Brain Oscillator Intrinsic Dimension marks consciousness, recovery mode, and failed downshifting in coma

Alexandra Bernadotte *[1] and Ivan Menshikov [1]

[1]Aicumene†

A Preprint
April 12, 2026

One-sentence summary

Conscious engagement is high-dimensional, recovery is low-dimensional, and poor-prognosis coma reflects pathological failure of the injured brain to downshift into restorative low-dimensional dynamics.

Abstract

Consciousness is usually defined by behavior, reportability, or spectral brain rhythms, but the dimensionality of the underlying brain dynamics remains poorly specified. Here we introduce brain oscillator intrinsic dimension (BOID): an intrinsic dimension of brain time series defined as the number of active exposinusoidal oscillatory degrees of freedom recovered from multichannel Hankel-delay embeddings by singular-spectrum analysis and ESPRIT pole extraction. We test a regime-switching hypothesis of consciousness: conscious engagement is high-dimensional and resource-demanding; healthy recovery and efficient control are low-dimensional; and poor-prognosis coma reflects pathological failure to downshift from high-dimensional overactivation into restorative low-dimensional dynamics.

Across independent datasets spanning human sleep, mouse sleep, and propofol anesthesia, reversible suppression of consciousness was accompanied by a robust reduction in BOID. In Sleep-EDF, oscillator count decreased monotonically from wakefulness to deep slow-wave sleep (median 18 to 4 modes per 4-s window, $p = 1.0 \times 10^{-8}$), a pattern reproduced across individual nights and confirmed

*Corresponding author. Email: bernadotte@aicumene.com
†Aicumene. https://research.aicumene.com

by independent fractal, graph-based, Hankel-rank, and topological estimators. Mouse non-REM sleep replicated the low-dimensional sleep state across species. Propofol sedation similarly reduced BOID relative to wakefulness (30.5 to 21.5 modes, $p = 3.2 \times 10^{-3}$), despite the organized propofol alpha signature, showing that BOID measures the size of the active oscillator repertoire rather than the power of any single frequency band.

Low dimensionality was not pathological. Expert meditators showed lower intrinsic dimension than novices, indicating that trained mental states can be maintained with fewer active dynamical degrees of freedom. We therefore interpret deep sleep, anesthesia, and expert meditation as organized low-dimensional recovery or efficiency modes: states in which conscious access narrows while the brain enters a more synchronized, economical, and controllable regime.

Post-anoxic coma exposed the pathological boundary of this mechanism. In 575 I-CARE patients, poor neurological outcome was associated with elevated BOID and a pathological mode-proliferation index combining high mode count, high-frequency poles, burst intermittency, and mode-frequency entropy (PMPI AUC 0.67, $p < 10^{-3}$). Thus, high dimensionality in coma did not indicate preserved consciousness; it marked hyperexcitable, fragmented dynamics and failure of the injured brain to enter a protective low-dimensional recovery state. Incorporating BOID into a multimodal prognostic pipeline with intrinsic-dimension, spectral/burst, and clinical features yielded leakage-free coma-outcome prediction comparable to state-of-the-art models (AUC 0.81), while preserving mechanistic interpretability.

These results identify BOID as a geometric marker of brain-state regime. Conscious wakefulness is high-dimensional; healthy recovery and expert control are low-dimensional; and poor-prognosis coma is a failure of dimensional regulation. Brain oscillator dimensionality therefore provides a single interpretable read-out linking sleep, anesthesia, meditation, and coma.

