the arithmetic mean of sensitivity (true positive rate) and specificity (false positive rate), giving equal weight to the

1188
 1189
 1190
 1191
 1192
 1193
 1194
 1195
 1196
 1197
 1198
 1199
 1200
 1201
 1202
 1203
 1204
 1205
 1206
 1207
 1208
 1209
 1210
 1211
 1212
 1213
 1214
 1215
 1216
 1217
 1218
 1219
 1220
 1221
 1222
 1223
 1224
 1225
 1226
 1227
 1228
 1229
 1230
 1231
 1232
 1233
 1234
 1235
 1236
 1237
 1238
 1239
 1240
 1241

performance on both positive and negative classes to mitigate the bias toward the majority class.

- **AUROC** is the acronym for the area under the receiver operating characteristic (ROC) curve that measures a model’s ability to distinguish between binary classes. The ROC curve plots sensitivity against specificity, which is a robust metric against class imbalance.
- **Weighted F1 Score (W-F1)** is an extension of F1 score, a harmonic mean of precision and recall. It extends F1 score by calculating the F1 score for each class independently and then averaging them, weighted by the number of true instances for each class. This ensures that the contribution of each class to the final score is proportional to its size, making it a reliable metric for datasets with multi-classes.

F MORE ABLATIONS

Since the accuracy difference between our band selection design and the no-band-limit setting is small (0.3%) in Table 5, we conduct an additional band selection analysis on the CHB-MIT dataset, as shown in Table 10, to further justify our band selection design. Following the same protocol as in Table 5, we report the average of linear probing and fine-tuning scores. The difference on CHB-MIT is notably larger (3.5%). Although our design does not yield significant advantage across every dataset, the fact that it achieves performance gain even without lower-frequency information in pre-training is a key finding that opens promising directions for future research.

The significant improvement observed on CHB-MIT can be explained by the presence of focal high-frequency oscillations sometimes detected on the scalp in epilepsy contexts. Learning localized spatial patterns allows the model to capture these very high-frequency, spatially confined activities, rather than broadly coherent ones observed at lower frequencies (Besio et al., 2014; Chen et al., 2021). This also helps explain why our approach outperforms foundation models on CHB-MIT. Nevertheless, further investigation is needed to better characterize the dataset conditions under which our band selection design provides significant improvements.

Table 10: Band selection analysis on average performance (%)

Method	STEW	CHB-MIT
	Accuracy	B-Accuracy
Our design	76.1	86.1
No band limit (All Hz)	75.8 (↓0.3)	82.6 (↓3.5)

In our framework, we adopt a default decoder with attention pooling as the main configuration. Although the result with average pooling on STEW (“no attention pool”) is already reported in Table 5, additional results are provided here to show that our SPR framework learns robust and efficient representations and improves supervised performance regardless of the decoder choice. We conduct further experiments using a decoder with global average pooling and a decoder with learned weighted pooling on the STEW and CHB-MIT datasets, and compare these three common decoder designs in Table 11.

For each decoder, we first report the performance obtained with purely supervised training, followed by the performance when combined with our SPR framework. The accuracy gains are 3%, 4.4%, and 5.4% on STEW, and 34.9%, 34.9%, and 32% on CHB-MIT, respectively. These results indicate that our method consistently improves each supervised baseline by a substantial margin. This ablation on the decoders clearly demonstrates the effectiveness of our SSL approach, which does not rely on the specific capacity or design of the downstream model architecture.

It is also informative to keep the encoder fixed and replace our SSL task with traditional SSL tasks. We report results for RP (Banville et al., 2021), SimCLR (Chen et al., 2020), and VICReg (Bardes et al., 2021) on STEW in Table 12. Our encoder preserves the C channels rather than collapsing them into a single channel, but these SSL methods require a single-channel representation. Therefore, it is necessary to introduce an additional pooling module to reduce the C channels to one. To make comprehensive comparisons, we also analyze different choices for this pooling operation.

1242

1243

1244

1245

1246

1247

1248

1249

1250

1251

1252

1253

1254

1255

1256

1257

1258

1259

1260

1261

1262

1263

1264

1265

1266

Table 11: Performance (%) of SPR with different decoders

Method	STEW		CHB-MIT	
	Accuracy	AUROC	B-Accuracy	AUROC
Attention (Default)	75.3(0.9)	83.0(0.8)	54.7(1.2)	76.3(4.1)
+SPR	78.3(1.9)	86.3(2.5)	89.6(0.8)	94.1(1.4)
Global average	74.3(0.4)	82.2(0.5)	53.9(1.5)	76.6(2.7)
+SPR	78.7(0.7)	86.6(0.4)	88.8(1)	95.1(1.1)
Learned weighted	72.4(1)	80.1(1.1)	55.6(1.9)	74.8(2.7)
+SPR	77.8(1.5)	85.7(1.5)	87.6(0.9)	92.9(1)

These results clearly show that our encoder is not intrinsically more powerful than those used in prior SSL work. The supervised method obtains 75.3% accuracy and 83.0% AUROC in Table 4, whereas all traditional SSL methods underperform this baseline. The best variant (SimCLR with average pooling) achieves 74.9% accuracy and 80.4% AUROC. When keeping the encoder fixed and simply replacing our task with other SSL tasks, performance is worse than that of the supervised baseline.

These methods degrade performance because they require a single-channel encoder design and our encoder is not suitable for their purposes. Previous SSL approaches focus on single-channel objectives, while our multi-channel SSL objective explicitly captures spatial information, which makes it essential to use a multi-channel-in, multi-channel-out encoder in our framework. However, the results show that this encoder choice alone is not the main driver of our SOTA gains.

Table 12: Performance (%) of different SSL methods on STEW

Method	Attention		Average		Weighted	
	Accuracy	AUROC	Accuracy	AUROC	Accuracy	AUROC
RP	66.5(4.4)	72.0(6.1)	69.1(1.9)	75.8(2.2)	66.9(3.3)	72.4(3.5)
SimCLR	70.1(4.6)	75.1(5.0)	74.9(1.1)	80.4(2.2)	72.0(4.1)	77.6(3.9)
VICReg	71.0(3.5)	76.1(3.9)	71.0(4.5)	77.7(2.9)	72.0(3.9)	78.0(4.4)
SPR	Accuracy: 78.3(1.9)		AUROC: 86.3(2.5)			

1273

1274

1275

1276

1277

1278

1279

1280

1281

1282

1283

1284

1285

1286

1287

1288

1289

1290

1291

1292

1293

1294

1295