LEARNING WHAT TO REMEMBER: SELECTIVE MEMORY FOR LONG-CONTEXT EEG REPRESENTATION LEARNING

Anonymous authors
Paper under double-blind review

ABSTRACT

Long EEG recordings contain brief informative events within extended background activity. Learning from longer inputs therefore requires a representation that preserves useful evidence as context accumulates. We introduce SPARK, an encoder combining recurrent memory across temporal chunks, selective compression, and utility-guided controllers. Self-supervised pretraining learns to propagate and summarize context; subsequent utility adaptation refines the representation using curated diagnostic labels and physiological objectives. Under matched 120-second pretraining, SPARK reaches $87.67 \pm 0.56\%$ recording-level balanced accuracy on CHB-MIT, compared with $76.19 \pm 1.76\%$ for a near-parameter-matched dense-attention control. Utility adaptation raises performance to $92.14 \pm 1.00\%$, exceeding a cross-entropy control with shared auxiliary objectives and matched adaptation data and updates ($90.57 \pm 0.31\%$). Evaluation on three recording-level datasets and ten short-window benchmarks shows the strongest adaptation gain on CHB-MIT, alongside task-dependent context preferences and transfer behavior. These results identify recurrent memory and utility adaptation as complementary contributors to long-context EEG representation learning.

1 INTRODUCTION

Long EEG recordings combine extended background activity with brief intervals that can determine a clinical decision. A useful representation must retain such intervals while also capturing their surrounding context. During unlabeled pretraining, their locations are unknown: an event-centered short crop assumes information that a screening system does not yet have. Longer inputs make more context available, but a compact embedding must still decide which information survives compression. This motivates studying long-context EEG learning as a problem of selective evidence retention.

Related work and the open question. EEG foundation models have made substantial progress through large-scale pretraining. BIOT accommodates heterogeneous biosignal inputs Yang et al. (2023a); LaBraM, CBraMod, REVE, and LUNA develop reconstruction-based objectives and architectures for transferable EEG representations Jiang et al. (2024); Wang et al. (2025); El Ouahidi et al. (2025); Döner et al. (2025). Reconstruction can learn both local morphology and temporal structure, but uniform reconstruction error does not explicitly prioritize a rare interval according to its downstream value. We refer to this potential mismatch as a *reconstruction trap*: faithfully modeling abundant background does not necessarily ensure that the evidence most useful for a later decision survives compression.

SPARK addresses this problem through recurrent evidence accumulation and utility-directed adaptation. Recurrent memory carries a compact working state through ordered chunks, allowing later context to update earlier evidence. A utility objective then shapes the representation toward diagnostic separation and physiological structure. This separates learning to accumulate context from learning which properties to emphasize for downstream use.

Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

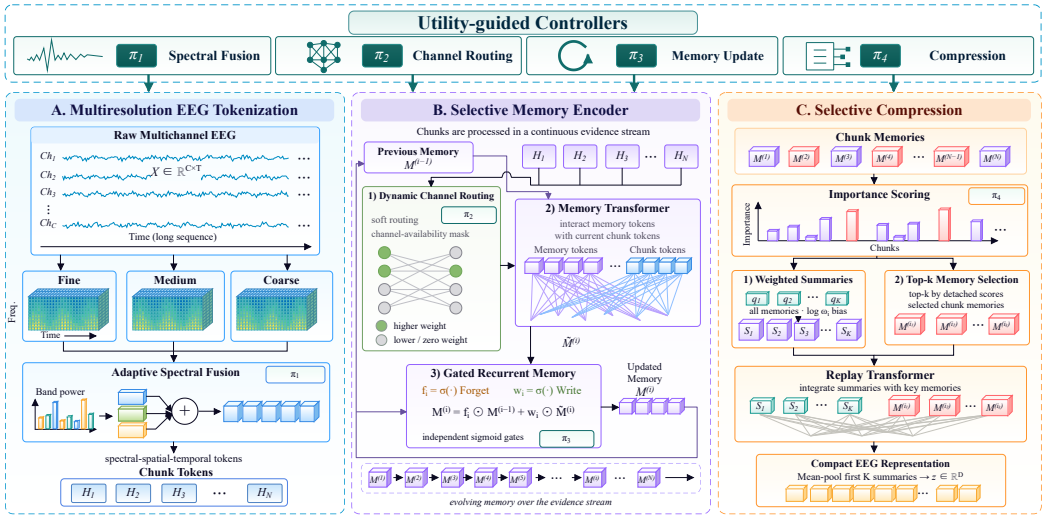


Figure 1: **SPARK with utility-guided controllers.** (A) Controller π_1 modulates spectral fusion. (B) Controller π_2 routes channels, and π_3 modulates independent memory gates. (C) Controller π_4 biases importance-weighted summaries and replay selection. Recurrence operates across chunks within an input window; compression produces the window embedding.

We introduce SPARK, which combines a multi-resolution spectral frontend, a recurrent memory encoder, and importance-weighted compression with replay. Stage 1 learns recurrent propagation and compression from unlabeled EEG. Stage 2 activates four deterministic controllers for spectral fusion, channel routing, memory updates, and compression. Their continuous outputs are optimized by backpropagation using diagnostic margins, physiological feature preservation, representation consistency, and boundary continuity. Estimators are recalibrated between adaptation rounds to account for the changing representation.

Our evaluation connects these design choices to three findings. First, the recurrent Stage 1 encoder outperforms a near-parameter-matched dense-attention control by 11.48 percentage points on CHB-MIT under matched 120-second pretraining. Second, utility adaptation adds 4.47 points over Stage 1 and exceeds conventional CE adaptation with shared auxiliaries by 1.57 points under matched data and update budgets. Reward-component ablations further show complementary downstream benefits from the utility terms. Third, evaluation across clinical, sleep, affective, and motor-imagery tasks reveals how context duration and controller use shape transfer: CHB-MIT benefits most from 120-second pretraining, whereas TUAB and TUSZ favor 60 seconds. Together, these experiments show how recurrent accumulation and utility-directed refinement improve EEG representations while exposing the task dependence of adaptation.

2 Methodology

2.1 Overview

Given an EEG input window $\mathbf{X} \in \mathbb{R}^{C\times T}$, SPARK produces an embedding $\mathbf{z} \in \mathbb{R}^D$ for downstream classification. Figure 1 shows the frontend, recurrent encoder, and compression module. Controllers π_1–π_4 modulate these computations using deterministic continuous outputs. Stage 1 learns recurrent propagation and compression with neutral controls; Stage 2 adapts the representation using a supervised utility objective. Memory is reset between independently evaluated windows, whose predicted probabilities are averaged for recording-level classification.

2.2 Multi-resolution EEG tokenization

Three short-time Fourier transforms provide fine, medium, and coarse spectral views. Each view retains its temporal axis, is projected to dimension D, and is aligned to a common token grid,

yielding $\mathbf{Z}^{(s)}$. The spectral controller takes five log band-power statistics $\mathbf{p} \in \mathbb{R}^5$ computed over the input window: delta, theta, alpha, beta, and low gamma. A $5 \to 64 \to 3$ MLP with a GELU hidden activation produces a bias for the frontend fusion logits $\boldsymbol{\ell}_{\text{front}} \in \mathbb{R}^3$:

$$\boldsymbol{\alpha} = \text{softmax}(\boldsymbol{\ell}_{\text{front}} + \pi_1(\mathbf{p})), \qquad \mathbf{Z} = \sum_s \alpha_s \mathbf{Z}^{(s)}. \tag{1}$$

The weights apply to the whole window; temporal variation remains in the spectral tokens. We add temporal, channel, and channel–time interaction embeddings and mask unavailable channels on a canonical montage. The sequence is split into N temporal chunks with segment embeddings, $\mathbf{H}_i = \text{Chunk}(\mathbf{Z}, i) + \text{SegEmb}(i)$, for $i = 1, \ldots, N$.

2.3 Selective memory architecture

Channel routing. Within each chunk, channel features receive continuous importance weights $\boldsymbol{\gamma}_i$. Controller π_2 modulates the softness of this distribution, and the availability mask excludes missing channels. We use $\mathbf{H}_i$ below for the resulting routed tokens.

Recurrent memory. The encoder jointly processes the current chunk and previous memory $\mathbf{M}^{(i-1)} \in \mathbb{R}^{M \times D}$:

$$[\widetilde{\mathbf{M}}^{(i)}; _] = \text{TransformerEnc}([\mathbf{M}^{(i-1)}; \mathbf{H}_i]). \tag{2}$$

The first M outputs form a candidate memory. Let $\mathbf{q}_i$ concatenate mean-pooled previous-memory and current-chunk features. Controller $\pi_3(\mathbf{q}_i)$ supplies biases Δ_i^f, Δ_i^w to independent sigmoid gates:

$$\begin{aligned} \mathbf{f}_i &= \sigma(\mathbf{W}_f \mathbf{q}_i + \mathbf{b}_f + \Delta_i^f), \qquad \mathbf{w}_i = \sigma(\mathbf{W}_w \mathbf{q}_i + \mathbf{b}_w + \Delta_i^w), \\ \mathbf{M}^{(i)} &= \mathbf{f}_i \odot \mathbf{M}^{(i-1)} + \mathbf{w}_i \odot \widetilde{\mathbf{M}}^{(i)}. \end{aligned} \tag{3}$$

Independent gates permit simultaneous retention and writing without requiring their weights to sum to one.

Compression and replay. After the last chunk, the model scores the stored memories:

$$s_i = \text{ScoreMLP}(\bar{\mathbf{M}}^{(i)}) + \Delta_i^{\text{comp}}, \qquad \boldsymbol{\omega} = \text{softmax}(\mathbf{s}), \tag{4}$$

where $\bar{\mathbf{M}}^{(i)}$ is the mean of the M memory tokens and π_4 supplies Δ_i^{comp}. Each of K learned queries attends to all memory tokens with attention-logit bias $\log \omega_i$ for tokens from chunk i. This differentiable path produces K summaries. In parallel, detached importance scores select the top-k chunk memories. A replay Transformer processes the concatenated summaries and selected memories; its first K outputs are mean-pooled into $\mathbf{z}$. Gradients pass through the soft summary path and selected memory contents, while the discrete selection indices are detached.

More explicitly, let j index the tokens in the concatenated memory history, $c(j)$ identify the originating chunk, and ℓ_{rj} denote the unmodulated attention logit between summary query r and token j. The weighted summary uses

$$A_{rj} = \frac{\exp(\ell_{rj})\, \omega_{c(j)}}{\sum_{j'} \exp(\ell_{rj'})\, \omega_{c(j')}}, \qquad \mathbf{S}_r = \sum_j A_{rj} \mathbf{V}_j, \quad r = 1, \ldots, K, \tag{5}$$

where $\mathbf{V}_j$ is the memory token's attention value. Thus, compression combines query-specific relevance with a shared chunk-level importance prior. The soft path can route gradients to every stored chunk; the replay path supplies the selected memory contents to the final Transformer. With neutral $\boldsymbol{\omega}$, the summary attention remains query-dependent, so disabling the controllers preserves the encoder's learned attention and memory computations.

The recurrent state occupies $O(MD)$ storage. Compression retains all N chunk memories, requiring $O(NMD)$ history storage in addition to the encoder's activations. The default configuration uses $M = 8$, $D = 256$, $K = 32$, and $k = 5$, with 20 chunks per 120-second pretraining window (Appendix E.3).

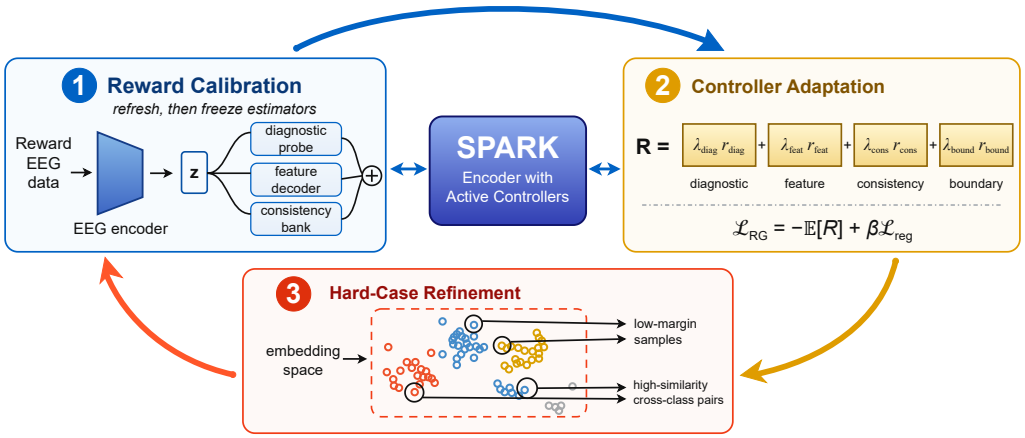


Figure 2: **Alternating recalibration and adaptation.** Utility estimators are recalibrated, frozen during encoder adaptation, and refreshed using mined hard cases in the next round.

2.4 Stage 1: Self-Supervised Initialization

Stage 1 uses unlabeled EEG and neutral controller values: uniform spectral weights, uniform routing over available channels, zero memory-gate biases, and uniform compression weights. The learned memory gates and recurrent encoder remain operative. Token masking is applied after spectral projection and fusion. The objective $\mathcal{L}_{\text{SSL}} = \mathcal{L}_{\text{mask}} + \lambda_{\text{tc}}\mathcal{L}_{\text{tc}}$ combines masked spectral-token reconstruction with mean-squared prediction of the next chunk's mean representation. This stage produces the SSL checkpoint $\theta^{(0)}$.

2.5 Stage 2: Supervised Utility Adaptation

The utility pool contains 4,786 normal/abnormal windows from TUSZ training data, OpenNeuro AD Miltiadous et al. (2023), and Siena, excluding downstream validation and test subjects. It is partitioned into 4,154 training, 240 validation, and 392 audit windows. Only the training subset is used to fit utility estimators, update the encoder, and mine hard cases. Its windows are mixed with unlabeled TUH windows at approximately 1:3. For labeled inputs, the utility score is

$$R = \lambda_{\text{diag}} r_{\text{diag}} + \lambda_{\text{feat}} r_{\text{feat}} + \lambda_{\text{cons}} r_{\text{cons}} + \lambda_{\text{bound}} r_{\text{bound}}. \tag{6}$$

A calibrated probe supplies the correct-class logit margin r_{diag}; a decoder of intermediate features scores band-power preservation through r_{feat}; a class-conditioned representation bank supplies r_{cons}; and neighboring chunk boundaries supply r_{bound}. The diagnostic term is a logit-margin objective,

$$r_{\text{diag}}(\mathbf{z}, y) = g_\phi(\mathbf{z})_y - \max_{c \neq y} g_\phi(\mathbf{z})_c. \tag{7}$$

This objective increases separation from the strongest competing class. We assess its effect through downstream classification and a held-out audit that tracks diagnostic-probe BAcc, feature recovery, and consistency (Section 3.3).

Unlabeled inputs use only feature and boundary terms with renormalized weights. Appendix E.2 defines these terms and their weights.

Training alternates three operations (Figure 2). First, the current encoder provides features for fitting the diagnostic probe and feature decoder and rebuilding the consistency bank. Second, these estimators are frozen while the model minimizes $\mathcal{L}_{\text{RG}} = -\mathbb{E}[R] + \beta\mathcal{L}_{\text{reg}}$. The regularizer penalizes squared deviations from neutral controls. Backpropagation updates controllers and adapters; the backbone is frozen in Cycle 1 and progressively unfrozen in later rounds. Third, low-margin examples and similar cross-class pairs are mined to reweight probe calibration and refresh the next consistency bank. Algorithm 1 specifies the schedule.

Downstream controller modes. `Off` overrides controls with the neutral values above. `Active` enables the learned controls while freezing controller parameters. `Finetune` also updates controller parameters using the downstream objective. For the recording-level comparison, Stage 1 uses $\theta^{(0)}$ with `Off`, and full SPARK uses the Cycle 2 checkpoint with `Active`. Short-window results retain the task-specific modes of their original evaluation campaign. Applying `Off` to a Stage 2 checkpoint preserves the other weights learned during adaptation and therefore defines a different condition from Stage 1 alone.

3 EXPERIMENTS

3.1 EVALUATION SETUP

Stage 1 used unlabeled TUH EEG resampled to 200 Hz and mapped to 19 channels. The Stage 1 checkpoint was pretrained on 120 s windows for 30 epochs, corresponding to 17,040 optimizer updates. Stage 2 uses the subject-partitioned labeled pool described in Section 2.5, with reward-validation data selecting Cycle 2. Appendices E and G give full training and audit settings.

We evaluated recording-level classification on CHB-MIT, TUAB, and TUSZ, and short-window transfer on ten benchmarks spanning clinical, sleep, affective, and motor-imagery tasks. Models were fine-tuned on task-specific windows. Recording scores average probabilities over non-overlapping windows; each seed selects a validation threshold and fixes it for testing. CHB-MIT metrics are computed per subject and averaged; threshold selection maximizes subject-averaged balanced accuracy (BAcc; Appendix G.3). The 120 s setting specifies pretraining context, while downstream windows follow each dataset's protocol. CHB-MIT follows CBraMod/BIOT-aligned preprocessing Wang et al. (2025); Yang et al. (2023a). Appendix F.6 gives dataset partitions and preprocessing, including SEED's session partition.