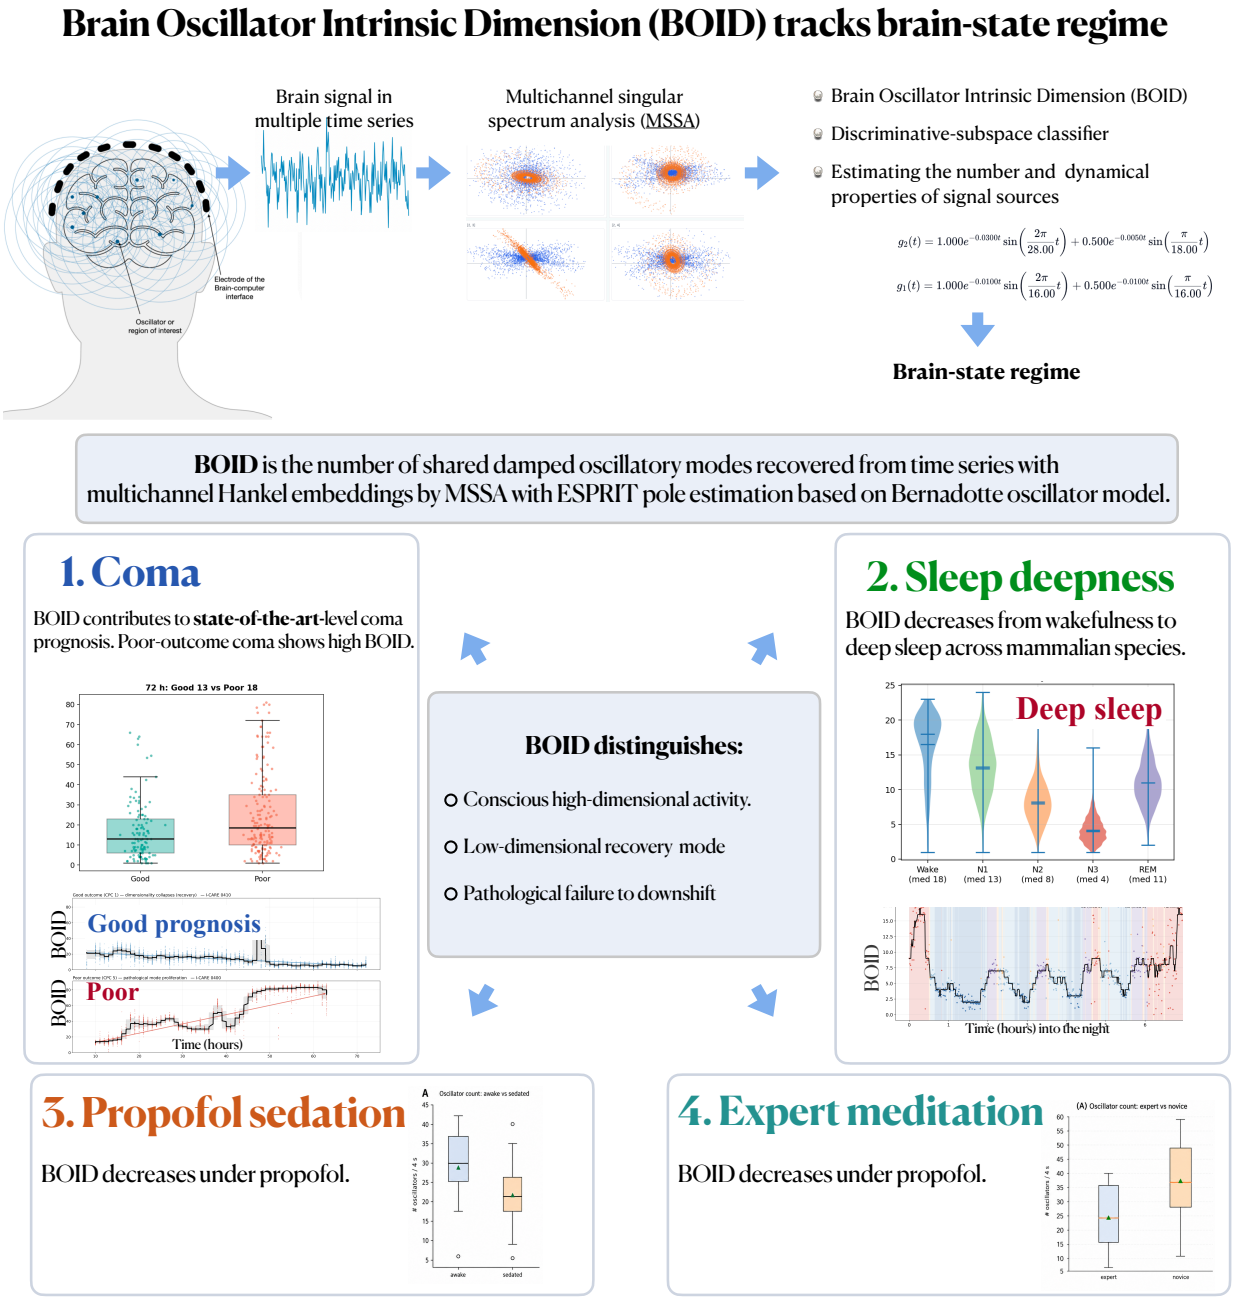

1 Introduction

Consciousness is not simply the presence of brain activity. It is an organized, resource-demanding dynamical regime in which the brain maintains a broad repertoire of possible perceptions, actions, memories, predictions, and internal simulations. This repertoire should have a measurable geometric signature. We propose that conscious engagement is characterized by high oscillator dimensionality:

a large number of dynamically active oscillatory degrees of freedom are required to generate the observed brain signal.

The converse is equally important. Low dimensionality is not necessarily failure. Healthy brains repeatedly enter low-dimensional states. Deep sleep, pharmacological sedation, and some forms of trained meditative stabilization reduce conscious access, but they also reduce metabolic and dynamical load. These states are slower, more synchronized, and more economical. We therefore interpret low oscillator dimensionality as a recovery or efficiency mode of the healthy brain: a regime in which the conscious repertoire narrows so that the system can stabilize, restore, and prepare for future high-dimensional activity.

This view differs from monotonic complexity theories in which "more complexity" is implicitly better. High dimensionality can be adaptive or pathological depending on biological context. In the intact brain, high-dimensional dynamics support conscious interaction, cognitive flexibility, and mental effort. In the injured brain, especially after anoxic damage, high dimensionality may instead reflect pathological mode proliferation: epileptiform activity, burst-like fragmentation, unstable high-frequency poles, myoclonus, and metabolically costly hyperexcitability. Poor-prognosis coma may therefore result not from unconsciousness itself, but from failure of the injured brain to downshift into an organized low-dimensional recovery state.

Sleep and anesthesia provide the physiological basis for this regime-switching hypothesis. Deep non-REM sleep is dominated by slow synchronized activity, and propofol anesthesia replaces awake high-frequency, low-amplitude EEG with structured alpha and slow-delta oscillations. Excessive anesthetic depth can produce burst suppression, a pattern also observed in anoxic coma and epilepsy [1, 2]. Thus, reversible loss of consciousness is not merely a reduction in activity. It is an organized dynamical downshift. The central question is whether the brain can enter and leave this low-dimensional state appropriately.

To formalize this question, we introduce brain oscillator intrinsic dimension (BOID), extending our previous topology-driven framework for time-series classification, where the intrinsic dimension and geometry of delay embeddings were proposed as invariants of the underlying dynamics and were linked to the number of latent dynamical components [3, 4, 5, 6]. Here, we specialize this principle to brain time series by defining BOID as the number of active exposinusoidal oscillatory degrees of freedom recovered from multichannel Hankel embeddings.

BOID is based on the Bernadotte oscillator model of brain signals [3, 5, 6]. A multichannel recording $X(t) \in \mathbb{R}^m$ is embedded as a Hankel trajectory $\mathcal{H}_{L,\tau}(X) \subset \mathbb{R}^{mL}$ [7, 4, 3, 5, 6]. In the oscillatory model, the observed signal is represented as a finite superposition of damped oscillatory components,

$$X(t) = \sum_{k=1}^{K} a_k e^{\lambda_k t} \sin(\omega_k t + \varphi_k) + \varepsilon(t), \tag{1}$$

where each component has a spatial projection a_k, frequency ω_k, phase φ_k, damping or growth parameter λ_k, and noise or unresolved activity $\varepsilon(t)$. Introduced Multichannel singular-spectrum

analysis (MSSA) followed by ESPRIT pole extraction [8] estimates these components. We interpret the recovered mode count as BOID: an intrinsic dimension of the reconstructed trajectory and an estimate of the active oscillatory degrees of freedom generating the brain state.

The Bernadotte oscillator model is complementary to classical oscillator, spiking-neuron, and information-theoretic frameworks. In relation to Kuramoto-type models, it treats large-scale brain activity as an observable superposition of coupled oscillatory modes, but instead of prescribing only phase synchronization within a predefined oscillator population [9, 10], it estimates from data the active damped exposinusoidal components, including their frequencies, phases, damping parameters, and shared multichannel projections. In relation to Izhikevich-type models, it does not attempt to reproduce the full microscopic spike-generating dynamics of individual neurons [11, 12]; rather, it provides a mesoscopic inverse model of the collective oscillatory degrees of freedom visible in EEG and neurointerface signals. From the viewpoint of information theory, BOID can be interpreted through the principle of minimum description length [13, 14]: among competing reconstructions of the signal, the relevant descriptor is the smallest set of oscillatory modes that explains the observed dynamics with sufficient accuracy. Thus, BOID measures the effective complexity of a brain state as the minimal oscillatory description required to encode its dynamics.

Oscillator dimensionality is distinct from spectral power

BOID is conceptually distinct from conventional spectral-band measures such as alpha, theta, or delta power. Band power quantifies how much signal energy falls within a predefined frequency range. Alpha power, for example, measures the amplitude or energy of EEG activity near 8–13 Hz. BOID instead estimates how many dynamically independent oscillatory modes are required to generate the observed brain signal. It is therefore a measure of the size of the active oscillator repertoire, not the strength of any single rhythm.

This distinction is essential for interpreting sleep and anesthesia. A brain state can show strong organized alpha or slow-delta activity while having low BOID, because the signal is dominated by a small number of synchronized oscillatory generators. Conversely, active wakefulness or cognitive effort may suppress alpha power while increasing BOID, because the brain recruits a broader set of independent oscillatory degrees of freedom. Thus, BOID does not ask which frequency band is strongest; it asks how many oscillatory degrees of freedom are active.

In this sense, BOID provides a geometric and dynamical descriptor that is complementary to spectral analysis. Spectral power describes the distribution of energy across frequencies. BOID describes the dimensionality of the latent oscillatory system that generates the signal. This allows BOID to capture regime changes that are not reducible to alpha power, delta dominance, or any single frequency-band amplitude.

BOID as a marker of consciousness and recovery-mode switching

We test this framework across open brain time-series datasets spanning physiological sleep, cross-species sleep, pharmacological anesthesia, expert meditation, and post-anoxic coma. Human sleep tests whether loss of conscious access during normal sleep is accompanied by dimensional downshifting. Mouse sleep tests whether this low-dimensional recovery mode is conserved across mammalian species. Propofol anesthesia tests whether pharmacological suppression of consciousness reduces oscillator dimensionality in the intact brain. Meditation tests whether trained mental states can enter a low-dimensional efficient regime while awake. Coma tests the pathological boundary of the mechanism: whether the injured brain fails to downshift into a restorative low-dimensional state, and whether this failure carries prognostic information.

Our results support a regime-switching model of consciousness. Wakefulness is high-dimensional. Sleep and anesthesia are low-dimensional recovery states. Expert meditation is a low-dimensional efficient waking state, consistent with reduced resource expenditure by a trained brain. Poor-outcome coma is characterized by pathological high-dimensional overactivation and failure of recovery-mode downshifting. Finally, we show that BOID is not only interpretable but clinically useful: when fused with intrinsic-dimension, spectral, burst-suppression, and clinical features, it contributes to a high-performing, leakage-free model of coma outcome.

2 Results

2.1 Conscious wakefulness is high-dimensional and deep sleep is a low-dimensional recovery mode

Sleep provides a natural within-brain experiment in dimensional regime switching. Across the Sleep-EDF cohort, brain oscillator intrinsic dimension (BOID) reproducibly decreased with sleep depth. Wakefulness showed the highest oscillator dimensionality, progressively deeper NREM stages showed a stepwise reduction, and N3 slow-wave sleep reached the lowest values. This monotonic decline was observed across subjects rather than being driven by isolated recordings, indicating that BOID behaves as a robust marker of consciousness level during physiological sleep. In this interpretation, the transition from wakefulness to deep sleep corresponds to a reproducible downshift from a high-dimensional conscious regime into a low-dimensional recovery mode.

We analyzed Sleep-EDF Expanded recordings using MSSA/ESPRIT extraction across expert-scored sleep stages. The analyzed subset comprised 10 subjects, 5,535 four-second windows, and 56,282 extracted oscillatory modes. The oscillator population changed systematically with sleep stage (Figs. 1, 2). Subject-level median BOID decreased from Wake (18) to N1 (13), N2 (8), and N3 (4), whereas REM occupied an intermediate position (11). N1 denotes the lightest transition from wakefulness to sleep, N2 is stable light sleep characterized by sleep spindles and K-complexes, and N3 is deep slow-wave sleep dominated by delta activity. The stage effect was highly significant

(Kruskal–Wallis $H = 43.0$, $p = 1.0 \times 10^{-8}$).

This result supports the first thesis of the paper: conscious wakefulness is high-dimensional in terms of BOID. As sleep deepens, conscious access is progressively reduced and the oscillator population contracts. This low-dimensionality is not pathological. In healthy sleep, it marks a recovery mode. N3 oscillators were concentrated at 1–2 Hz, consistent with delta-dominant slow-wave sleep, whereas wakefulness showed the broadest frequency distribution and extended into the 15–30 Hz range. Oscillatory power during sleep was dominated by slow activity, especially in N3 and REM. Thus, deepening sleep produced a coordinated reduction in active oscillatory degrees of freedom and a redistribution of activity toward slow restorative dynamics.

This pattern was reproducible across individual whole-night recordings rather than being driven by a single subject or pooled-window statistics (Fig. S11).