Recording results report mean $\pm$ sample SD over downstream seeds 0, 1, and 2, with one fixed pretrained checkpoint per model or context length. Stage 1 uses neutral controllers (`Off`); full SPARK uses the Cycle 2 checkpoint with frozen enabled controllers (`Active`). Short-window comparisons also report mean and SD; dataset-specific preprocessing and controller modes are given in the appendix.

3.2 RECORDING PERFORMANCE, MEMORY, AND CONTEXT LENGTH

SPARK achieved the highest recording BAcc among the evaluated models on all three datasets (Table 1). On CHB-MIT, Stage 1 reached $87.67 \pm 0.56\%$, 7.25 percentage points above REVE ($80.42 \pm 0.55\%$), while full SPARK reached $92.14 \pm 1.00\%$, 4.47 points above Stage 1. The paired full-minus-Stage 1 difference was 4.47 ± 1.38 points across downstream seeds; mean differences on TUAB and TUSZ were 0.42 and 0.74 points. Thus, the largest CHB-MIT performance difference was already present after self-supervised Stage 1, with a further increase after utility-supervised adaptation.

Matched architecture controls. We compared two non-recurrent 120 s models with Stage 1 using the same TUH windows, local encoder, self-supervised objective, and 17,040 optimizer updates. Flat-120 independently encodes twenty 6 s chunks and averages their summaries. Dense-Attn-120 adds a five-layer global Transformer over the ordered summaries. Dense-Attn-120 has 10.81M parameters, close to Stage 1's 11.11M; Flat-120 has 6.70M. Dense attention improved CHB-MIT BAcc over flat pooling by 12.17 points, but remained 11.48 points below recurrent Stage 1 (Figure 3a). Recurrent processing therefore improves on both uniform pooling and global attention over chunk summaries under the same local encoding and optimization protocol.

Preferred context length is task-dependent. We varied Stage 1 context from 30 to 60 to 120 s, using 5, 10, and 20 chunks and effective batches of 512, 256, and 128 to match optimizer updates and cumulative EEG exposure. CHB-MIT BAcc increased from 78.45 to 82.12 to 87.67%. TUAB peaked at 60 s (83.62% versus 82.07% at 120 s), as did TUSZ (67.34% versus 63.39%). Extended context was most beneficial on CHB-MIT; the other two tasks favored the intermediate duration. Appendix G.1 gives the complete settings and scores.

Table 1: **Recording-level BAcc (%).** Mean ± sample SD over three downstream seeds and one pretrained checkpoint per condition. CHB-MIT BAcc is averaged across subjects. Blue and green rows distinguish self-supervised Stage 1 and utility-adapted full SPARK. The bottom row reports paired seed differences in percentage points (mean ± SD). Appendix G.8, Table 26, provides BAcc, AUROC, and PR-AUC with three-seed means and SDs.

Model / adaptation	CHB-MIT	TUAB	TUSZ
PRETRAINED EEG BASELINES			
BIOT	72.71 ±0.52	79.52 ±0.50	60.97 ±0.55
LaBraM	76.88 ±0.58	80.33 ±0.62	61.42 ±0.60
CBraMod	77.51 ±0.41	77.90 ±0.45	61.93 ±0.42
REVE	80.42 ±0.55	81.28 ±0.55	62.68 ±0.55
SPARK Stage 1 SSL	**87.67** ±0.56	**82.07** ±0.40	**63.39** ±0.53
Full SPARK utility adaptation	**92.14** ±1.00	**82.49** ±0.60	**64.13** ±0.65
Paired full − Stage 1 (pp)	**+4.47** ±1.38	**+0.42** ±1.00	**+0.74** ±1.18

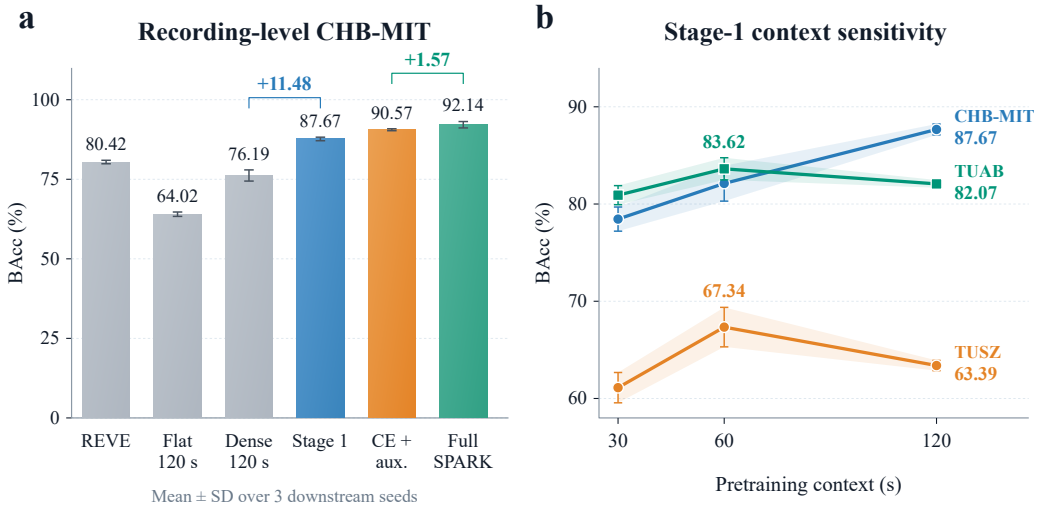


Figure 3: **Memory, adaptation, and pretraining context.** (a) CHB-MIT recording BAcc, averaged across subjects. Flat-120, Dense-Attn-120, and Stage 1 share pretraining data and update count. CE with shared auxiliaries and full SPARK share Stage 1 initialization, adaptation data, and adaptation-update budget (Appendix G.4); REVE is a reference baseline. (b) Stage 1 context scaling with 6 s chunks, matched updates and EEG exposure. Error bars and shading indicate sample SD over three downstream seeds, with one pretrained checkpoint per condition.

3.3 What does utility-supervised adaptation contribute?

Comparison with supervised adaptation. Utility adaptation outperformed a cross-entropy (CE) control initialized from the same Stage 1 checkpoint and trained with the same adaptation data, shared auxiliary objectives, and adaptation-update budget. On CHB-MIT, CE achieved $90.57 \pm 0.31\%$ recording-level BAcc, whereas full SPARK achieved $92.14 \pm 1.00\%$ (Figure 3a). The paired gain was 1.57 ± 0.71 percentage points across three downstream seeds, with higher AUROC and PR-AUC as well (Appendix G.4).

Reward components and sources. In the matched short-window campaign, full reward exceeded feature-only adaptation by 5.29 points on CHB-MIT, 2.41 on TUAB, and 3.45 on Sleep-EDFx (Figure 4). Removing diagnostic reward produced the largest CHB-MIT decrease, from 78.49 to 69.18%. Appendix G.5 gives all component and source results, their campaign-specific references, and the separate boundary-reward comparison.

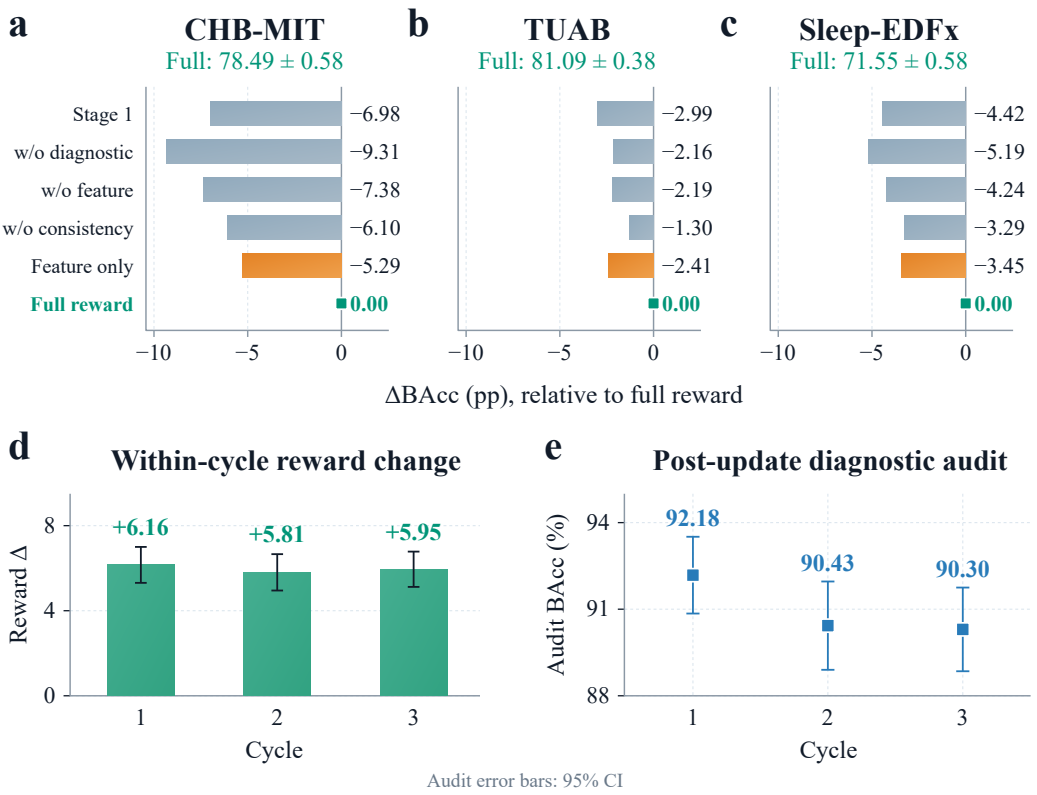


Figure 4: **Reward components and held-out utility audit.** (a–c) Short-window BAcc differences (variant minus full; pp) within one matched campaign, without error bars. Headers give full-reward mean ± SD; Appendix G.5 gives all three-seed summaries. (d) Within-cycle composite reward change. (e) Post-update diagnostic BAcc. Audit error bars: 95% CIs from 2,000 cluster-bootstrap resamples of 21 subjects (392 windows). Each cycle uses its own frozen estimators for pre/post comparisons.

The advantage over feature-only adaptation on five-class Sleep-EDFx also supports transfer beyond the binary normal/abnormal reward-label partition.

To examine reward-source dependence, we retrained Stage 2 with AD-only reward data, removing TUSZ and Siena while matching optimizer updates and labeled exposures by sampling AD windows with replacement. Relative to this campaign's full reward pool, paired short-window BAcc changed by $+0.50 \pm 0.10$ points on CHB-MIT, -1.35 ± 0.25 on TUSZ, and -0.48 ± 0.03 on Siena. Seizure-aligned reward sources therefore benefited TUSZ and Siena in this comparison, while CHB-MIT short-window transfer remained effective without them.

Held-out utility audit. The reward audit excluded 21 subjects and 392 windows from estimator fitting, controller adaptation, mining, and model selection. Composite reward increased within all three cycles under their respective frozen estimators: $+6.16$, $+5.81$, and $+5.95$. The accompanying post-update diagnostic BAcc values were 92.18, 90.43, and 90.30%, while feature R^2 was 0.584, 0.746, and 0.756; within-cycle consistency changes were negative in Cycles 2 and 3. The audit shows that adaptation improves the composite objective while trading off diagnostic accuracy, feature recovery, and consistency. Appendix G.6 reports the complete audit and separates within-cycle changes from the effects of estimator recalibration.

Table 2: **Short-window BAcc (%) across ten benchmarks.** Cells report mean ± SD. **Bold** and underlining mark the highest and second-highest means among the displayed model columns under the local task protocol; entries with different preprocessing or adaptation (†) are excluded from these rankings. The green column reports full SPARK after utility adaptation. Full protocols and metrics appear in Appendix G.7.

Dataset	BIOT	LaBraM	CBraMod	CSBrain	REVE	SPARK Full
MUMTAZ	93.58 ±0.52	94.09 ±0.79	95.60 ±0.56	96.43 ±1.55	96.44† ±0.97	**97.04** ±1.21
TUAB	79.59 ±0.57	81.40 ±1.25	78.91 ±1.43	81.72 ±1.56	83.15† ±0.14	**82.29** ±2.14
CHB-MIT	70.68 ±4.57	70.75 ±3.58	73.98 ±2.84	72.62 ±1.15	74.62 ±0.48	**79.22** ±1.85
STRESS	68.75 ±1.86	69.09 ±1.25	72.56 ±1.32	75.58 ±1.06	76.60† ±3.55	**77.00** ±1.98
Siena	73.52 ±6.69	70.82 ±3.29	73.17 ±6.47	76.62 ±4.71	78.41 ±0.45	**78.61** ±2.67
TUSZ	59.27 ±1.51	60.20 ±1.82	58.50 ±1.46	57.80 ±1.08	63.12 ±0.48	**65.46** ±3.45
Sleep-EDFx	66.40 ±0.70	67.70 ±3.00	67.10 ±1.23	65.40 ±2.33	67.88 ±0.62	**70.62** ±1.54
TUEV	52.81 ±2.25	66.16 ±1.70	66.71 ±1.07	69.03 ±0.59	67.59† ±2.29	**69.89** ±2.89
SEED	62.62 ±0.94	67.71 ±0.55	**72.25** ±1.90	71.92 ±0.39	71.34 ±0.55	72.02 ±1.24
BCIC-IV-2a*	47.48 ±0.93	48.69 ±0.85	51.38 ±0.66	56.57 ±0.71	**63.96** ±0.95	56.87 ±0.26

† REVE uses different preprocessing/adaptation on these four datasets El Ouahidi et al. (2025). * BCIC-IV-2a follows the motor-imagery protocol in Table 25; SPARK Stage 1 obtains $\mathbf{57.08 \pm 0.83}\%$ BAcc, 0.21 pp above full SPARK.

3.4 Short-window transfer and controller modes

Table 2 summarizes ten short-window benchmarks, including motor imagery. Comparators include BIOT Yang et al. (2023a), LaBraM Jiang et al. (2024), CBraMod Wang et al. (2025), CSBrain Zhou et al. (2025), and REVE El Ouahidi et al. (2025). SPARK exceeded REVE on CHB-MIT, Sleep-EDFx, SEED, Siena, and TUSZ under the aligned downstream settings. The Siena difference was small (78.61 versus 78.41%). REVE preprocessing and adaptation differ on MUMTAZ, TUAB, STRESS, and TUEV, so comparisons on these four tasks are descriptive. The motor-imagery results complement these clinical, sleep, and affective tasks by testing transfer to a different EEG paradigm.

On BCIC-IV-2a, Stage 1 achieved $57.08 \pm 0.83\%$ BAcc and full SPARK achieved $56.87 \pm 0.26\%$, compared with $63.96 \pm 0.95\%$ for REVE. Stage 1 exceeded BIOT, LaBraM, CBraMod, and CSBrain, demonstrating transfer from long-context pretraining to four-second motor-imagery trials. Clinical utility adaptation yielded similar BAcc and lower Cohen's κ and weighted F1 than Stage 1 (Appendix G.7).

The controller-mode comparison likewise showed task dependence (Figure 5). Frozen enabled controllers performed best in the reported comparisons on CHB-MIT, MUMTAZ, Siena, and Sleep-EDFx; downstream controller fine-tuning was preferable on TUAB, STRESS, and TUSZ, while neutral modulation was preferable on TUEV. These choices distinguish reusing the learned modulation, adapting it to a new label structure, and retaining the encoder with neutral controls. On TUEV, neutral modulation reaches 68.1% BAcc versus 61.4% with frozen enabled controllers, a 6.7-point gap. On TUAB, downstream controller fine-tuning improves on frozen enabled controllers by 1.2 points (82.9 versus 81.7%). These differences make controller mode a practical adaptation choice, selected on validation data for each task.

An example of spatial selection. In the TUSZ example (Figure 6), π_2 shifts from broadly distributed background weights to higher weights over F7, T3, and T5 during the seizure interval. The example links input-dependent routing to changing EEG activity. Appendix D shows the corresponding spectral preferences and compression patterns.

Scope and limitations. SPARK's strongest demonstrated gains concern long-context EEG pretraining and CHB-MIT recording classification. The matched CE comparison establishes an adaptation advantage on CHB-MIT; broader transfer depends on utility-task alignment. Recording inference averages independently scored windows, so persistent memory across multi-hour recordings remains untested. Three downstream seeds quantify fine-tuning variability for fixed checkpoints and partitions; CHB-MIT has two held-out test identities. P300, stimulus decoding, and a controlled LUNA comparison remain future evaluations. Appendix C discusses these settings and the observed utility trade-offs.

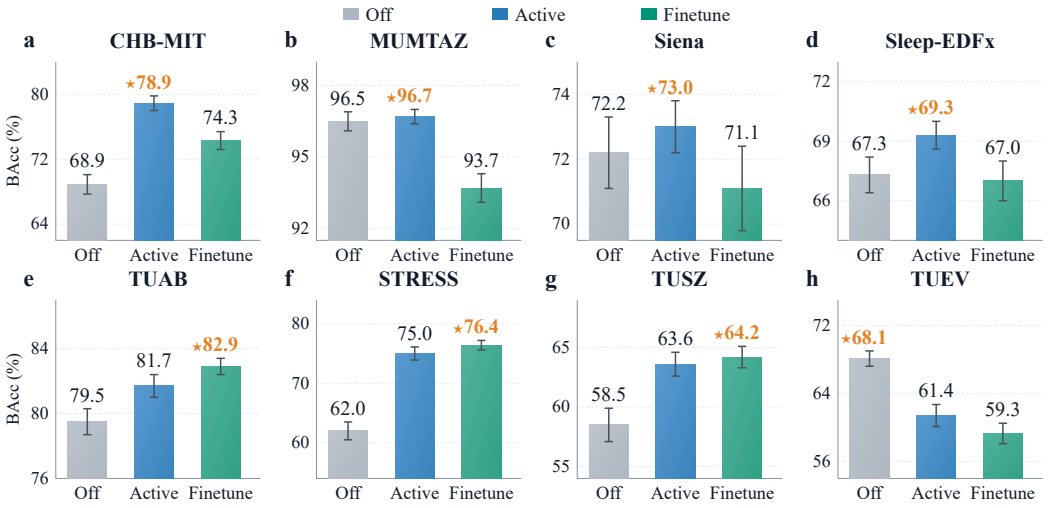


Figure 5: **Controller use depends on the task.** Short-window BAcc (mean ± SD) for neutral controllers (Off), frozen enabled controllers (Active), and downstream-adapted controllers (Finetune). Stars mark the highest mean for each task. Axes differ across tasks. Appendix G.9 gives all values from this separate mode-ablation campaign.