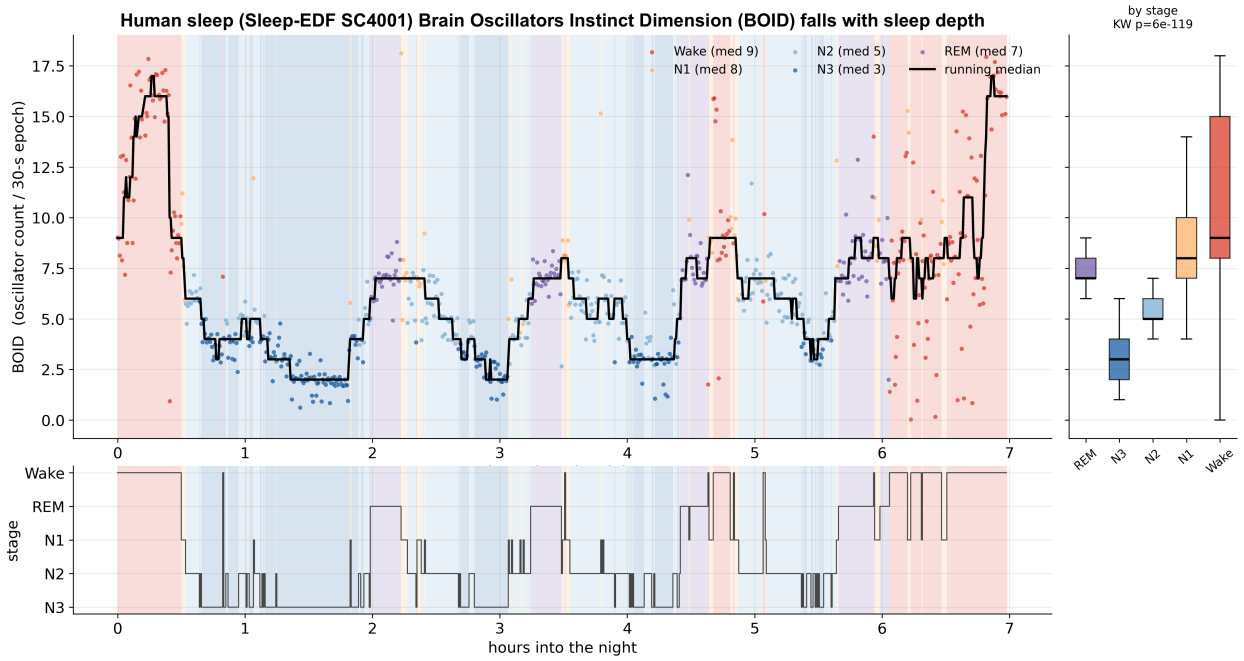

Figure 1: BOID follows sleep architecture across a whole night. Representative Sleep-EDF recording SC4001. The upper panel shows brain oscillator intrinsic dimension (BOID), defined as oscillator count per 30-s epoch, across approximately 7 h of sleep. Points are colored by manually scored sleep stage; the black curve shows the running median. Background shading indicates the hypnogram-derived sleep stage at each time point. BOID is highest during wakefulness at sleep onset and final awakening, decreases during non-REM sleep, and reaches its lowest values during N3 slow-wave sleep. REM and N1 occupy intermediate dimensional regimes. The lower panel shows the corresponding hypnogram, confirming that dimensional downshifts align with transitions into deeper non-REM sleep and that dimensional increases align with REM or wake transitions. The right panel summarizes BOID by sleep stage in the same recording, showing the expected ordering: Wake highest, N3 lowest, and REM/N1/N2 intermediate. This representative night illustrates that oscillator dimensionality is a continuous state variable tracking sleep depth and the transition between high-dimensional conscious activity and low-dimensional recovery mode.

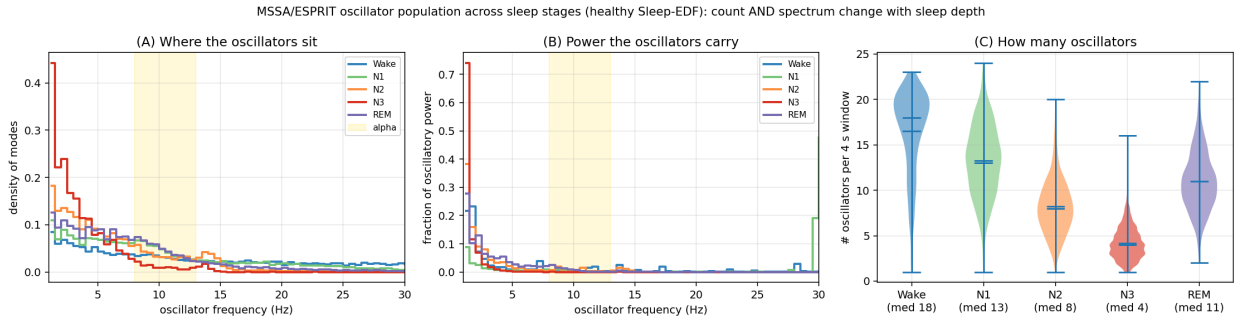

Figure 2: Healthy sleep downshifts in BOID. BOID analysis of Sleep-EDF recordings. **(A)** Oscillator-frequency distributions shift toward slow frequencies with sleep depth. **(B)** Oscillatory power is dominated by slow activity, especially in deep sleep (N3) and REM. **(C)** BOID decreases from wakefulness to deep sleep (N3), with REM intermediate.

2.2 Independent estimators confirm the low-dimensional state of deep sleep

To test whether the sleep-depth effect depended specifically on the Hankel-spectral oscillator representation of BOID, we computed an independent intrinsic-dimension battery on the same Sleep-EDF subjects [15, 16, 17, 18]. Scalar estimators from distinct mathematical families converged on the same conclusion: N3 slow-wave sleep was the lowest-dimensional regime (Fig. 3; Table 1). Correlation dimension, false-nearest-neighbor embedding dimension, Higuchi fractal dimension, minimum-spanning-tree length, and Brito–Quiroz–Yukich degree all reached their minima in N3. Because these estimators probe different representations of the data, including raw single-channel signals and multichannel delay-embedded trajectories, their agreement provides convergent evidence that sleep-related dimensional downshifting is not an artifact of the BOID extraction pipeline.

Topological descriptors revealed a complementary axis of sleep organization. First-homology loop count and persistence entropy peaked in REM rather than in N3. REM therefore occupied an intermediate scalar-dimensional regime but showed richer cyclic trajectory structure, consistent with internally active sleep. Thus, scalar dimensionality and topology separate two different properties of sleep dynamics: N3 is low-dimensional and restorative, whereas REM is topologically richer and internally structured.

Table 1: Intrinsic-dimension battery across Sleep-EDF stages. Values are subject-level medians. Scalar estimators are minimized in N3 slow-wave sleep, whereas the topological H_1 loop count peaks in REM.