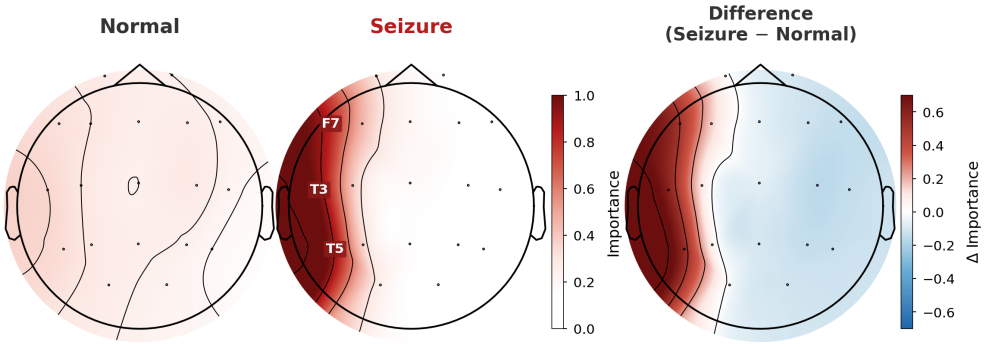


Figure 6: **Channel-routing example.** Background, seizure, and difference maps for the illustrated TUSZ windows. The spatial shift describes controller weights, without spatial localization supervision.

4 Conclusion

SPARK studies long-context EEG representation learning as a problem of selective evidence retention, combining recurrent evidence accumulation with utility-guided adaptation. Matched pretraining controls show that recurrent memory improves recording classification over uniform pooling and dense attention across chunk summaries, while a matched supervised control further isolates the benefit of utility-guided adaptation beyond conventional CE adaptation with shared auxiliary objectives. Context scaling shows that more temporal context is not uniformly beneficial: the preferred pretraining duration depends on the target task. Transfer experiments likewise show that utility adaptation is most effective when its learned preferences align with the downstream setting, with limited benefit on motor imagery and task-dependent controller modes. Together, these results suggest a practical division of roles: self-supervised selective memory builds representations that propagate and retain evidence across temporal chunks, while subsequent utility adaptation provides task-directed refinement. Effective long-context EEG learning therefore depends not only on how much context is available, but also on what information is retained and how the resulting representation is aligned with downstream use.

CODE OF ETHICS

This work uses publicly available EEG datasets and follows the corresponding dataset access and usage requirements. No new human-subject data were collected as part of this study. The proposed method is intended for research on EEG representation learning and has not been prospectively validated for clinical deployment. In particular, the controller weights and visualizations reported in this work should not be interpreted as clinical localization or diagnostic evidence. We discuss additional limitations concerning dataset coverage, evaluation protocols, and clinical generalization in the paper and appendix.

REPRODUCIBILITY STATEMENT

We provide the methodological and experimental details required to reproduce the main results in the paper and supplementary material. These include the model architecture, self-supervised and utility-adaptation objectives, pretraining and downstream hyperparameters, dataset preprocessing and partitions, checkpoint-selection procedures, and individual downstream-seed results. The appendix further reports matched architecture controls, context scaling, a supervised adaptation control, reward-component and reward-source ablations, held-out utility audits, and additional evaluation metrics. We intend to release the implementation and trained checkpoints to facilitate reproduction of the reported experiments.

REFERENCES

Aydar Bulatov, Yury Kuratov, and Mikhail Burtsev. Recurrent memory transformer. In S. Koyejo, S. Mohamed, A. Agarwal, D. Belgrave, K. Cho, and A. Oh (eds.), *Advances in Neural Information Processing Systems*, volume 35, pp. 11079–11091. Curran Associates, Inc., 2022. URL https://proceedings.neurips.cc/paper_files/paper/2022/file/47e288629a6996a17ce50b90a056a0e1-Paper-Conference.pdf.

Paolo Detti. Siena Scalp EEG Database. *PhysioNet*, August 2020. doi: 10.13026/5d4a-j060. URL https://doi.org/10.13026/5d4a-j060. Version 1.0.0.

Berkay Döner, Thorir Mar Ingolfsson, Luca Benini, and Yawei Li. LUNA: Efficient and topology-agnostic foundation model for EEG signal analysis. In D. Belgrave, C. Zhang, H. Lin, R. Pascanu, P. Koniusz, M. Ghassemi, and N. Chen (eds.), *Advances in Neural Information Processing Systems*, volume 38, Main Conference, pp. 70682–70708. Curran Associates, Inc., 2025. doi: 10.52202/085713-2377. URL https://proceedings.neurips.cc/paper_files/paper/2025/file/66969a9e6bd7a26dfeccea7227178ca7-Paper-Conference.pdf.

Yassine El Ouahidi, Jonathan Lys, Philipp Thölke, Nicolas Farrugia, Bastien Pasdeloup, Vincent Gripon, Karim Jerbi, and Giulia Lioi. REVE: A foundation model for EEG - adapting to any setup with large-scale pretraining on 25,000 subjects. In D. Belgrave, C. Zhang, H. Lin, R. Pascanu, P. Koniusz, M. Ghassemi, and N. Chen (eds.), *Advances in Neural Information Processing Systems*, volume 38, Main Conference, pp. 22541–22577. Curran Associates, Inc., 2025. doi: 10.52202/085713-0760. URL https://proceedings.neurips.cc/paper_files/paper/2025/file/20a917f77773ac0fa8bea2bdd6606b66-Paper-Conference.pdf.

John Guttag. CHB-MIT Scalp EEG Database. *PhysioNet*, June 2010. doi: 10.13026/C2K01R. URL https://doi.org/10.13026/C2K01R. Version 1.0.0.

Amir Harati, Meysam Golmohammadi, Silvia Lopez, Iyad Obeid, and Joseph Picone. Improved EEG Event Classification Using Differential Energy. In *2015 IEEE Signal Processing in Medicine and Biology Symposium (SPMB)*, pp. 1–4. IEEE, 2015. doi: 10.1109/SPMB.2015.7405421.

Wei-Bang Jiang, Li-Ming Zhao, and Bao-Liang Lu. Large brain model for learning generic representations with tremendous EEG data in BCI. In *International Conference on Learning Representations*, 2024. URL https://openreview.net/forum?id=QzTpTRVtrP.

Jin Jing, Wendong Ge, Shenda Hong, Marta Bento Fernandes, Zhen Lin, Chaoqi Yang, Sungtae An, Aaron F. Struck, Aline Herlopian, Ioannis Karakis, et al. Development of expert-level classification of seizures and rhythmic and periodic patterns during EEG interpretation. *Neurology*, 100 (17):e1750–e1762, 2023. doi: 10.1212/WNL.0000000000207127.

B. Kemp, A. H. Zwinderman, B. Tuk, H. A. C. Kamphuisen, and J. J. L. Oberyé. Analysis of a sleep-dependent neuronal feedback loop: The slow-wave microcontinuity of the EEG. *IEEE Transactions on Biomedical Engineering*, 47(9):1185–1194, 2000. doi: 10.1109/10.867928.

Vernon J. Lawhern, Amelia J. Solon, Nicholas R. Waytowich, Stephen M. Gordon, Chou P. Hung, and Brent J. Lance. EEGNet: A compact convolutional neural network for EEG-based brain–computer interfaces. *Journal of Neural Engineering*, 15(5):056013, 2018. doi: 10.1088/1741-2552/aace8c.

Hongli Li, Man Ding, Ronghua Zhang, and Chunbo Xiu. Motor imagery EEG classification algorithm based on CNN-LSTM feature fusion network. *Biomedical Signal Processing and Control*, 72:103342, 2022. doi: 10.1016/j.bspc.2021.103342.

Livia Livint Popa, Hanna Dragos, Cristina Pantelemon, Olivia Verisezan Rosu, and Stefan Strilciuc. The role of quantitative EEG in the diagnosis of neuropsychiatric disorders. *Journal of Medicine and Life*, 13(1):8–15, 2020. doi: 10.25122/jml-2019-0085.

Andreas Miltiadous, Katerina D. Tzimourta, Theodora Afrantou, Panagiotis Ioannidis, Nikolaos Grigoriadis, Dimitrios G. Tsalikakis, Pantelis Angelidis, Markos G. Tsipouras, Euripidis Glavas, Nikolaos Giannakeas, and Alexandros T. Tzallas. A dataset of scalp EEG recordings of alzheimer's disease, frontotemporal dementia and healthy subjects from routine EEG. *Data*, 8(6): 95, 2023. doi: 10.3390/data8060095. URL https://www.mdpi.com/2306-5729/8/6/95.

Wajid Mumtaz. MDD Patients and Healthy Controls EEG Data (New), November 2016. URL https://doi.org/10.6084/m9.figshare.4244171.v2.

Iyad Obeid and Joseph Picone. The Temple University Hospital EEG Data Corpus. *Frontiers in Neuroscience*, 10:196, 2016. doi: 10.3389/fnins.2016.00196.

Wei Yan Peh, Yuanyuan Yao, and Justin Dauwels. Transformer convolutional neural networks for automated artifact detection in scalp EEG. In *Proceedings of the 2022 44th Annual International Conference of the IEEE Engineering in Medicine & Biology Society*, pp. 3599–3602, 2022. doi: 10.1109/EMBC48229.2022.9871916.

Vinit Shah, Eva von Weltin, Silvia Lopez, James Riley McHugh, Lillian Veloso, Meysam Golmohammadi, Iyad Obeid, and Joseph Picone. The Temple University Hospital Seizure Detection Corpus. *Frontiers in Neuroinformatics*, 12:83, 2018. doi: 10.3389/fninf.2018.00083.

Yonghao Song, Xueyu Jia, Lie Yang, and Longhan Xie. Transformer-based spatial–temporal feature learning for EEG decoding. *arXiv preprint arXiv:2106.11170*, 2021. doi: 10.48550/arXiv.2106.11170.

Yonghao Song, Qingqing Zheng, Bingchuan Liu, and Xiaorong Gao. EEG conformer: Convolutional transformer for EEG decoding and visualization. *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, 31:710–719, 2023. doi: 10.1109/TNSRE.2022.3230250.

Michael Tangermann, Klaus-Robert Müller, Ad Aertsen, Niels Birbaumer, Christoph Braun, Clemens Brunner, Robert Leeb, Carsten Mehring, Kai J. Miller, Gernot R. Müller-Putz, Guido Nolte, Gert Pfurtscheller, Hubert Preissl, Gerwin Schalk, Alois Schlögl, Carmen Vidaurre, Stephan Waldert, and Benjamin Blankertz. Review of the BCI competition IV. *Frontiers in Neuroscience*, 6:55, 2012. doi: 10.3389/fnins.2012.00055.

Jiquan Wang, Sha Zhao, Zhiling Luo, Yangxuan Zhou, Haiteng Jiang, Shijian Li, Tao Li, and Gang Pan. CBraMod: A criss-cross brain foundation model for EEG decoding. In *International Conference on Learning Representations*, 2025. URL https://openreview.net/forum?id=NPNUHgHF2w.

Chaoqi Yang, M. Brandon Westover, and Jimeng Sun. BIOT: Biosignal transformer for cross-data learning in the wild. In *Advances in Neural Information Processing Systems*, volume 36, pp. 78240–78260, 2023a. doi: 10.52202/075280-3420. URL https://proceedings.neurips.cc/paper_files/paper/2023/hash/f6b30f3e2dd9cb53bbf2024402d02295-Abstract-Conference.html.

Chaoqi Yang, Cao Xiao, M. Brandon Westover, and Jimeng Sun. Self-supervised electroencephalogram representation learning for automatic sleep staging: Model development and evaluation study. *JMIR AI*, 2:e46769, 2023b. doi: 10.2196/46769.

Wei-Long Zheng and Bao-Liang Lu. Investigating Critical Frequency Bands and Channels for EEG-based Emotion Recognition with Deep Neural Networks. *IEEE Transactions on Autonomous Mental Development*, 7(3):162–175, 2015. doi: 10.1109/TAMD.2015.2431497.

Yuchen Zhou, Jiamin Wu, Zichen Ren, Zhouheng Yao, Weiheng Lu, Kunyu Peng, Qihao Zheng, Chunfeng Song, Wanli Ouyang, and Chao Gou. CSBrain: A cross-scale spatiotemporal brain foundation model for EEG decoding. In *Advances in Neural Information Processing Systems*, volume 38, pp. 87150–87195, 2025. doi: 10.52202/085713-2915. URL https://openreview.net/forum?id=agcXjEHmyW.

Igor Zyma, Ivan Seleznov, Anton Popov, Mariia Chernykh, and Oleksii Shpenkov. EEG During Mental Arithmetic Tasks, 2018. URL https://doi.org/10.13026/C2JQ1P. Version 1.0.0.

A LLM USAGE STATEMENT

Large language models were used as assistive tools during manuscript preparation, including language editing, improving clarity and organization, and assisting with LaTeX and code-related drafting. They were not used to generate experimental data, dataset labels, or reported numerical results. All scientific claims, methodological decisions, experimental procedures, analyses, and final manuscript content were reviewed and verified by the authors, who take full responsibility for the work.

Guide to the supplementary material. The links below connect the main claims to their implementation and supporting results.

Main-text item	Supporting material
Introduction (Section 1)	Related work: Appendix B.
Method and Figures 1–2	Implementation: Appendix E; algorithm: Algorithm 1.
Recording results (Table 1)	Protocol and seeds: Appendices G.1–G.2; additional metrics: Table 26.
Memory, adaptation, and context (Figure 3)	Architecture/context: Tables 17–18; CE control: Appendix G.4.
Reward evidence (Figure 4)	Component/source results: Appendix G.5; held-out audit: Appendix G.6.
Short-window transfer (Table 2)	Dataset metrics: Appendix F.6; BCIC-IV-2a: Table 25.
Controller modes (Figure 5)	Full means and SDs: Table 27.
Channel example (Figure 6)	Spectral and compression examples: Appendix D.
Limitations	Extended discussion: Appendix C.

B RELATED WORK

EEG representation learning. Large-scale EEG pretraining includes BIOT, LaBraM, CBraMod, and CSBrain Yang et al. (2023a); Jiang et al. (2024); Wang et al. (2025); Zhou et al. (2025). LaBraM, CBraMod, REVE, and LUNA use reconstruction-based objectives, with different tokenization and spatial architectures El Ouahidi et al. (2025); Döner et al. (2025). These objectives can learn useful local and temporal structure. Our controls examine how a recurrent architecture aggregates ordered chunks during pretraining, and how an additional supervised utility objective changes transfer. Quantitative EEG features provide clinically motivated targets Livint Popa et al. (2020); their preservation complements diagnostic supervision in the adaptation objective.

Memory and attention. Recurrent memory Transformers propagate state between segments Bulatov et al. (2022). Standard self-attention can integrate jointly presented chunks. SPARK combines a recurrent state with importance-weighted compression and replay. The Dense-Attn-120 control tests a concrete global-attention alternative at a similar parameter count. The recurrent state has fixed capacity, while compression stores the growing chunk history.

Utility-guided adaptation. The four SPARK controllers are deterministic differentiable modules optimized by backpropagation. Their supervision specifies diagnostic, spectral, consistency, and boundary preferences. Alternating estimator fitting and encoder adaptation keeps the estimator fixed within an update phase, but changes the proxy between cycles. Held-out auditing tracks these utility trade-offs, and downstream ablations measure their effects on classification.

C LIMITATIONS AND FUTURE WORK

Utility supervision and domain alignment. Stage 2 depends on curated labels and introduces additional optimization. The Stage 1 comparison measures the combined effect of that adaptation procedure. The matched CE control establishes an additional utility-adaptation gain on CHB-MIT (Appendix G.4); extending this comparison across tasks would characterize its transfer range. Removing seizure reward sources tests reward-source dependence, while broader reward-design bias remains possible. The motor-imagery results show no BAcc advantage from utility adaptation and lower scores on the other two metrics.

Proxy trade-offs. A higher composite reward can coexist with worse diagnostic accuracy or consistency. The held-out audit shows this trade-off on subjects excluded from fitting and adaptation. The window-level spectral controller also cannot express channel- or chunk-specific fusion weights. Localized spectral control and downstream-aligned utility calibration are directions for further study.

Context and uncertainty. The experiments test pretraining windows of up to 120 seconds and aggregate independent downstream window predictions for recording classification. They do not test persistent state across a multi-hour recording. The recurrent state has fixed capacity, but retained memory history requires storage linear in the number of chunks. Reported repeated runs use one pretrained checkpoint per condition and a fixed split. CHB-MIT excludes chb01 from training and uses two validation and two test identities. The small held-out sample limits conclusions about other subject assignments (Appendix G.3).

Evaluation scope. The evaluated clinical, sleep, affective, and motor-imagery tasks demonstrate varied transfer behavior. P300 and stimulus decoding remain untested, and there is no controlled LUNA reproduction. REVE preprocessing and adaptation are not fully matched on four datasets. Context duration and controller mode should therefore be selected on validation data for the target setting.

Preprocessing and deployment. The experiments use dataset-specific preprocessing and channel mappings. Transfer to new acquisition systems and clinical deployment requires evaluation under the corresponding montage, artifact, and patient conditions.

D Additional Ablation and Visualization Results

Adaptation cycles. Figure 7 reports the four-task cycle comparison. These point estimates are separate from the three-seed recording experiment in Table 1 and the mode comparison in Figure 5. Differences across cycles combine recalibration, continued optimization, and the backbone-unfreezing schedule. Cycle 2 is the validation-selected default. The within-cycle audit in Appendix G.6 evaluates utility changes under explicit estimator identities.

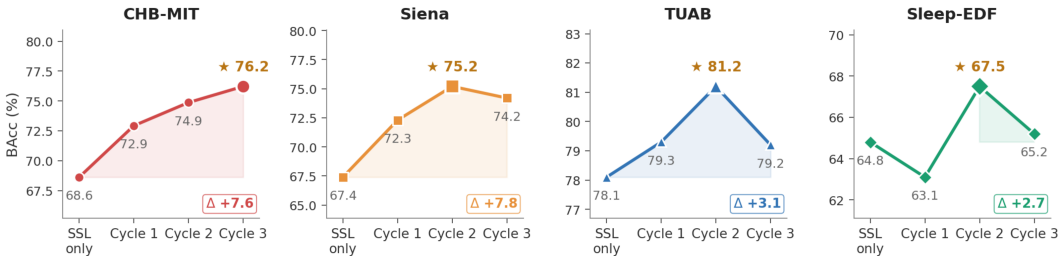


Figure 7: **Downstream performance across adaptation cycles.** Four-task point estimates from a separate cycle-ablation campaign.

Window-level spectral fusion. Figure 8 compares the three fusion weights in the illustrated seizure and background windows. In the heavy-seizure example, seizure windows assign more weight to the coarse branch (0.52 versus 0.42). In the sparse-seizure example, seizure windows assign more weight to the fine branch (0.34 versus 0.14), while background windows favor the medium branch (0.50). The two examples show different spectral preferences for seizure and background activity.

Compression weights and event annotations. Figure 9 places compression scores alongside seizure annotations and predicted probabilities for two recordings. The horizontal axis indexes separately processed window starts. Elevated weights occur near some seizure intervals, with imperfect alignment and nonzero background weights. The examples show how compression allocation varies with the input.

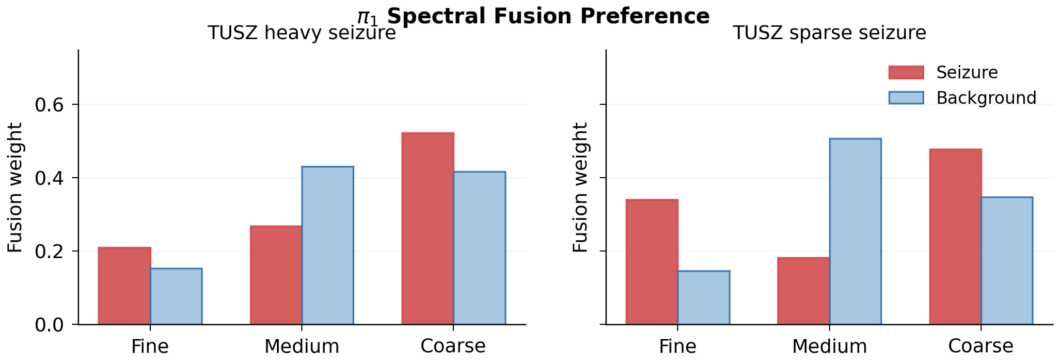


Figure 8: **Spectral-fusion weights for two TUSZ examples.** Each weight vector applies to an entire input window. The temporal token sequence is retained separately.

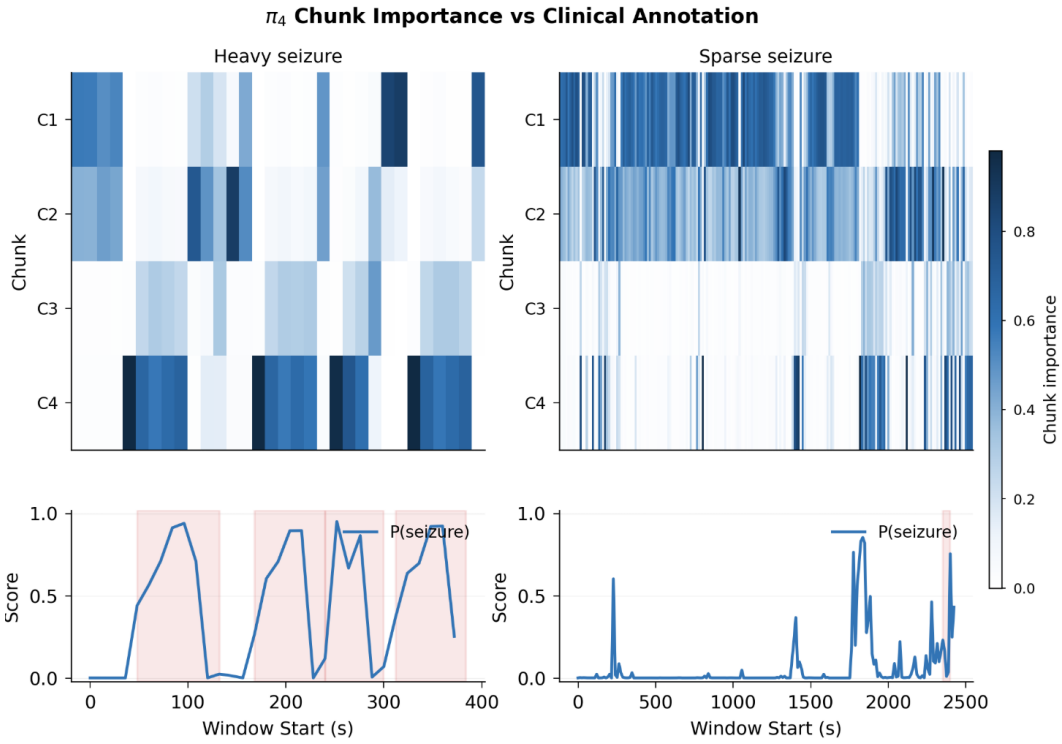


Figure 9: **Compression scores across independently processed windows.** Top: four displayed chunk scores per window; bottom: seizure probabilities and annotated seizure intervals. Memory is reset between windows.

Training reward traces. Figure 10 shows the training reward traces. Reward estimators are refreshed between cycles and component scales differ, so vertical shifts across cycles do not measure improvement under a single fixed evaluator. The feature term shifts from a negative to a positive regime after recalibration. The $\Delta + 1.3$ annotation summarizes this training trace. Figure 4 reports held-out within-cycle changes; component ablations in Appendix G.5 evaluate the downstream contribution of each term.

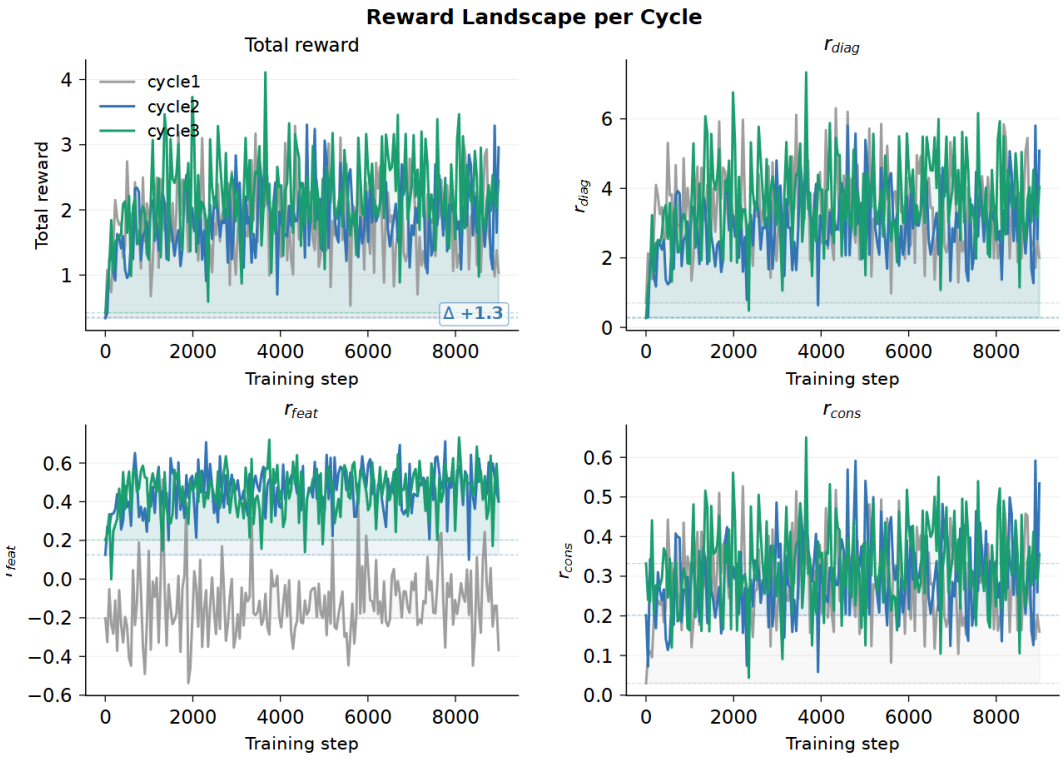


Figure 10: **Training rewards across adaptation cycles.** Training traces under the respective refreshed estimators. Shading is not used as an uncertainty estimate here. The boundary term is omitted in the original visualization.

E Pretraining Implementation Details

E.1 Pipeline Overview

The complete pretraining procedure contains two stages. In **Stage 1**, SPARK is pretrained on unlabeled long-context EEG using masked reconstruction and temporal consistency regularization, producing a general-purpose EEG encoder $\theta^{(0)}$. In **Stage 2**, the SSL-pretrained encoder is further adapted through an alternating reward-guided cycle. Each cycle alternates between reward calibration, selective controller adaptation, and hard-case refinement. The backbone configuration and Stage 1 SSL settings are summarized in Table 3, while the Stage 2 reward-guided optimization settings and runtime setup are summarized in Tables 4 and 5.

E.2 Reward Definitions and Controller Regularization

This appendix gives the explicit definitions of the four reward terms used during utility-supervised adaptation, together with the regularization term applied to the continuous controller outputs.

Diagnostic reward. Let $\mathbf{z} \in \mathbb{R}^D$ denote the pooled representation of the current input and let g_ϕ be the frozen diagnostic probe calibrated on the current representation space. For a labeled sample with class label y, the diagnostic reward is defined as the classification margin

$$r_{\text{diag}}(\mathbf{z}, y) = g_\phi(\mathbf{z})_y - \max_{c \neq y} g_\phi(\mathbf{z})_c. \tag{8}$$

Thus, r_{diag} is large when the correct class logit exceeds the strongest competing class logit by a wide margin.

Physiological reward. Let $\mathbf{u} \in \mathbb{R}^D$ denote the intermediate feature representation used for physiological decoding, and let d_ψ be the frozen feature decoder. Given target band-power features

$\mathbf{b} \in \mathbb{R}^5$, the decoder predicts

$$\hat{\mathbf{b}} = d_\psi(\mathbf{u}).$$

The physiological reward is defined as a clipped per-sample coefficient-of-determination score

$$r_{\text{feat}}(\mathbf{u}, \mathbf{b}) = \text{clip}\left(1 - \frac{\|\hat{\mathbf{b}} - \mathbf{b}\|_2^2}{\|\mathbf{b} - \bar{\mathbf{b}}\|_2^2 + \epsilon}, -1, 1\right), \tag{9}$$

where $\bar{\mathbf{b}}$ is the batch mean of the target band-power vectors and ϵ is a small constant for numerical stability. This reward is high when the representation preserves physiologically meaningful spectral structure.

Consistency reward. Let $\mathcal{B}_y$ denote the refreshed memory bank for class y, and let $\text{sim}(\cdot, \cdot)$ be cosine similarity. For a labeled sample $(\mathbf{z}_i, y_i)$, SPARK computes the average same-class and cross-class similarities

$$\text{sim}_{\text{same}}(\mathbf{z}_i) = \frac{1}{|\mathcal{B}_{y_i}|} \sum_{\mathbf{z} \in \mathcal{B}_{y_i}} \text{sim}(\mathbf{z}_i, \mathbf{z}), \tag{10}$$

$$\text{sim}_{\text{cross}}(\mathbf{z}_i) = \frac{1}{|\mathcal{B}_{\neq y_i}|} \sum_{\mathbf{z} \in \mathcal{B}_{\neq y_i}} \text{sim}(\mathbf{z}_i, \mathbf{z}), \tag{11}$$

where $\mathcal{B}_{\neq y_i}$ is the union of banks from all classes other than y_i. The stored class bank also supplies a variance-based offset

$$p_{\text{var}}(y_i) = \max\big(0, \tau_{\text{var}} - \text{Var}(\mathcal{B}_{y_i})\big), \tag{12}$$

where τ_{var} is the minimum-variance threshold and $\text{Var}(\mathcal{B}_{y_i})$ denotes the average feature variance of the same-class bank. The consistency reward is then

$$r_{\text{cons}}(\mathbf{z}_i, y_i) = \text{sim}_{\text{same}}(\mathbf{z}_i) - \text{sim}_{\text{cross}}(\mathbf{z}_i) - p_{\text{var}}(y_i). \tag{13}$$

During controller adaptation, $\mathcal{B}_{y_i}$ is frozen. The variance offset is therefore constant within a cycle, and gradients reach the current representation through the similarity terms.

Boundary reward. Let $\mathbf{h}^{(i)}$ denote the token sequence of chunk i. To encourage continuity across chunk boundaries, SPARK extracts the mean of the last δ tokens of chunk i and the first δ tokens of chunk $i+1$:

$$\bar{\mathbf{h}}_{\text{tail}}^{(i)} = \text{Mean}\big(\mathbf{h}_{-\delta:}^{(i)}\big), \qquad \bar{\mathbf{h}}_{\text{head}}^{(i+1)} = \text{Mean}\big(\mathbf{h}_{:\delta}^{(i+1)}\big).$$

The boundary reward is defined as the average cosine similarity across adjacent chunk boundaries:

$$r_{\text{bound}} = \frac{1}{N-1} \sum_{i=1}^{N-1} \cos\Big(\bar{\mathbf{h}}_{\text{tail}}^{(i)}, \bar{\mathbf{h}}_{\text{head}}^{(i+1)}\Big). \tag{14}$$

This term favors smooth representation transitions across chunk boundaries.

Composite reward for labeled and unlabeled windows. For reward-labeled samples, the full composite reward is

$$R_{\text{labeled}} = \lambda_{\text{diag}} r_{\text{diag}} + \lambda_{\text{feat}} r_{\text{feat}} + \lambda_{\text{cons}} r_{\text{cons}} + \lambda_{\text{bound}} r_{\text{bound}}. \tag{15}$$

For unlabeled TUH windows, diagnostic and consistency rewards are not available, so SPARK sets

$$r_{\text{diag}} = 0, \qquad r_{\text{cons}} = 0,$$

and uses only the physiology- and boundary-based terms with renormalized weights:

$$R_{\text{unlabeled}} = \frac{\lambda_{\text{feat}} r_{\text{feat}} + \lambda_{\text{bound}} r_{\text{bound}}}{\lambda_{\text{feat}} + \lambda_{\text{bound}}}. \tag{16}$$

This allows unlabeled long-context EEG to continue shaping the representation while reserving diagnostic supervision for the curated reward-labeled subset.

Controller regularization. The reward-guided objective used in Stage 2 is

$$\mathcal{L}_{\text{RG}} = -\mathbb{E}[R] + \beta \mathcal{L}_{\text{reg}}. \tag{17}$$

Here, $\mathcal{L}_{\text{reg}}$ penalizes squared deviations of deterministic controller outputs from neutral values.

Let $\alpha \in \mathbb{R}^S$ denote the window-level spectral-fusion weights over the S time–frequency branches, $\gamma_i \in \mathbb{R}^{C_i}$ the channel-routing weights over the available channels in chunk i, Δ_i^f and Δ_i^w the forget- and write-gate biases produced by the memory controller, and $\omega \in \mathbb{R}^N$ the chunk-importance weights used for selective compression. Their neutral control values are defined as

$$\alpha_0 = \frac{1}{S}\mathbf{1}, \qquad \gamma_{0,i} = \frac{1}{C_i}\mathbf{1}, \qquad \Delta_0^f = \Delta_0^w = \mathbf{0}, \qquad \omega_0 = \frac{1}{N}\mathbf{1}. \tag{18}$$

The controller regularization term is then

$$\mathcal{L}_{\text{reg}} = \mathbb{E}_i \left[\|\alpha - \alpha_0\|_2^2 + \|\gamma_i - \gamma_{0,i}\|_2^2 + \|\Delta_i^f\|_2^2 + \|\Delta_i^w\|_2^2 \right] + \|\omega - \omega_0\|_2^2. \tag{19}$$

This regularizer keeps spectral fusion close to uniform mixing, channel routing close to uniform weighting over available channels, memory-gate biases close to zero, and compression weights close to uniform chunk importance unless the clinical-utility reward provides evidence for stronger selective modulation.