Family	Estimator	Wake	N2	N3	REM	p
Fractal	Correlation dimension	6.25	5.75	4.56	6.53	6×10^{-6}
Fractal	False-nearest-neighbor	8.0	4.0	4.0	4.0	9×10^{-8}
Fractal	Higuchi FD	1.71	1.53	1.32	1.59	2×10^{-8}
Graph	MST edge length	13.1	7.1	5.1	8.4	4×10^{-9}
Graph	BQY degree	11.0	9.5	5.0	10.3	1×10^{-7}
Topology	H_1 loop count	8.5	12.5	14.5	21.8	2×10^{-4}

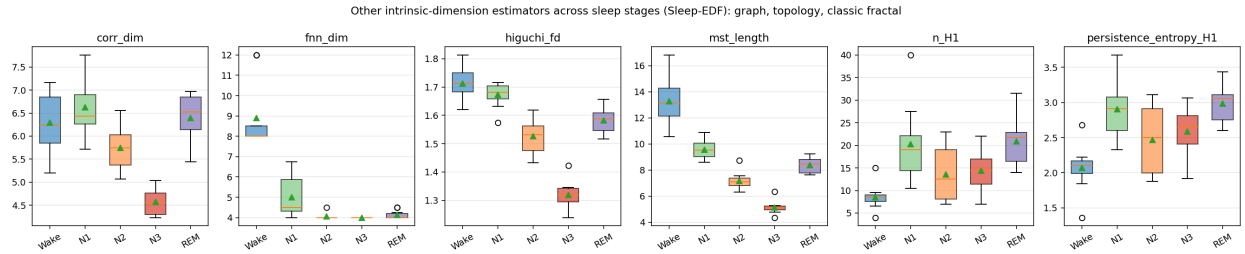

Figure 3: Independent estimators confirm the deep-sleep dimensional minimum. Fractal and graph-based estimators reach their lowest values in N3 slow-wave sleep, confirming the BOID-defined dimensional downshift with sleep depth. In contrast, topological loop descriptors peak in REM, showing that scalar dimension and trajectory topology capture complementary aspects of sleep-state organization.

2.3 Sleep-related dimensional downshifting is conserved across species

Having established a graded dimensional downshift across human sleep stages, we next asked whether the same effect generalizes beyond humans. We applied the BOID pipeline to OpenNeuro ds006366, the Mouse Sleep Staging Validation dataset, which contains recordings from 92 mice across five laboratories with single-channel cortical EEG and EMG sampled at 128 Hz [19]. Because the integer stage codes were not documented, we identified stages physiologically: the code with maximal delta power was assigned to non-REM sleep, and the code with maximal theta and minimal delta power was assigned to REM.

Despite differences in species, brain size, recording montage, and the use of a single cortical EEG channel, mouse sleep reproduced the human pattern (Fig. 4). BOID was lowest in non-REM sleep, whereas wakefulness and REM occupied higher-dimensional regimes. Median oscillator count decreased from 7 in wakefulness to 6 in non-REM sleep and returned to 7 in REM, with a highly significant stage effect (Kruskal–Wallis $p = 1.7 \times 10^{-11}$; paired Wake versus non-REM $p = 2.6 \times 10^{-12}$). Independent single-channel estimators agreed: Higuchi fractal dimension, singular-spectrum entropy, and participation ratio all reached their minima in non-REM sleep.

This cross-species replication indicates that sleep-related dimensional downshifting is not a peculiarity

of human EEG montage, Sleep-EDF scoring, or a single dataset. Rather, it reflects a conserved dynamical transition of the mammalian brain into a low-dimensional sleep state. Together with the Sleep-EDF results, the mouse data support the interpretation of non-REM sleep as an organized low-dimensional recovery mode across species.

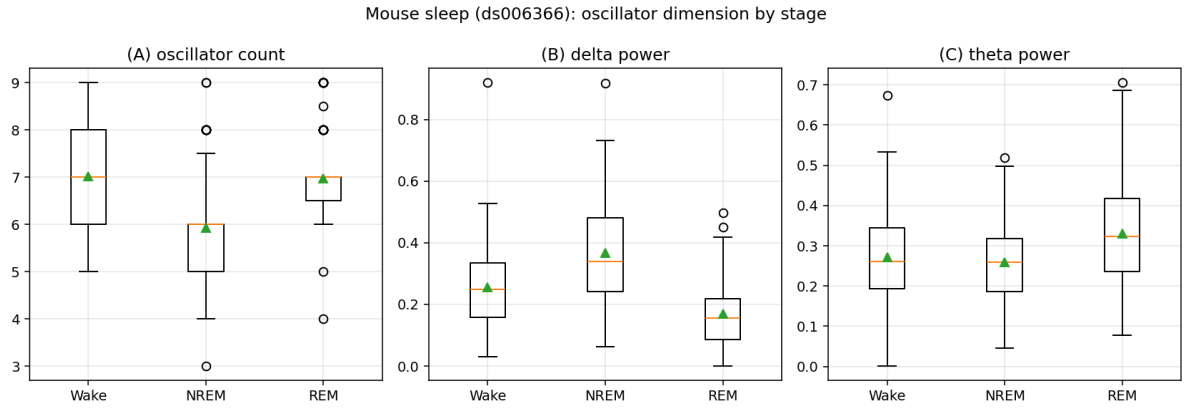

Figure 4: Cross-species validation in mouse sleep. Mouse sleep EEG reproduces the human sleep-dimensionality pattern. BOID is lowest in non-REM sleep, whereas wakefulness and REM occupy higher-dimensional regimes. Delta power peaks in non-REM and theta power peaks in REM, supporting the physiological identification of the otherwise undocumented sleep-stage codes.

This cross-species pattern was reproducible across individual mouse recordings rather than being driven by a single animal or pooled-epoch statistics (Fig. S12).

2.4 Anesthesia lowers BOID during pharmacological suppression of consciousness

Propofol sedation provided a pharmacological test of the same dimensional-switching principle. We analyzed OpenNeuro ds005620, a repeated-awakening propofol study with 21 healthy adults and 65-channel EEG [20]. Within subjects, BOID fell from a median of 30.5 active oscillators per window during wakefulness to 21.5 under sedation (paired Wilcoxon $p = 3.2 \times 10^{-3}$, $n = 20$; Fig. 5). This reduction occurred together with the canonical propofol EEG signature: anteriorized alpha activity, increased peak alpha frequency, and an increased alpha-to-theta ratio.