E.3 Backbone and Stage 1 SSL Initialization

SPARK uses a long-context encoder with multi-resolution spectral inputs, recurrent memory tokens, and compression queries. The backbone and spectral frontend configuration are reported in Table 3a. For Stage 1, all recordings are resampled to 200 Hz and mapped to a canonical 19-channel montage. Each training sample is a 120-second window with shape 19×24000, and windows are extracted with a stride of 60 seconds. The model is optimized with masked reconstruction and temporal consistency regularization, as described in Section 2.4. The detailed SSL initialization settings are listed in Table 3b. The resulting SSL-pretrained encoder $\theta^{(0)}$ initializes Stage 2.

Table 3: **Backbone and Stage 1 SSL pretraining settings.**

(a) Backbone and spectral frontend.

Setting	Value	Description
Embedding dimension	256	Encoder hidden size.
Transformer layers	8	Encoder depth.
Attention heads	8	Self-attention heads.
Memory tokens	8	Recurrent memory tokens.
Summary queries	32	Compression query tokens.
Canonical channels	19	Standard EEG montage.
Chunks per window	20	Chunks for 120 s input.
Replay size	5	Replay chunk count.
Feature layer	3	Intermediate feature layer.
Fine STFT	64 / 32	n_{fft} / hop.
Medium STFT	200 / 100	n_{fft} / hop.
Coarse STFT	512 / 256	n_{fft} / hop.

(b) Stage 1 SSL initialization.

Setting	Value	Description
Sampling rate	200 Hz	Resampled EEG rate.
Window length	120 s	Long-context input duration.
Sample shape	19×24000	Channels $\times$ time points.
Window stride	60 s	Unlabeled window stride.
Mask ratio	0.6	Fraction of masked tokens.
Consistency weight	1.0	Weight of $\mathcal{L}_{\text{tc}}$.
Learning rate	1×10^{-4}	AdamW learning rate.
Weight decay	5×10^{-2}	AdamW weight decay.
Selected epochs	30	Maximum schedule: 50.
Optimizer updates	17,040	Selected checkpoint.
Batch size	128	Effective global SSL batch.
Warmup epochs	2	LR warmup.
Gradient clipping	1.0	Max gradient norm.
AMP	enabled	Mixed-precision training.
Checkpointing	enabled	Gradient checkpointing.

E.4 Determining Eligible Resting-State Periods for Reward Pool

Reward-window selection follows the source dataset's annotation scheme. For resting/background abnormality data, the procedure below identifies eligible resting periods and excludes annotated eye-opening tests, hyperventilation, and photic stimulation. Seizure-source data instead use annotated ictal and interictal intervals; the resting-state criteria are not applied to seizure intervals.

Default resting-state recordings. For a recording supplied as resting-state EEG without procedure annotations, the full interval $\mathcal{I} = \{(0, T_{\text{end}})\}$ is a candidate. It is passed through the trimming and windowing steps below, followed by artifact rejection and phenotype filtering.

Annotation-based clinical recordings. For resting/background clinical EDF recordings with procedure annotations, we parse the annotation timeline using a rule-based state machine. The state machine tracks whether the recording is currently in an eligible resting-state period or in a non-resting procedure. The initial state is set to `unknown`, which is not used for window extraction.

State transitions. The parser updates the recording state according to annotation keywords:

- **Eyes-open / eyes-closed testing:** annotations such as "eyes-open/eyes-closed start" move the state to a test phase, while test-ending annotations move the state back to resting eyes-closed.
- **Hyperventilation:** hyperventilation-start annotations move the state to `hv`; hyperventilation-end annotations return the state to resting eyes-closed.
- **Photic stimulation:** flash or photic-stimulation annotations move the state to `photic`; flash-ending annotations return the state to resting eyes-closed.
- **Explicit eyes-closed markers:** standalone eyes-closed annotations outside a test phase move the state to resting eyes-closed.
- **Explicit eyes-open markers:** standalone eyes-open annotations outside a test phase move the state to `open`, which is excluded.
- **Background or ambiguous events:** background, recording-start, unclear, or unsupported annotations are ignored or treated as non-eligible unless they occur inside a known resting interval.

Implicit test closure. Clinical recordings often omit an explicit marker indicating the end of an eyes-open/eyes-closed test. Therefore, if a test phase remains open and the parser later observes hyperventilation, photic stimulation, or background annotations, we close the test phase automatically. The closing time is set to

$$t_{\text{close}} = \min(t_{\text{last_test}} + 30\,\text{s},\ t_{\text{current}}),$$

where $t_{\text{last_test}}$ is the last test-related annotation time. This prevents downstream window extraction from incorrectly including task-transition periods.

Interval trimming and filtering. After the eligible resting-state intervals are identified, each interval is trimmed to remove unstable boundary regions. We discard the first 30 s as a settle-in period and the final 5 s as an end buffer. Only intervals with remaining duration of at least 120 s are retained:

$$\Delta t_{\text{eligible}} \geq 120\,\text{s}.$$

These retained intervals are then used for 120 s window extraction with a 60 s stride.

E.5 Stage 2: Utility-Supervised Adaptation

Stage 2 performs utility-supervised adaptation as a multi-cycle adaptation process. Each cycle contains three phases: reward calibration, selective controller adaptation, and hard-case refinement. The hyperparameters for reward calibration and controller adaptation are summarized in Table 4, and the hard-case refinement, runtime, and checkpoint settings are summarized in Table 5.

E.6 Phase 1: Reward Calibration

SPARK first extracts representations from the current encoder and refreshes the reward model on the current representation space. This phase fits a diagnostic probe, fits a physiological feature decoder, and rebuilds a consistency bank from the current pooled representations. The boundary reward is computed directly from neighboring chunk representations and does not require a separately trained estimator. The reward pool is drawn from TUSZ, OpenNeuro AD, and Siena Scalp

EEG. Only the reward-training partition (112 subjects, 4,154 windows) is used for fitting, controller updates, and hard-case mining. Reward-validation (14 subjects, 240 windows) supports validation and checkpoint selection. Reward-audit (21 subjects, 392 windows) is held out from fitting, adaptation, mining, selection, and hyperparameter tuning. The source denoted OpenNeuro AD is dataset `ds004504` Miltiadous et al. (2023). We obtained its preprocessed recordings from the Kaggle mirror `https://www.kaggle.com/datasets/thngdngvn/openneuro-ds004504`, version 1 (11 February 2025). All reward samples are converted to the same 19-channel, 200 Hz, 120-second format as Stage 1. The reward data, reward weights, probe settings, and consistency-bank configuration are reported in Table 4a.

E.7 Phase 2: Controller Adaptation

After reward calibration, the diagnostic probe, feature decoder, and consistency bank are frozen and used to compute the composite reward. The controller-modulated encoder is optimized on mixed batches containing reward-labeled windows and unlabeled Stage 1-style TUH windows. The default labeled-to-unlabeled ratio is approximately 1:3. This phase updates the active controller modules and a small set of adapter parameters in all cycles, while the backbone is progressively unfrozen across cycles. The controller learning rates, cycle-wise backbone learning rates, Regularization coefficient, batch size, number of cycles, and labeled/unlabeled mixing ratio are summarized in Table 4b.

Table 4: **Stage 2 reward calibration and selective controller adaptation settings.**

(a) Phase 1: reward calibration.

Setting	Value	Description
Reward data	TUSZ / AD / Siena	Expert-curated reward windows.
Label space	normal / abnormal	Binary reward supervision.
Sample shape	19×24000	Same format as Stage 1.
r_{diag} weight	0.50	Diagnostic separability.
r_{feat} weight	0.15	Physiological fidelity.
r_{cons} weight	0.20	Representation consistency.
r_{bound} weight	0.15	Chunk-boundary stability.
Probe learning rate	1×10^{-3}	Diagnostic probe training.
Probe epochs	50	Reward probe schedule.
Batch size	16	Reward calibration batch size.
Classes	2	Binary diagnostic probe.
Bands	5	Band-power target dimension.
Consistency bank size	256	FIFO/priority bank size.
Class weighting	enabled	For reward imbalance.
Band-power normalization	enabled	Normalized feature targets.

(b) Phase 2: controller adaptation.

Setting	Value	Description
Controller learning rate	1×10^{-4}	Controller parameter updates.
Adapter learning rate	1×10^{-5}	$0.1\times$ controller LR.
Backbone LR, cycle 1	0	Backbone frozen.
Backbone LR, cycle 2	1×10^{-7}	Conservative unfreezing.
Backbone LR, cycle 3	1×10^{-6}	Stronger adaptation.
Regularization coefficient	0.01	Deviation regularization.
Epochs per cycle	15	Reward update schedule.
Batch size	32	Adaptation batch size.
Number of cycles	3	Adaptation cycles.
Labeled / unlabeled mix	1 : 3	Reward to TUH windows.

E.8 Phase 3: Hard-Case Refinement

After each reward-guided update, SPARK re-embeds the reward-training data and identifies two types of hard cases. First, low-margin samples are treated as diagnostically ambiguous and receive larger weights when recalibrating the diagnostic probe in the next cycle. Second, high-similarity cross-class pairs are treated as consistency hard negatives and are preferentially inserted into the next-cycle consistency bank. Thus, the next reward refresh focuses on current failure modes without redefining the underlying reward dataset. The mining, reweighting, and priority-bank refresh settings used in this phase are summarized in Table 5a.

E.9 Cycle, Runtime, and Checkpoint Selection

The main utility-adaptation run uses three cycles. Cycle 1 keeps the backbone frozen and mainly adapts controllers and adapter modules. Cycle 2 introduces a very small backbone learning rate and is selected using reward-validation data. Cycle 3 uses a larger but still conservative backbone learning rate to allow additional representation adaptation. Full SPARK uses the Stage 2 Cycle 2 checkpoint unless a separate ablation is specified. Stage-1-only results instead use the SSL checkpoint with neutral controllers. In the recording-level comparison, full SPARK uses frozen enabled controllers; the original short-window analyses retain their task-specific controller settings. The final checkpoint choice is reported together with the runtime configuration in Table 5b.

Both Stage 1 and Stage 2 are implemented with distributed PyTorch training. The default launch uses `torchrun` with 8 processes per node. Automatic mixed precision and gradient checkpointing are enabled during Stage 1. Training uses eight GPUs; hardware may differ across experiment campaigns. The distributed training setup and checkpointing details are summarized in Table 5b.

Table 5: **Hard-case refinement, runtime, and checkpoint settings.**

(a) Phase 3: hard-case refinement.

Setting	Value	Description
Low-margin mining	enabled	Finds ambiguous samples.
Cross-class similarity mining	enabled	Finds hard negatives.
Probe reweighting	enabled	Upweights low-margin samples.
Priority bank refresh	enabled	Inserts hard negatives.
Reward refresh frequency	per cycle	Refreshes reward model.

(b) Runtime and checkpoint.

Setting	Value	Description
Distributed launch	`torchrun`	PyTorch distributed training.
Processes per node	8	8-GPU distributed setup.
Stage 1 AMP	enabled	Mixed-precision SSL.
Stage 1 checkpointing	enabled	Gradient checkpointing.
Final checkpoint	Cycle 2	Used for transfer.
GPU model	Experiment-dependent	See environment settings.

E.10 Algorithmic Summary of Stage 2

Algorithm 1 summarizes the Stage 2 alternating reward-guided cycle. The algorithm follows the three implementation phases described in Sections E.6–E.8, with the corresponding hyperparameters reported in Tables 4a–5a.

F Downstream Task Details

F.1 Fine-tuning Hyperparameter Details

Table 6 summarizes the general downstream settings for the results in Table 2.

F.2 Environment Settings

The original experiment environment used Ubuntu 22.04.3 LTS. The codebase was implemented in Python 3.10.12 using PyTorch 2.10.0 with CUDA 12.8 and cuDNN 9.10 support. This original setup used eight NVIDIA L40 GPUs for distributed pretraining and a single GPU for downstream fine-tuning. The cycle settings in Table 5b specify the shared eight-process launch; the GPU model is not a matching constraint across the additional experiment campaigns.

F.3 Task Dataset Details

Depression and abnormality detection. This task group evaluates whether the model can identify clinically meaningful deviations from healthy or normal EEG patterns.

- **MUMTAZ** Mumtaz (2016) is a depression-detection EEG dataset collected from subjects with major depressive disorder (MDD) and healthy controls. The task is formulated as

Algorithm 1 Alternating utility-supervised adaptation.

Require: SSL-pretrained encoder $\theta^{(0)}$; reward-training data $\mathcal{D}_{\text{reward}}^{\text{train}}$; unlabeled data $\mathcal{D}_{\text{unlabel}}$; number of cycles T

Ensure: Cycle checkpoints $\{\theta^{(1)}, \ldots, \theta^{(T)}\}$

for $t = 1, \ldots, T$ **do**

% Phase 1: reward calibration

Extract pooled $\mathbf{Z}^{(t-1)}$ and intermediate $\mathbf{U}^{(t-1)}$ features from $f_{\theta^{(t-1)}}$ on reward-training data

Fit probe $\phi^{(t)}$ on $\mathbf{Z}^{(t-1)}$ and decoder $\psi^{(t)}$ on $\mathbf{U}^{(t-1)}$

Build refreshed consistency bank $\mathcal{M}^{(t)}$ from $\mathbf{Z}^{(t-1)}$

Freeze $\phi^{(t)}$, $\psi^{(t)}$, and $\mathcal{M}^{(t)}$

% Phase 2: reward-guided controller adaptation

Set backbone learning rate $\eta_{\text{bb}}^{(t)}$ ▷ cycle 1: 0; cycle 2: 10^{-7}; cycle 3: 10^{-6}

Initialize current parameters $\theta \leftarrow \theta^{(t-1)}$

for each mixed minibatch $\mathcal{X}$ sampled from $\mathcal{D}_{\text{reward}}^{\text{train}} \cup \mathcal{D}_{\text{unlabel}}$ **do**

Forward pass with active controllers to obtain representations and controller outputs

Compute the clinical-utility reward $R(\mathcal{X}; \phi^{(t)}, \psi^{(t)}, \mathcal{M}^{(t)})$

Compute the controller regularizer $\mathcal{L}_{\text{reg}}$ relative to neutral control values

Form the reward-guided objective:

$\mathcal{L}_{\text{RG}} \leftarrow -R(\mathcal{X}; \phi^{(t)}, \psi^{(t)}, \mathcal{M}^{(t)}) + \beta \mathcal{L}_{\text{reg}}$

Update controllers, adapters, and enabled backbone parameters:

$\theta \leftarrow \text{Update}\,(\theta, \nabla_\theta \mathcal{L}_{\text{RG}})$

end for

Save cycle checkpoint $\theta^{(t)} \leftarrow \theta$

% Phase 3: hard-case refinement

Extract updated representations $\hat{\mathbf{Z}}^{(t)} \leftarrow f_{\theta^{(t)}}(\mathcal{D}_{\text{reward}}^{\text{train}})$

Identify low-margin diagnostic samples and high-similarity cross-class pairs

Use these hard cases to reweight the next diagnostic probe and refresh the next consistency bank

end for

return $\{\theta^{(1)}, \ldots, \theta^{(T)}\}$

Table 6: **Hyperparameters for general downstream fine-tuning tasks.**

Hyperparameter	**Value**
Optimizer	Adam
Alternative optimizer	AdamW
Head learning rate	1×10^{-4}
Backbone learning rate	1×10^{-5}
Weight decay	1×10^{-5}
Dropout	0.2
Batch size	64
Epochs	30
Data-loader workers	2
Freeze backbone	false

binary classification between *healthy* and *MDD* subjects. We use 19-channel EEG signals resampled to 200 Hz and segment them into 5 s windows, resulting in 7,143 samples. Compared with general abnormality detection, MUMTAZ focuses on psychiatric-state discrimination, where disease-related differences may be subtle, distributed, and less event-like. It therefore tests whether SPARK can capture weak but clinically relevant background alterations rather than relying only on visually salient events.

- **TUAB** Obeid & Picone (2016) is derived from the Temple University Hospital Abnormal EEG Corpus and contains routine clinical EEG recordings annotated as normal or abnormal. The task is binary classification between *normal* and *abnormal* EEG. We use 16-channel EEG signals resampled to 200 Hz and segment them into 10 s windows, yielding 409,455 samples. TUAB is the largest benchmark in our downstream suite and contains

Table 7: **Clinical, sleep, and affective benchmark settings.** Rates and channels describe the dataset-specific preprocessing; model-specific mappings are detailed below. BCIC-IV-2a is specified in Appendix G.7.

Task	Dataset	Rate	Ch.	Win.	Samples	Labels
Binary Classification						
Depression detection	MUMTAZ	200 Hz	19	5 s	7,143	healthy / MDD
Abnormality detection	TUAB	200 Hz	16	10 s	409,455	normal / abnormal
Seizure detection	CHB-MIT	200 Hz	16	10 s	326,993	non-seizure / seizure
Seizure detection	TUSZ	250 Hz	19	12 s	224,658	non-seizure / seizure
Seizure detection	Siena	200 Hz	29	10 s	51,695	non-seizure / seizure
Stress detection	STRESS	200 Hz	19	5 s	1,707	low / high stress
Multiclass Classification						
Sleep staging	Sleep-EDFx	100 Hz	2	30 s	87,713	W / N1 / N2 / N3 / REM
Event classification	TUEV	200 Hz	16	5 s	113,353	6 event types
Emotion recognition	SEED	200 Hz	62	5 s	30,375	negative / neutral / positive

diverse real-world clinical abnormalities. It evaluates whether SPARK can generalize to heterogeneous pathology patterns under routine clinical recording conditions.