The repeated-awakening design allowed us to distinguish the global level of consciousness from the momentary content of experience. Among sedated awakenings, BOID did not separate reports with experience from reports without experience (medians 23 versus 26; Mann–Whitney $p = 0.98$). Thus, BOID tracked the global dynamical state of the brain "awake versus sedated" more robustly than the presence or absence of reportable experience within sedation.

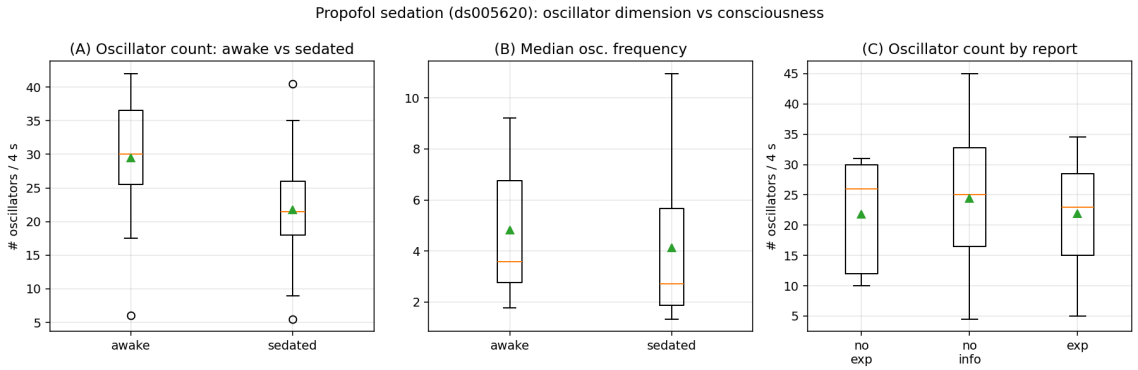

Figure 5: Propofol sedation reduces BOID during pharmacological suppression of consciousness. BOID decreases from wakefulness to sedation in healthy adults. Within sedated awakenings, BOID does not distinguish reports with experience from reports without experience, indicating that oscillator dimensionality tracks the global level of consciousness more robustly than momentary reportable content.

These results support the interpretation that high BOID marks conscious, effortful, high-resource brain activity in the intact brain. Propofol suppresses consciousness by shifting the brain into a lower-dimensional, slower, more synchronized regime. As in deep sleep, this low-dimensional state should not be interpreted as pathological collapse; rather, it represents an organized pharmacological downshift of consciousness.

2.5 Concentration did not change geometry, but expert meditation reveals low-dimensional efficient waking control

We examined the Brandmeyer-Delorme experience-sampling dataset (OpenNeuro ds001787) [21], in which 12 expert and 12 novice meditators were interrupted during meditation and rated their mind-wandering.

Within subjects, reported momentary concentration did not change geometry. None of the estimators differed across mind-wandering levels, and BOID was nearly unchanged. Intrinsic dimension in this dataset therefore tracked the trait of expertise rather than momentary attentional content. Experienced meditators showed lower intrinsic dimension than novices across the estimator battery: correlation dimension (6.3 versus 8.9, $p = 2 \times 10^{-3}$), false-nearest-neighbor dimension (4.9 versus 7.4, $p = 0.017$), Levina-Bickel estimate (5.9 versus 6.5, $p = 0.030$), singular-spectrum entropy (145 versus 172, $p = 0.035$), and BOID at trend level (Fig. 6).

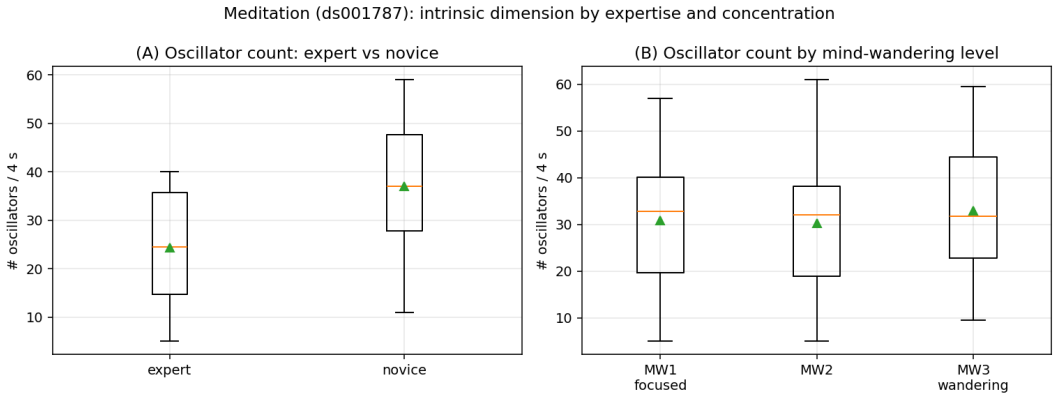

Figure 6: Meditation expertise lowers intrinsic dimension. (A) Oscillator count (BOID) is lower in experts than in novices. (B) Oscillator count (BOID) does not differ across reported mind-wandering levels.

This result is essential for the theory. High dimensionality reflects mental tension and resource-demanding conscious activity; trained brains can achieve specialized mental states with fewer active degrees of freedom (a low-dimensional, synchronized control regime). Professional meditators therefore appear able to enter a low-dimensional waking recovery or efficiency mode (almost a sleep-like state during wakefulness) rather than a pathological collapse.

2.6 Poor-prognosis coma reflects failure to enter low-BOID recovery mode

Healthy sleep and propofol sedation show that an intact brain can downshift into low-BOID organized states. Acute post-anoxic coma tests what happens when this regulation fails. We analyzed EEG from all available comatose survivors of cardiac arrest in PhysioNet I-CARE [22]: 575 patients and 826 recordings at 24–72 hours post-arrest, dichotomized into Good (CPC 1–2) and Poor (CPC 3–5) neurological outcome (Fig. 7).