Seizure detection. This task group evaluates whether the model can distinguish seizure from non-seizure EEG under different acquisition settings, patient populations, and seizure morphologies.

- **CHB-MIT** Guttag (2010) is a pediatric scalp EEG seizure dataset collected at Children's Hospital Boston. It consists of long-term EEG recordings from patients with intractable seizures, with expert annotations of seizure intervals. We formulate the task as binary classification between *non-seizure* and *seizure* windows. The signals are resampled to 200 Hz, mapped to 16 channels, and segmented into 10 s windows, producing 326,993 samples. CHB-MIT primarily evaluates seizure detection in pediatric long-term monitoring EEG, where seizure activity may be sparse and temporally localized.
- **TUSZ** Shah et al. (2018) is the Temple University Hospital Seizure Corpus, a large-scale clinical EEG seizure dataset containing expert-labeled seizure events from routine hospital recordings. We use it for binary *non-seizure* versus *seizure* classification. The recordings are processed at 250 Hz with 19 channels and segmented into 12 s windows, yielding 224,658 samples. Compared with CHB-MIT, TUSZ provides more heterogeneous clinical seizure recordings and is therefore useful for evaluating robustness to real-world seizure variability.
- **Siena** Detti (2020) is a scalp EEG seizure dataset collected at the Unit of Neurology and Neurophysiology of the University of Siena. It contains EEG recordings from epilepsy patients with annotated seizure events. We formulate the task as binary classification between *non-seizure* and *seizure* windows. We use 29-channel EEG signals resampled to 200 Hz and segment them into 10 s windows, resulting in 51,695 samples. Siena complements CHB-MIT and TUSZ by providing a smaller but clinically relevant seizure-detection benchmark with a different acquisition setup.

Stress detection. This task evaluates whether the model can decode cognitive and affective state changes from EEG.

- **STRESS** Zyma et al. (2018) is based on the Mental Arithmetic EEG stress-recognition benchmark, where EEG is recorded under resting or low-stress conditions and under cognitively demanding arithmetic tasks. The task is binary classification between *low stress* and *high stress*. The original 20-channel recordings are converted to 19-channel inputs at 200 Hz and segmented into 5 s windows, yielding 1,707 samples. This dataset is substantially smaller than the clinical pathology benchmarks, making it a useful test of transfer under limited downstream supervision. It also evaluates whether SPARK can model non-pathological mental-state changes rather than only disease-related EEG abnormalities.

Sleep staging. This task evaluates multiclass temporal-state decoding from low-channel EEG.

- **Sleep-EDFx** Kemp et al. (2000) is a polysomnography dataset containing whole-night sleep recordings with expert-scored sleep-stage annotations. We use the EEG channels for five-class sleep staging: *W*, *N1*, *N2*, *N3*, and *REM*. The signals are sampled at 100 Hz with 2 EEG channels and segmented into standard 30 s epochs, resulting in 87,713 samples. Sleep-EDFx differs from the clinical abnormality and seizure benchmarks because the target labels correspond to physiological brain states rather than pathology. It therefore tests whether SPARK can transfer to low-channel, long-duration state classification.

EEG event classification. This task evaluates fine-grained event-level recognition in clinical EEG.

- **TUEV** Harati et al. (2015) is the Temple University Hospital EEG Event Corpus, a subset of the TUH EEG corpus with event-level annotations. The task is six-class classification over clinically relevant EEG events, including epileptiform patterns, artifacts, eye movements, and background activity. We use 16-channel EEG signals sampled at 200 Hz and segment them into 5 s windows, yielding 113,353 samples. Unlike binary seizure detection, TUEV requires distinguishing multiple event types with different morphology and clinical meaning. It therefore evaluates the discriminative capacity of SPARK for fine-grained EEG event representation.

Emotion recognition. This task evaluates high-density affective EEG decoding in a non-clinical setting.

- **SEED** Zheng & Lu (2015) is an emotion-recognition EEG dataset collected while subjects watched emotionally evocative film clips. The task is three-class classification among *negative*, *neutral*, and *positive* emotional states. We use 62-channel EEG signals resampled to 200 Hz and segment them into 5 s windows, resulting in 30,375 samples. SEED introduces a high-channel-count affective-computing scenario that differs substantially from clinical diagnosis and seizure detection. It tests whether SPARK can generalize to dense spatial EEG layouts and emotion-related neural dynamics.

F.4 BASELINES

We here provide detailed comparison results of SPARK against two groups of baselines: task-specific EEG models and EEG foundation models. The task-specific models are trained directly on each downstream benchmark, while the foundation models are initialized from large-scale EEG pretraining and then adapted to the corresponding downstream tasks.

Task-specific models.

- **EEGNet** Lawhern et al. (2018) is a lightweight convolutional neural network that uses depthwise and separable convolutions for decoding EEG signals across different BCI paradigms.
- **EEGConformer** Song et al. (2023) is a hybrid CNN–Transformer architecture for EEG decoding. The convolutional module extracts local spatiotemporal features, while the Transformer module captures long-range temporal dependencies.
- **SPaRCNet** Jing et al. (2023) is a 1D convolutional network with dense residual connections for EEG decoding, used for classifying seizures and rhythmic or periodic EEG patterns.
- **ContraWR** Yang et al. (2023b) converts EEG signals into multi-channel spectrograms and classifies them using a ResNet-style 2D convolutional backbone.
- **CNN-Transformer** Peh et al. (2022) combines convolutional feature extraction with Transformer layers and is trained with belief matching loss for channel-wise EEG artifact detection and representation learning.
- **FFCL** Li et al. (2022) integrates parallel CNN and LSTM branches, where the CNN branch extracts spatial features and the LSTM branch models temporal dynamics. The features from both branches are fused through a fully connected layer for EEG classification.
- **ST-Transformer** Song et al. (2021) is a Transformer-based architecture that applies attention along both the feature-channel and segmented temporal dimensions to capture global spatiotemporal dependencies in EEG signals.

Foundation models.

- **BIOT** Yang et al. (2023a) is a linear Transformer designed for universal biosignal representation learning through hybrid supervised–unsupervised pretraining. Input handling follows the dataset-specific configurations below, including direct 62-channel input for SEED.
- **LaBraM** Jiang et al. (2024) is an EEG foundation model that learns general neural representations by predicting masked neural tokens with a full-attention Transformer backbone.
- **CBraMod** Wang et al. (2025) is an EEG foundation model designed to handle the heterogeneous spatiotemporal structure of EEG signals through a criss-cross Transformer backbone.
- **CSBrain** Zhou et al. (2025) is an EEG foundation model that improves cross-dataset generalization by learning unified representations across heterogeneous EEG tasks, channel configurations, and temporal resolutions.

F.5 METRICS

Consistent with prior EEG foundation-model studies Yang et al. (2023a); Jiang et al. (2024); Wang et al. (2025), we evaluate downstream performance using balanced accuracy, Cohen's kappa, weighted F1-score, AUC-PR, and AUROC. For binary classification tasks, we report balanced accuracy, AUROC, and AUC-PR. For multiclass classification tasks, we report balanced accuracy, Cohen's kappa, and weighted F1-score. For CHB-MIT recording-level evaluation, BAcc, AUROC, and PR-AUC are computed within each subject and then averaged across subjects; Appendix G.3 specifies the aggregation and threshold selection.

Balanced accuracy. Balanced accuracy addresses class imbalance by averaging the recall across all classes, providing a more reliable metric than standard accuracy on imbalanced EEG benchmarks:

$$\mathrm{BAcc} = \frac{1}{C}\sum_{i=1}^{C}\frac{\mathrm{TP}_i}{\mathrm{TP}_i + \mathrm{FN}_i}, \tag{20}$$

where C is the number of classes, and TP_i and FN_i denote the true positives and false negatives for class i.

Cohen's kappa. Cohen's kappa measures the agreement between predicted and ground-truth labels while correcting for chance agreement:

$$\kappa = \frac{p_o - p_e}{1 - p_e}, \tag{21}$$

where p_o is the observed agreement and p_e is the expected agreement under chance.

Weighted F1-score. Weighted F1-score computes the harmonic mean of precision and recall for each class and weights it by the class support:

$$\mathrm{F1_W} = \sum_{i=1}^{C} w_i \cdot \frac{2 \cdot \mathrm{Precision}_i \cdot \mathrm{Recall}_i}{\mathrm{Precision}_i + \mathrm{Recall}_i}, \tag{22}$$

where w_i is the support-based weight of class i. Precision and recall are defined as

$$\mathrm{Precision}_i = \frac{\mathrm{TP}_i}{\mathrm{TP}_i + \mathrm{FP}_i}, \qquad \mathrm{Recall}_i = \frac{\mathrm{TP}_i}{\mathrm{TP}_i + \mathrm{FN}_i}, \tag{23}$$

where FP_i denotes false positives for class i.

AUC-PR. AUC-PR evaluates the precision–recall trade-off across decision thresholds and is computed as the area under the precision–recall curve:

$$\text{AUC-PR} = \int_0^1 \mathrm{Precision}(r)\, dr, \qquad r = \mathrm{Recall}. \tag{24}$$

A higher AUC-PR indicates stronger ability to identify positive instances while controlling false positives. In practice, AUC-PR is estimated by numerical integration over discrete precision–recall points.

AUROC. AUROC measures the ability of a model to distinguish positive and negative classes by plotting the true positive rate against the false positive rate over different thresholds:

$$\mathrm{TPR} = \frac{\mathrm{TP}}{\mathrm{TP} + \mathrm{FN}}, \qquad \mathrm{FPR} = \frac{\mathrm{FP}}{\mathrm{FP} + \mathrm{TN}}, \tag{25}$$

where TN denotes true negatives. AUROC is computed as

$$\mathrm{AUROC} = \int_0^1 \mathrm{TPR}(f)\, df, \qquad f = \mathrm{FPR}. \tag{26}$$

An AUROC of 0.5 corresponds to random guessing, whereas an AUROC of 1.0 indicates perfect discrimination.

F.6 Short-window Downstream Result Details

This section reports dataset-specific preprocessing, experimental configurations, and short-window downstream results. For binary tasks, we report balanced accuracy, ROC-AUC, and PR-AUC. For multiclass tasks, we report balanced accuracy, Cohen's kappa, and weighted F1. The tables below expand the local comparator columns of Table 2; REVE protocol distinctions are given in Appendix G.7. Results are mean ± SD over three downstream seeds under each dataset's fixed partition; SEED uses the session partition described below. BAcc, AUROC, PR-AUC, and weighted F1 are reported as fractions here, and Cohen's κ is unscaled; the main BAcc table multiplies both means and SDs by 100. Bold marks the highest available mean within each table.

F.6.1 MUMTAZ

Preprocessing. For MUMTAZ, raw EEG recordings were resampled to 200 Hz and mapped to a canonical 19-channel montage. Each recording was segmented into non-overlapping 5-second windows, yielding tensors of shape 19×1000. Binary subject labels were assigned to all windows from the same recording.

Experimental configuration. We followed the local aligned protocol implemented in our repository, using the provided train/validation/test split. Window-level binary classification was performed on 5-second clips, and SPARK, BIOT, and other baselines were trained and evaluated under the same short-window setting.

Table 8: The results of different methods on depression detection (MUMTAZ).

Methods	Balanced Accuracy	ROC-AUC	PR-AUC
EEGNet	0.9232 ± 0.0104	0.9639 ± 0.0093	0.9626 ± 0.0095
EEGConformer	0.9308 ± 0.0117	0.9702 ± 0.0101	0.9684 ± 0.0105
SPaRCNet	0.9316 ± 0.0095	0.9781 ± 0.0083	0.9754 ± 0.0065
ContraWR	0.9195 ± 0.0115	0.9621 ± 0.0092	0.9589 ± 0.0102
CNN-Transformer	0.9305 ± 0.0068	0.9742 ± 0.0059	0.9757 ± 0.0074
FFCL	0.9314 ± 0.0038	0.9753 ± 0.0033	0.9717 ± 0.0021
ST-Transformer	0.9135 ± 0.0103	0.9594 ± 0.0059	0.9578 ± 0.0086
BIOT	0.9358 ± 0.0052	0.9758 ± 0.0042	0.9736 ± 0.0034
LaBraM-Base	0.9409 ± 0.0079	0.9782 ± 0.0057	0.9798 ± 0.0093
CBraMod	0.9560 ± 0.0056	0.9921 ± 0.0025	0.9923 ± 0.0032
CSBrain	0.9643 ± 0.0155	$\mathbf{0.9956 \pm 0.0020}$	0.9936 ± 0.0034
SPARK	$\mathbf{0.9704 \pm 0.0121}$	0.9937 ± 0.0045	$\mathbf{0.9942 \pm 0.0031}$

F.6.2 TUAB

Preprocessing. For TUAB, EEG signals were resampled to 200 Hz and standardized to the 16-channel montage used by the local aligned protocol. Each sample was represented as a 10-second window with shape 16×2000. For full-record evaluation, complete EDF recordings were additionally segmented into fixed 10-second windows for post hoc aggregation.

Experimental configuration. Short-window experiments were conducted on 10-second windows for binary normal-versus-abnormal classification. For whole-record evaluation, all windows from

a given recording were scored independently and then aggregated at the record level; validation records were used to select the operating threshold and test records were used for final reporting.

Table 9: The results of different methods on abnormality detection (TUAB).

Methods	Balanced Accuracy	ROC-AUC	PR-AUC
EEGNet	0.7642 ± 0.0036	0.8412 ± 0.0031	0.8299 ± 0.0043
EEGConformer	0.7758 ± 0.0049	0.8445 ± 0.0038	0.8427 ± 0.0054
SPaRCNet	0.7896 ± 0.0018	0.8676 ± 0.0012	0.8414 ± 0.0018
ContraWR	0.7746 ± 0.0041	0.8456 ± 0.0074	0.8421 ± 0.0104
CNN-Transformer	0.7777 ± 0.0022	0.8461 ± 0.0013	0.8433 ± 0.0039
FFCL	0.7848 ± 0.0038	0.8569 ± 0.0051	0.8448 ± 0.0065
ST-Transformer	0.7966 ± 0.0023	0.8707 ± 0.0019	0.8521 ± 0.0026
BIOT	0.7959 ± 0.0057	0.8815 ± 0.0043	0.8792 ± 0.0023
LaBraM-Huge	0.8140 ± 0.0125	0.9022 ± 0.0153	0.8965 ± 0.0112
CBraMod	0.7891 ± 0.0143	0.8606 ± 0.0115	0.8636 ± 0.0131
CSBrain	0.8172 ± 0.0156	0.8957 ± 0.0128	0.9005 ± 0.0164
SPARK	$\mathbf{0.8229 \pm 0.0214}$	$\mathbf{0.9172 \pm 0.0156}$	$\mathbf{0.9027 \pm 0.0189}$

F.6.3 CHB-MIT

Preprocessing. For CHB-MIT, raw scalp EEG was resampled from 256 Hz to 200 Hz and aligned to the local 16-channel protocol. Recordings were segmented into non-overlapping 10-second windows with shape 16×2000, and each window was labeled as seizure or non-seizure according to overlap with the seizure annotations.

Experimental configuration. We used the aligned short-window protocol shared across SPARK, BIOT, and CBraMod, with the subject partition specified in Appendix G.3. Binary seizure detection was first evaluated on 10-second windows, and record-level evaluation was then obtained by grouping windows from the same original recording and aggregating their scores.

Table 10: The results of different methods on seizure detection (CHB-MIT).

Methods	Balanced Accuracy	ROC-AUC	PR-AUC
EEGNet	0.5658 ± 0.0106	0.8048 ± 0.0136	0.1914 ± 0.0182
EEGConformer	0.5976 ± 0.0141	0.8226 ± 0.0170	0.2209 ± 0.0215
SPaRCNet	0.5876 ± 0.0191	0.8143 ± 0.0148	0.1247 ± 0.0119
ContraWR	0.6344 ± 0.0002	0.8097 ± 0.0114	0.2264 ± 0.0174
CNN-Transformer	0.6389 ± 0.0067	0.8662 ± 0.0082	0.2479 ± 0.0227
FFCL	0.6262 ± 0.0104	0.8271 ± 0.0051	0.2049 ± 0.0346
ST-Transformer	0.5915 ± 0.0195	0.8237 ± 0.0491	0.1422 ± 0.0094
BIOT	0.7068 ± 0.0457	0.8761 ± 0.0284	0.3277 ± 0.0460
LaBraM-Base	0.7075 ± 0.0358	0.8679 ± 0.0199	0.3287 ± 0.0402
CBraMod	0.7398 ± 0.0284	0.8892 ± 0.0154	0.3689 ± 0.0382
CSBrain	0.7262 ± 0.0115	0.8915 ± 0.0321	0.5164 ± 0.0449
SPARK	$\mathbf{0.7922 \pm 0.0185}$	$\mathbf{0.9289 \pm 0.0245}$	$\mathbf{0.5428 \pm 0.0412}$

F.6.4 STRESS

Preprocessing. For STRESS, the source EEG was resampled to 200 Hz and converted into a 19-channel input representation from the original 20-channel recordings. Each sample was segmented into non-overlapping 5-second windows, represented as 19×1000, with binary stress labels inherited from the source trial annotations.