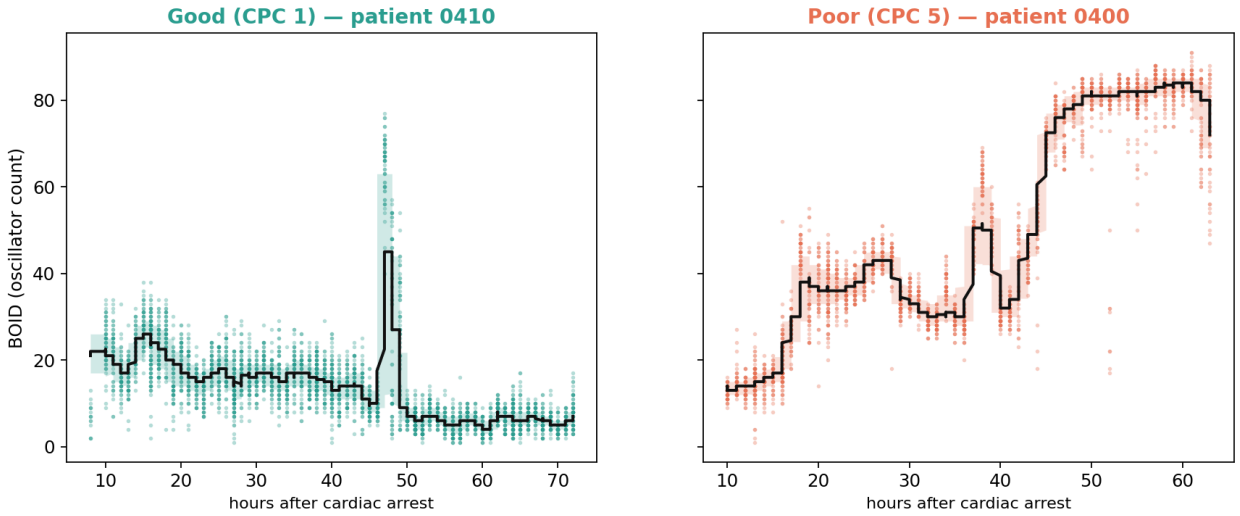

Figure 7: Brain Oscillator Intrinsic Dimension (BOID) of the post-anoxic coma EEG tracks neurological outcome. Brain Oscillator Intrinsic Dimension (BOID) computed in every non-overlapping 4-s EEG window and plotted against time since cardiac arrest (8–72 h) for two representative I-CARE patients. (Top) Good outcome (patient 0410, Cerebral Performance Category 1): BOID is low throughout and declines over the recording (linear fit, blue; 20 to 6 modes), a progressive collapse of the oscillatory state space accompanying recovery (the transient excursion near 47 h is an artifactual burst). (Bottom) Poor outcome (patient 0400, Cerebral Performance Category 5): BOID rises 6-fold (14 to 80 modes; trend arrow, red), a proliferation of pathological oscillatory modes. Points, individual 4-s windows; black line, 10-min running median; shaded band, interquartile range.

In this injured-brain context, higher BOID marked worse outcome (Fig. 8). The strongest single descriptor was a pre-specified pathological mode-proliferation index (PMPI) combining high mode count, high-frequency poles, burst intermittency, and mode-frequency entropy. At the patient level, PMPI was higher in poor-outcome patients (AUC 0.67, one-sided Mann-Whitney $p < 10^{-3}$), as was bare BOID (AUC 0.60, $p < 10^{-3}$). The direction was present at both 24 hours and 72 hours.

This is not a contradiction of the sleep and anesthesia results. It is the pathological version of the same switching framework. In healthy physiology, low dimensionality accompanies reversible loss of consciousness and recovery. After anoxic injury, recovery may require the brain to stabilize into a sleep-like low-dimensional unconscious regime before it can re-emerge into high-dimensional conscious activity. Poor-prognosis coma is instead marked by failure of this downshift: the EEG remains high-dimensional because the injured brain is overheated, hyperexcitable, and fragmented. Epileptiform discharges, myoclonus, burst-suppression bursts, and unstable high-frequency poles increase the detected mode count without representing preserved consciousness. Thus, high dimensionality in coma is a poor prognostic sign because it reflects the inability of the injured brain to enter a restorative low-dimensional unconscious state.

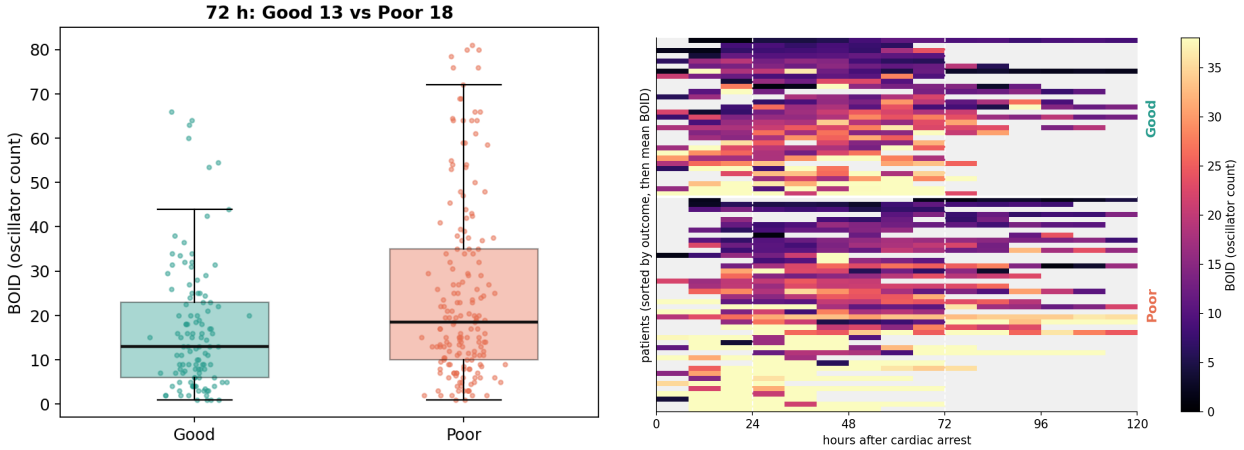

Figure 8: Poor-prognosis coma is marked by pathological high-dimensional proliferation. In PhysioNet I-CARE, PMPI and oscillator count (BOID) are elevated in poor-outcome patients. The high-dimensional pattern is interpreted as pathological overactivation and failed downshifting into a restorative low-dimensional unconscious state, not as preserved consciousness.