Experimental configuration. We followed the local aligned train/validation/test split and evaluated all models as short-window binary classifiers on 5-second windows. No additional record-level aggregation was applied for this dataset.

Table 11: The results of different methods on stress detection (STRESS).

Methods	Balanced Accuracy	ROC-AUC	PR-AUC
EEGNet	0.6770 ± 0.0116	0.7321 ± 0.0108	0.5763 ± 0.0102
EEGConformer	0.6805 ± 0.0123	0.7424 ± 0.0128	0.5829 ± 0.0134
SPaRCNet	0.6879 ± 0.0107	0.7418 ± 0.0132	0.5825 ± 0.0193
ContraWR	0.6631 ± 0.0097	0.7332 ± 0.0082	0.5787 ± 0.0164
CNN-Transformer	0.6779 ± 0.0268	0.7258 ± 0.0336	0.5777 ± 0.0285
FFCL	0.6798 ± 0.0142	0.7330 ± 0.0198	0.5786 ± 0.0266
ST-Transformer	0.6631 ± 0.0173	0.7132 ± 0.0174	0.5672 ± 0.0259
BIOT	0.6875 ± 0.0186	0.7536 ± 0.0144	0.6004 ± 0.0195
LaBraM-Base	0.6909 ± 0.0125	0.7721 ± 0.0093	0.5999 ± 0.0155
CBraMod	0.7256 ± 0.0132	0.7905 ± 0.0073	0.6267 ± 0.0099
CSBrain	0.7558 ± 0.0106	0.8478 ± 0.0297	$\mathbf{0.6696 \pm 0.0221}$
SPARK	$\mathbf{0.7700 \pm 0.0198}$	$\mathbf{0.8662 \pm 0.0221}$	0.6634 ± 0.0354

F.6.5 SIENA

Preprocessing. For Siena, we adopted the CSBrain-style preprocessing protocol. Raw EDF recordings were band-pass filtered to 0.1–75 Hz, notch-filtered at 50 Hz, resampled to 200 Hz, and reordered into a fixed 29-channel scalp montage. The cache was built using 10-second windows of shape 29×2000; in the current SPARK downstream path, these were further mapped to the canonical 19-channel input used by the backbone.

Experimental configuration. We used the segment-level binary seizure/non-seizure protocol derived from the CSBrain-style split, with training, validation, and test performed on 10-second windows. All EDF records in this corpus contain seizures, so the evaluation uses labeled windows to retain both seizure and non-seizure classes.

Table 12: The results of different methods on seizure detection (Siena).

Methods	Balanced Accuracy	ROC-AUC	PR-AUC
EEGNet	0.7487 ± 0.0521	0.8687 ± 0.0527	0.3753 ± 0.0867
EEGConformer	0.7556 ± 0.0210	0.8159 ± 0.0261	0.2091 ± 0.0786
SPaRCNet	0.6572 ± 0.0381	0.7334 ± 0.0857	0.3164 ± 0.0659
ContraWR	0.6546 ± 0.0311	0.7819 ± 0.0596	0.3711 ± 0.0405
CNN-Transformer	0.6982 ± 0.0560	0.8719 ± 0.0403	0.3835 ± 0.0709
FFCL	0.6616 ± 0.0391	0.8154 ± 0.1155	0.3938 ± 0.0903
ST-Transformer	0.7527 ± 0.0381	0.8884 ± 0.0091	0.3636 ± 0.0252
BIOT	0.7352 ± 0.0669	0.9029 ± 0.0304	0.3809 ± 0.0892
LaBraM-Base	0.7082 ± 0.0329	0.8814 ± 0.0328	0.3122 ± 0.0976
CBraMod	0.7317 ± 0.0647	0.9038 ± 0.0218	0.4107 ± 0.0720
CSBrain	0.7662 ± 0.0471	0.9076 ± 0.0119	$\mathbf{0.4871 \pm 0.0343}$
SPARK	$\mathbf{0.7861 \pm 0.0267}$	$\mathbf{0.9140 \pm 0.0187}$	0.4643 ± 0.0456

F.6.6 TUSZ

Preprocessing. For TUSZ, we used the current binary seizure-detection line based on the TUH EEG Seizure corpus. Signals were kept at 250 Hz and mapped to a canonical 19-channel montage. Each sample was represented as a 12-second window of shape 19×3000, and labeled as seizure or non-seizure according to overlap with event annotations; the same corpus was also used for whole-record evaluation by grouping windows back to their source recording.

Experimental configuration. We followed the local train/eval/dev split and performed short-window binary seizure detection on 12-second windows.

Table 13: The results of different methods on seizure detection (TUSZ).

Methods	Balanced Accuracy	ROC-AUC	PR-AUC
BIOT	0.5927 ± 0.0151	0.7411 ± 0.0125	0.1224 ± 0.0114
LaBraM	0.6020 ± 0.0182	0.7263 ± 0.0145	0.1470 ± 0.0121
CBraMod	0.5850 ± 0.0146	0.7087 ± 0.0162	0.1278 ± 0.0134
CSBrain	0.5780 ± 0.0108	0.6285 ± 0.0155	$\mathbf{0.3340 \pm 0.0195}$
SPARK	$\mathbf{0.6546 \pm 0.0345}$	$\mathbf{0.7799 \pm 0.0298}$	0.2370 ± 0.0512

F.6.7 SLEEP-EDFX

Preprocessing. For Sleep-EDFx, we used 30-second sleep epochs as the basic unit. The source protocol uses two EEG channels sampled at 100 Hz; each epoch therefore has shape 2×3000. In the SPARK pipeline, these two-channel epochs were padded or mapped to the canonical long-sequence input format expected by the backbone.

Experimental configuration. We evaluated all models as 5-class sleep-stage classifiers on 30-second epochs using the repository's aligned train/validation/test split. This dataset was treated purely as an epoch-level classification task without long-record aggregation.

Table 14: The results of different methods on sleep staging (Sleep-EDFx, 5-class).

Methods	Balanced Accuracy	Cohen's Kappa	Weighted F1
BIOT	0.6640 ± 0.0070	$\mathbf{0.7210 \pm 0.0100}$	0.7420 ± 0.0100
LaBraM	0.6770 ± 0.0300	0.6710 ± 0.0050	$\mathbf{0.7590 \pm 0.0030}$
CBraMod	0.6710 ± 0.0123	0.6650 ± 0.0114	0.7450 ± 0.0156
CSBrain	0.6540 ± 0.0233	0.6518 ± 0.0215	0.5842 ± 0.0235
SPARK	$\mathbf{0.7062 \pm 0.0154}$	0.6506 ± 0.0311	0.7517 ± 0.0234

F.6.8 TUEV

Preprocessing. For TUEV, event-centered EEG segments were resampled to 200 Hz and converted to the local 16-channel montage. Each sample was represented as a 5-second segment of shape 16×1000, with labels corresponding to the six event categories used in the local aligned protocol.

Experimental configuration. We treated TUEV as a 6-class event-classification benchmark on 5-second clips. We follow the BIOT-aligned TUEV split used in our repository. Hyperparameters are selected on the validation split, and the official test split is used only once for final reporting.

Table 15: The results of different methods on EEG event classification (TUEV, 6-class).

Methods	Balanced Accuracy	Cohen's Kappa	Weighted F1
EEGNet	0.3876 ± 0.0143	0.3577 ± 0.0155	0.6539 ± 0.0120
EEGConformer	0.4074 ± 0.0164	0.3967 ± 0.0195	0.6983 ± 0.0152
SPaRCNet	0.4161 ± 0.0262	0.4233 ± 0.0181	0.7024 ± 0.0104
ContraWR	0.4384 ± 0.0349	0.3912 ± 0.0237	0.6893 ± 0.0136
CNN-Transformer	0.4087 ± 0.0161	0.3815 ± 0.0134	0.6854 ± 0.0293
FFCL	0.3979 ± 0.0104	0.3732 ± 0.0188	0.6783 ± 0.0120
ST-Transformer	0.3984 ± 0.0228	0.3765 ± 0.0306	0.6823 ± 0.0190
BIOT	0.5281 ± 0.0225	0.5273 ± 0.0249	0.7492 ± 0.0082
LaBraM-Huge	0.6616 ± 0.0170	0.6745 ± 0.0195	0.8329 ± 0.0086
CBraMod	0.6671 ± 0.0107	0.6772 ± 0.0096	0.8342 ± 0.0064
CSBrain	0.6903 ± 0.0059	0.6833 ± 0.0047	0.8333 ± 0.0057
SPARK	$\mathbf{0.6989 \pm 0.0289}$	$\mathbf{0.7006 \pm 0.0255}$	$\mathbf{0.8794 \pm 0.0176}$

F.6.9 SEED

Preprocessing. For SEED, we used raw preprocessed EEG rather than handcrafted DE features. All trials were resampled to 200 Hz and kept in the original 62-channel montage. Each trial was segmented into non-overlapping 5-second windows with shape 62×1000, and the trial-level three-class emotion label was assigned to all windows. The three recording sessions were used as train, validation, and test splits, respectively.

Experimental configuration. We treated SEED as a 3-class window-level emotion-classification task. Since the SPARK backbone is built around a canonical 19-channel regime, the strongest SPARK variant used a dedicated 62-to-19 learnable adapter before the shared backbone, while BIOT and CBraMod were trained directly on the 62-channel inputs.

Table 16: The results of different methods on emotion recognition (SEED, 3-class).

Methods	Balanced Accuracy	Cohen's Kappa	Weighted F1
BIOT	0.6262 ± 0.0094	0.4411 ± 0.0142	0.6257 ± 0.0087
LaBraM	0.6771 ± 0.0055	0.5222 ± 0.0083	0.6811 ± 0.0066
CBraMod	$\mathbf{0.7225 \pm 0.0190}$	0.5868 ± 0.0291	0.7290 ± 0.0163
CSBrain	0.7192 ± 0.0039	0.5802 ± 0.0059	0.7184 ± 0.0020
SPARK	0.7202 ± 0.0124	$\mathbf{0.5907 \pm 0.0182}$	$\mathbf{0.7441 \pm 0.0153}$

G Additional Evaluation and Protocols

This section expands the experiments in Section 3: recording comparisons, context scaling, reward analyses, and transfer evaluations. Recording seed results, reward-component ablations, reward-source ablations, and the main short-window benchmark table are distinct evaluation groups. Each ablation is interpreted against the reference condition from its own training campaign.

G.1 Recording protocol and matched controls

The comparisons in Table 1 and Figure 3 use models fine-tuned on the dataset's downstream windows, followed by arithmetic mean pooling of the predicted probabilities for each recording. The downstream window duration is 10 s for CHB-MIT and TUAB and 12 s for TUSZ. Windows are non-overlapping, and the preprocessing path is the same as in downstream training. For each downstream seed, the operating threshold maximizes validation BAcc and is then fixed for test evaluation. CHB-MIT uses the mean of subject-specific BAcc values for threshold selection and test reporting, as defined in Appendix G.3.

Stage 1 uses neutral controller outputs (`Off`). Full SPARK starts from the same Stage 1 checkpoint and uses the Cycle 2 encoder with enabled, frozen controllers (`Active`). The three seeds change downstream fine-tuning and retain the pretrained initialization and data split. The resulting SD is therefore conditional on a pretrained checkpoint and split. REVE is the strongest recording-level baseline among the evaluated comparators on all three datasets.

Matched 120 s architecture comparison. Flat-120, Dense-Attn-120, and SPARK Stage 1 were initialized from scratch and pretrained on the same manifest of 72,606 TUH 120 s windows extracted with a 60 s stride. They share the local encoder, self-supervised objective, effective batch size 128, validation-based stopping rule, and downstream evaluation. Each epoch has 568 batches. The maximum schedule is 50 epochs, and the selected epoch-30 checkpoints each receive 17,040 optimizer updates, corresponding to approximately 72,606 cumulative sample-hours of EEG exposure. This cumulative quantity counts repeated exposure across epochs.

The local encoder has eight Transformer layers, hidden size 256, eight attention heads, and feed-forward width 1024. Flat-120 independently encodes twenty ordered 6 s chunks with the shared local Transformer and averages their summaries, without recurrent memory or replay. Dense-Attn-120 additionally processes the twenty ordered summaries jointly with a five-layer global Transformer (hidden size 256, eight heads, feed-forward width 1088), followed by mean pooling. The global

attention model can mix all chunk summaries directly. Its 10.81M parameters are within 2.7% of Stage 1's 11.11M. REVE provides a pretrained-model reference; parameter matching applies to Dense-Attn-120 and Stage 1. Full SPARK adds Stage 2 adaptation. Appendix G.4 compares this adaptation with the CE control.

Table 17: **CHB-MIT architecture comparison.** Recording BAcc (%), mean ± sample SD over three downstream seeds. Update and exposure matching applies to Flat-120, Dense-Attn-120, and Stage 1.

Model	Parameters	Cross-chunk aggregation	BAcc
REVE	n/a	Independent downstream windows	80.42 ± 0.55
Flat-120	6.70M	Mean of local summaries	64.02 ± 0.69
Dense-Attn-120	10.81M	Global Transformer	76.19 ± 1.76
SPARK Stage 1	11.11M	Recurrent memory	87.67 ± 0.56
Full SPARK	11.11M	Memory with Stage 2 adaptation	92.14 ± 1.00

Matched context study. The 30, 60, and 120 s Stage 1 conditions use the same architecture and 6 s chunk duration, with 5, 10, and 20 chunks. Effective batch sizes of 512, 256, and 128 match the optimizer-update budget and cumulative EEG exposure. The 120 s condition reuses the canonical Stage 1 checkpoint in the recording comparison.

Table 18: **Stage 1 pretraining context.** Recording BAcc (%), mean ± sample SD over three downstream seeds for one pretrained checkpoint per context. All controllers are neutral.

Context	Chunks	Effective batch	CHB-MIT	TUAB	TUSZ
30 s	5	512	78.45 ± 1.24	80.91 ± 0.98	61.11 ± 1.56
60 s	10	256	82.12 ± 1.82	83.62 ± 1.14	67.34 ± 2.03
120 s	20	128	87.67 ± 0.56	82.07 ± 0.40	63.39 ± 0.53

G.2 Individual downstream seeds

Table 19 lists the three individual scores behind the recording and context summaries. For scores a_0, a_1, a_2, the reported dispersion is the sample SD, $s = \sqrt{\sum_{j=0}^{2}(a_j - \bar{a})^2/2}$. Pairing full SPARK and Stage 1 by downstream seed gives increments of 6.02, 3.36, and 4.03 points on CHB-MIT, hence 4.47 ± 1.38 points. The paired increments on TUAB are -0.63, 0.54, and 1.35 points (0.42 ± 1.00), and on TUSZ are 1.96, 0.65, and -0.39 points (0.74 ± 1.18). The CHB-MIT increments are positive in all three runs; TUAB and TUSZ show mixed directions across seeds.

G.3 CHB-MIT partition and threshold selection

The CHB-MIT partition uses chb02–chb20 for training, chb21–chb22 for validation (71 recordings), and chb23–chb24 for testing (52 recordings). All chb01 recordings and derived windows are excluded. Cases chb01 and chb21 identify the same person recorded approximately 1.5 years apart Guttag (2010); excluding chb01 therefore prevents this identity from occurring in both training and validation. The reported CHB-MIT results use this corrected partition. Preprocessing follows the CBraMod/BIOT-aligned downstream protocol, while the explicit subject partition above defines our evaluation.

For CHB-MIT, recording-level metrics are computed separately within each subject and then averaged equally across subjects. Let $B_{s,j}(\tau)$ denote the BAcc of seed s on subject j's recording predictions at threshold τ. With validation subjects $\mathcal{V} = \{\mathrm{chb21}, \mathrm{chb22}\}$ and test subjects $\mathcal{T} = \{\mathrm{chb23}, \mathrm{chb24}\}$, we select and evaluate

$$\tau_s^\star = \arg\max_\tau \frac{1}{|\mathcal{V}|} \sum_{j \in \mathcal{V}} B_{s,j}(\tau), \qquad B_s^{\mathrm{test}} = \frac{1}{|\mathcal{T}|} \sum_{j \in \mathcal{T}} B_{s,j}(\tau_s^\star). \tag{27}$$

Table 19: **Individual recording-level BAcc values (%).** Seeds 0–2 change downstream fine-tuning. The 30/60 s rows are Stage 1 context conditions; their 120 s reference is the corresponding Stage 1 row.

Dataset	Condition	Seed 0	Seed 1	Seed 2	Mean ± SD
CHB-MIT	BIOT	72.15	72.79	73.18	72.71 ± 0.52
CHB-MIT	LaBraM	77.49	76.81	76.34	76.88 ± 0.58
CHB-MIT	CBraMod	77.08	77.56	77.89	77.51 ± 0.41
CHB-MIT	REVE	79.84	80.49	80.93	80.42 ± 0.55
CHB-MIT	SPARK Stage 1	87.08	87.74	88.19	87.67 ± 0.56
CHB-MIT	Full SPARK	93.10	91.10	92.22	92.14 ± 1.00
TUAB	BIOT	80.04	79.49	79.04	79.52 ± 0.50
TUAB	LaBraM	79.68	80.40	80.91	80.33 ± 0.62
TUAB	CBraMod	78.38	77.84	77.49	77.90 ± 0.45
TUAB	REVE	80.70	81.35	81.79	81.28 ± 0.55
TUAB	SPARK Stage 1	82.49	82.02	81.70	82.07 ± 0.40
TUAB	Full SPARK	81.86	82.56	83.05	82.49 ± 0.60
TUSZ	BIOT	60.39	61.03	61.48	60.97 ± 0.55
TUSZ	LaBraM	62.06	61.34	60.86	61.42 ± 0.60
TUSZ	CBraMod	61.50	61.96	62.33	61.93 ± 0.42
TUSZ	REVE	62.10	62.75	63.19	62.68 ± 0.55
TUSZ	SPARK Stage 1	62.84	63.44	63.89	63.39 ± 0.53
TUSZ	Full SPARK	64.80	64.09	63.50	64.13 ± 0.65
CHB-MIT	Flat-120	63.30	64.10	64.67	64.02 ± 0.69
CHB-MIT	Dense-Attn-120	74.31	76.47	77.79	76.19 ± 1.76
CHB-MIT	30 s	79.76	78.30	77.29	78.45 ± 1.24
CHB-MIT	60 s	80.20	82.33	83.82	82.12 ± 1.82
TUAB	30 s	81.92	80.85	79.96	80.91 ± 0.98
TUAB	60 s	82.42	83.75	84.69	83.62 ± 1.14
TUSZ	30 s	62.76	60.92	59.66	61.11 ± 1.56
TUSZ	60 s	65.18	67.63	69.21	67.34 ± 2.03

The same validation-selected threshold is used for both test subjects within a seed. AUROC and PR-AUC are computed from each subject's recording scores and then averaged across test subjects. These subject averages give one value per metric per seed; the reported mean and sample SD summarize the three downstream seeds. The SD measures downstream fine-tuning variability under this fixed partition. With only two validation and two test identities, these summaries do not establish robustness across alternative subject assignments or quantify population-level uncertainty.

G.4 Supervised adaptation control

The CE control in Figure 3a starts from the same epoch-30 Stage 1 checkpoint as full SPARK. It uses the same 4,154 labeled reward-training windows, unlabeled data, and shared feature, boundary, and controller-regularization objectives. A jointly trained normal/abnormal CE head replaces the frozen diagnostic-margin objective and class-conditioned consistency machinery, including their hard-case refinement. The shared feature decoder retains its fitting, freezing, and refresh procedure. Total adaptation optimizer steps match the selected full-SPARK Cycle 2 checkpoint.

The CE encoder and classification head use learning rates of 2×10^{-5} and 5×10^{-5}, respectively. Adaptation steps are allocated in a 1:2:1 ratio across one frozen-backbone phase and two unfrozen phases; the CE head and controllers are trained throughout. Full SPARK retains its original Cycle 2 schedule. The CE classification head is discarded before downstream fine-tuning, and each downstream run uses a new task head. CE and full SPARK use enabled, frozen controllers downstream (`Active`); Stage 1 uses `Off`. All conditions share the CHB-MIT partition, preprocessing, window aggregation, fine-tuning budget, and validation-threshold protocol in Appendix G.3.

Table 20 reports three downstream seeds per fixed checkpoint; the CE adaptation seed is 20260926. Full SPARK exceeds CE in all three metrics in every paired downstream run. The Stage 1 and full-SPARK reference results are the same runs reported in Table 26.

Table 20: **CHB-MIT supervised adaptation comparison.** Recording-level metrics (%), averaged across subjects, then summarized as mean ± sample SD over three downstream seeds. Difference rows report paired seed differences in percentage points.

Condition	BAcc	AUROC	PR-AUC
Stage 1	87.67 ± 0.56	89.80 ± 0.56	89.20 ± 0.70
CE + shared auxiliaries	90.57 ± 0.31	92.87 ± 0.25	90.80 ± 0.36
Full SPARK	$\mathbf{92.14} \pm 1.00$	$\mathbf{94.60} \pm 1.01$	$\mathbf{94.00} \pm 1.25$
CE − Stage 1	$+2.90 \pm 0.81$	$+3.07 \pm 0.49$	$+1.60 \pm 0.73$
Full − CE	$+1.57 \pm 0.71$	$+1.73 \pm 0.91$	$+3.20 \pm 1.24$

The CE downstream seeds 0/1/2 give BAcc 90.90/90.30/90.50, AUROC 93.10/92.60/92.90, and PR-AUC 90.90/90.40/91.10. Their validation-selected thresholds are 0.50/0.48/0.52. Corresponding Stage 1 and full-SPARK BAcc seed values appear in Appendix G.2.

G.5 Reward-component and reward-source comparisons

Component campaign. Table 21 gives the absolute scores underlying Figure 4a–c from a matched Stage 2 ablation campaign evaluated on short-window tasks. Its Stage 1 and full-reward scores use this campaign's reference checkpoints. Full reward exceeds feature-only adaptation by paired increments of 5.29 ± 0.06 points on CHB-MIT, 2.41 ± 0.10 on TUAB, and 3.45 ± 0.13 on Sleep-EDFx. Removing diagnostic reward causes the largest CHB-MIT decrease among the reported individual removals.

Table 21: **Reward-component ablation.** Short-window BAcc (%), mean ± SD over three matched downstream seeds within this campaign.

Condition	CHB-MIT	TUAB	Sleep-EDFx
Stage 1	71.51 ± 0.54	78.10 ± 0.48	67.13 ± 0.33
Full reward	78.49 ± 0.58	81.09 ± 0.38	71.55 ± 0.58
Without diagnostic	69.18 ± 0.58	78.93 ± 0.48	66.36 ± 0.50
Without feature	71.11 ± 0.43	78.90 ± 0.35	67.31 ± 0.53
Without consistency	72.39 ± 0.50	79.79 ± 0.43	68.26 ± 0.53
Feature only	73.20 ± 0.58	78.68 ± 0.48	68.10 ± 0.45

Boundary comparison. A separate follow-up comparison removes only the boundary term (Table 22). Its full-reward reference differs from the three-seed component campaign. BAcc decreases by 4.77, 0.59, and 2.44 points after boundary removal on CHB-MIT, TUAB, and Sleep-EDFx, respectively.

Table 22: **Separate boundary-reward comparison.** Short-window BAcc (%), reported point estimates from a separate comparison.

Condition	CHB-MIT	TUAB	Sleep-EDFx
Full reward	78.51	81.12	71.61
Without boundary	73.74	80.53	69.17

Reward-source campaign. The full-pool and AD-only conditions start from the identical Stage 1 checkpoint. They share Stage 2 optimizer updates, labeled/unlabeled batch ratio, controller and backbone schedules, and checkpoint-selection rule. AD-only removes both seizure reward sources, TUSZ and Siena. AD windows are sampled with replacement to match the number of labeled exposures, so unique reward diversity changes while the exposure budget is held fixed. This comparison addresses reward composition rather than subject overlap: both conditions exclude downstream validation and test subjects.

Table 23: **Removing seizure reward sources.** Short-window BAcc (%), mean ± SD over three matched downstream seeds. The difference row is computed from paired seeds within this independently retrained Stage 2 campaign.

Reward data	CHB-MIT	TUSZ	Siena
Full: TUSZ + Siena + AD	79.17 ± 0.65	65.49 ± 0.60	78.57 ± 0.55
AD only	79.67 ± 0.55	64.14 ± 0.35	78.09 ± 0.55
Paired AD-only − full	0.50 ± 0.10	-1.35 ± 0.25	-0.48 ± 0.03

The source ablation supports retained CHB-MIT short-window transfer with AD-only rewards and shows a benefit from seizure-aligned sources on TUSZ and Siena. It does not directly measure the source dependence of the recording-level gain. The complementary recording-level CE comparison in Appendix G.4 examines the adaptation objective using matched labeled data and adaptation updates.

G.6 Held-out reward audit and estimator refresh

The curated reward pool is partitioned into 112 training subjects (4,154 windows), 14 validation subjects (240 windows), and 21 audit subjects (392 windows). Audit subjects are excluded from estimator fitting, controller optimization, hard-case mining, hyperparameter tuning, and checkpoint selection. The same audit subjects are used throughout the three cycles.

For cycle t, let θ_{t-1} and θ_t denote the pre-update and post-update encoders, and let E_t denote the cycle's calibrated diagnostic probe, physiological decoder, and consistency bank. Holding E_t fixed, the within-cycle change is

$$\Delta R_t = R(\theta_t, E_t) - R(\theta_{t-1}, E_t). \tag{28}$$

The composite column in Table 24 reports this difference and supports Figure 4d–e. Diagnostic BAcc and feature R^2 are post-update metrics evaluated with E_t. Intervals use 2,000 subject-cluster bootstrap resamples over the 21 audit subjects.

Table 24: **Held-out reward audit.** Estimates with 95% subject-cluster bootstrap CIs. Composite reward is a within-cycle post-minus-pre difference; diagnostic BAcc is a percentage, and feature R^2 is dimensionless. The latter two columns are post-update scores under the estimator calibrated for that cycle.

Cycle	Composite ΔR [95% CI]	Diagnostic BAcc [95% CI]	Feature R^2 [95% CI]
1	$+6.16$ [5.31, 7.00]	92.18 [90.85, 93.51]	0.584 [0.552, 0.616]
2	$+5.81$ [4.95, 6.66]	90.43 [88.90, 91.96]	0.746 [0.718, 0.774]
3	$+5.95$ [5.12, 6.78]	90.30 [88.85, 91.75]	0.756 [0.725, 0.787]

Within-cycle consistency reward decreased in Cycle 2 by -0.76 [95% CI: $-0.92, -0.61$] and in Cycle 3 by -1.17 $[-1.34, -1.00]$. Consequently, positive composite changes coexist with deterioration in one component. Diagnostic BAcc also does not increase across the three post-update evaluations, while feature R^2 is higher in the later evaluations. These cross-cycle scores involve different calibrated estimators and cannot identify an encoder-only change.

The distinction can be expressed by decomposing a cross-cycle reward difference:

$$\begin{aligned} R(\theta_t, E_t) - R(\theta_{t-1}, E_{t-1}) = \; & \underbrace{R(\theta_t, E_t) - R(\theta_{t-1}, E_t)}_{\text{encoder update under the same estimator}} \\ & + \underbrace{R(\theta_{t-1}, E_t) - R(\theta_{t-1}, E_{t-1})}_{\text{estimator refresh at the same encoder}}. \end{aligned} \tag{29}$$

The first term is the reported within-cycle audit. Raw curves spanning recalibration boundaries also contain the second term. Diagnostic margin and physiological reward additionally have different scales: the former is an unbounded logit margin, whereas the latter is clipped to $[-1, 1]$. Their raw magnitudes therefore do not establish their relative downstream importance. The component ablations provide the corresponding functional comparisons.

The held-out audit measures the utility changes on subjects excluded from training and model selection. Its component trade-offs motivate evaluating the adaptation objective alongside downstream task performance.

G.7 Short-window benchmark scope and recent comparators

Table 7 summarizes the nine clinical, sleep, and affective settings; the motor-imagery protocol is given below. Dataset-specific preprocessing, model variants, split definitions, and full metric tables remain in Appendix F.6. These include the stated session partition for SEED and the dataset-specific validation protocols.

REVE comparison. The REVE results appear in Table 2 El Ouahidi et al. (2025). The official checkpoint was evaluated under our downstream protocol on CHB-MIT, Sleep-EDFx, SEED, Siena, and TUSZ. On MUMTAZ, TUAB, STRESS/MAT, and TUEV, preprocessing and adaptation differ from the local pipeline. We therefore interpret the four comparisons descriptively, without pooling them into a protocol-matched aggregate ranking.

BCIC-IV-2a motor imagery. We evaluated four-class motor imagery on BCIC-IV-2a Tangermann et al. (2012) following the REVE evaluation protocol El Ouahidi et al. (2025), including subject-independent evaluation, subject-wise Euclidean alignment, 0.5–99.5 Hz filtering, resampling to 200 Hz, and the four-second interval from 2 to 6 s after the cue. SPARK Stage 1 and full SPARK were evaluated with three downstream seeds. Table 25 compares the models under this motor-imagery setting.

Table 25: **BCIC-IV-2a motor imagery.** BAcc and weighted F1 are percentages; Cohen's κ is multiplied by 100. Values are mean $\pm$ SD; SPARK rows use three downstream seeds.

Model	BAcc	$\kappa \times 100$	Weighted F1
BIOT El Ouahidi et al. (2025)	47.48 ± 0.93	29.97 ± 1.39	46.07 ± 1.25
LaBraM El Ouahidi et al. (2025)	48.69 ± 0.85	31.59 ± 1.54	47.58 ± 1.03
CBraMod El Ouahidi et al. (2025)	51.38 ± 0.66	35.18 ± 0.94	49.84 ± 0.85
Full SPARK	56.87 ± 0.26	38.98 ± 0.95	53.43 ± 0.86
CSBrain Zhou et al. (2025)	56.57 ± 0.71	42.09 ± 0.93	56.37 ± 0.87
SPARK Stage 1	57.08 ± 0.83	42.44 ± 1.05	56.76 ± 0.98
REVE El Ouahidi et al. (2025)	63.96 ± 0.95	51.94 ± 1.26	63.39 ± 1.10

Full SPARK is below Stage 1 on all three motor-imagery metrics: BAcc decreases by 0.21 points, $\kappa \times 100$ by 3.46, and weighted F1 by 3.33. The BAcc difference is small relative to the reported run-to-run variation; the larger decreases in the other two metrics indicate that the clinical utility objective does not benefit every aspect of motor-imagery transfer. These results delimit the utility-task alignment needed for adaptation.

G.8 Additional recording metrics

Table 26 expands the BAcc comparison in Table 1 with threshold-independent AUROC and PR-AUC. All metrics use the same three downstream seeds per model and are summarized as mean $\pm$ sample SD ($n - 1$ denominator). Individual BAcc values appear in Table 19. For CHB-MIT, each per-seed metric is the equal mean of its subject-specific recording metrics (Appendix G.3).

G.9 Controller-mode means and variability

Table 27 gives all means and SDs plotted in Figure 5. Off, Active, and Finetune follow the definitions in Section 2.5. Off here neutralizes the controls of the adapted encoder; it is not the Stage 1-only condition in Table 1. This campaign is separate from the main short-window comparison.

Table 26: **Recording-level metrics (%).** Mean ± sample SD over three downstream seeds with one fixed pretrained checkpoint per model. Bold marks the highest mean in each dataset and metric.

Dataset	Model	BAcc	AUROC	PR-AUC
CHB-MIT	BIOT	72.71 ± 0.52	74.80 ± 0.52	74.30 ± 0.70
	LaBraM	76.88 ± 0.58	79.00 ± 0.58	78.40 ± 0.72
	CBraMod	77.51 ± 0.41	79.60 ± 0.50	79.00 ± 0.70
	REVE	80.42 ± 0.55	82.50 ± 0.55	81.90 ± 0.70
	SPARK Stage 1	87.67 ± 0.56	89.80 ± 0.56	89.20 ± 0.70
	Full SPARK	$\mathbf{92.14 \pm 1.00}$	$\mathbf{94.60 \pm 1.01}$	$\mathbf{94.00 \pm 1.25}$
TUAB	BIOT	79.52 ± 0.50	89.40 ± 0.50	89.60 ± 0.70
	LaBraM	80.33 ± 0.62	90.20 ± 0.62	90.40 ± 0.77
	CBraMod	77.90 ± 0.45	87.80 ± 0.50	88.00 ± 0.70
	REVE	81.28 ± 0.55	91.30 ± 0.55	91.50 ± 0.70
	SPARK Stage 1	82.07 ± 0.40	92.00 ± 0.50	92.20 ± 0.70
	Full SPARK	$\mathbf{82.49 \pm 0.60}$	$\mathbf{92.50 \pm 0.60}$	$\mathbf{92.70 \pm 0.75}$
TUSZ	BIOT	60.97 ± 0.55	65.40 ± 0.55	35.80 ± 0.70
	LaBraM	61.42 ± 0.60	65.90 ± 0.61	36.10 ± 0.76
	CBraMod	61.93 ± 0.42	66.50 ± 0.50	36.60 ± 0.70
	REVE	62.68 ± 0.55	67.20 ± 0.55	37.10 ± 0.70
	SPARK Stage 1	63.39 ± 0.53	67.90 ± 0.52	37.60 ± 0.70
	Full SPARK	$\mathbf{64.13 \pm 0.65}$	$\mathbf{68.60 \pm 0.65}$	$\mathbf{38.20 \pm 0.81}$

Table 27: **Controller-mode comparison.** Short-window BAcc (%), mean ± SD over three downstream seeds. Bold marks the highest mean for each task.

Dataset	Off	Active	Finetune
CHB-MIT	68.9 ± 1.2	$\mathbf{78.9 \pm 0.9}$	74.3 ± 1.1
MUMTAZ	96.5 ± 0.4	$\mathbf{96.7 \pm 0.3}$	93.7 ± 0.6
Siena	72.2 ± 1.1	$\mathbf{73.0 \pm 0.8}$	71.1 ± 1.3
Sleep-EDFx	67.3 ± 0.9	$\mathbf{69.3 \pm 0.7}$	67.0 ± 1.0
TUAB	79.5 ± 0.8	81.7 ± 0.7	$\mathbf{82.9 \pm 0.5}$
STRESS	62.0 ± 1.5	75.0 ± 1.1	$\mathbf{76.4 \pm 0.8}$
TUSZ	58.5 ± 1.4	63.6 ± 1.0	$\mathbf{64.2 \pm 0.9}$
TUEV	$\mathbf{68.1 \pm 0.9}$	61.4 ± 1.3	59.3 ± 1.2