2.7 A low-dimensional Bernadotte discriminant subspace carries coma outcome.

Beyond the scalar dimensionality markers, we asked whether the *geometry* of the early-EEG signal subspace itself separates outcomes, using the Bernadotte discriminative-subspace classifier (Algorithm 1). Across 500 patients (195 Good / 305 Poor; first $\leq 24\,\mathrm{h}$, $\sim 4\,\mathrm{min}$ of 19-channel EEG), the difference operator $\Delta = P_{\mathrm{Good}} - P_{\mathrm{Poor}}$ has a strikingly concentrated spectrum: of its 228 eigenvalues only a handful are appreciably non-zero—a few Good-characteristic and a few Poor-characteristic directions over an otherwise shared, near-zero bulk (Fig. 9C). Restricting the classifier to this low-dimensional discriminant subspace ($k \approx 3$–6 of 228) separates the outcomes (Fig. 9A), the signed-energy score being higher in Good than in Poor patients (illustrative full-cohort AUROC 0.72). The separation is genuinely multi-directional—no single eigendirection discriminates well (Fig. 9B)—consistent with the topology-driven classifier of [3] operating on a small but distributed set of trajectory-geometry features rather than any one mode. Under strict leakage control (nested (r, k) re-selection inside each outer fold, Δ rebuilt from training patients only, and a shuffled-feature null that collapses the score to chance) the standalone discriminant AUROC is 0.72, and the subspace adds a small but robust +0.009 AUROC over the full intrinsic-dimension battery and clinical variables (30-seed repeated cross-validation). The injured-brain outcome signal is therefore not only a matter of the scalar dimensionality *level* but also of the low-dimensional *orientation* of the EEG trajectory subspace—a complementary, geometrically interpretable contribution.

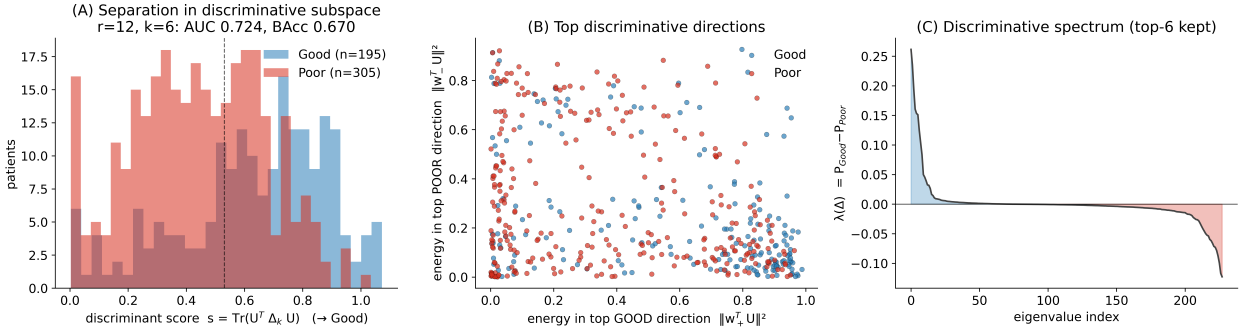

Figure 9: Bernadotte discriminative-subspace classifier of coma outcome (I-CARE, 500 patients, 195 Good / 305 Poor; first ≤ 24 h, $\sim$4 min of 19-channel EEG). **(A)** Distribution of the discriminant score $s = \mathrm{Tr}(U^\top \Delta_k U)$ for the two outcomes, using the leading $k = 6$ eigenvectors of $\Delta = P_{\mathrm{Good}} - P_{\mathrm{Poor}}$ at subspace rank $r = 12$: Good patients (blue) shift to higher s, Poor (red) to lower. **(B)** Each patient placed by the energy of its own subspace in the single most Good- and most Poor-characteristic direction; the classes separate only weakly on any one axis, so discrimination comes from the *combination* of directions. **(C)** Eigenvalue spectrum of Δ: the class difference is carried by a handful of extreme eigenvectors (blue $\lambda > 0$, Good-characteristic; red $\lambda < 0$, Poor-characteristic) over a large, near-zero shared bulk, so a low-dimensional discriminant subspace ($k \approx 3$–6 of 228 dimensions) suffices. The panel-A AUROC (0.72) is an in-sample, selected-(r, k) value; the leakage-free estimate (nested (r, k) selection, per-fold Δ, shuffled-feature null) is AUROC ≈ 0.72 standalone, and the subspace adds a verified +0.009 AUROC over the intrinsic-dimension battery + clinical fusion.

2.8 BOID contributes to state-of-the-art-level coma prognosis

The descriptor-level coma analysis showed that markers of failed downshifting and pathological overactivation carry partial and complementary outcome information.

We fused four feature families: BOID, the full intrinsic-dimension battery, spectral and burst-suppression descriptors, and routine clinical variables. Models were trained under outcome-stratified, patient-grouped five-fold cross-validation so that no patient contributed to both training and test folds; significance was assessed by patient-level label permutation.

Fusing feature families produced monotone gains (Fig. 10; Table 2). BOID alone reached AUC 0.68. The intrinsic-dimension battery alone reached 0.72. Their combination (BOID + intrinsic-dimension battery) reached 0.75, and adding spectral/burst descriptors reached 0.77 without clinical variables. The full model including clinical variables reached AUC 0.81, balanced accuracy 0.72, Brier score 0.176, and patient-permutation $p < 10^{-3}$. At a conservative false-positive rate of at most 5%, the fused model identified 33% of poor-outcome patients compared with 10% for the oscillator family alone. The result is best interpreted as an interpretable, state-of-the-art-level prognostic substrate for the failure of dimensional regulation, not as a stand-alone replacement for externally validated clinical EEG prognostication. The ablation shows that BOID makes a nontrivial contribution to the pipeline: it improves when combined with the broader intrinsic-dimension battery and again when fused with spectral, burst-suppression, and clinical variables.

Table 2: Leakage-free coma-outcome prognosis in I-CARE.

Feature set	AUC	Bal. acc.	Sens.@FPR≤5%	Brier
BOID	0.68	0.64	0.10	0.227
Spectral + burst-suppression	0.69	0.64	0.20	0.225
Clinical	0.72	0.68	0.17	0.214
Intrinsic-dimension battery	0.72	0.66	0.19	0.217
Intrinsic-dimension battery + BOID	0.75	0.68	0.23	0.203
Geometry + spectral	0.77	0.70	0.31	0.196
All + clinical	0.81	0.72	0.33	0.176

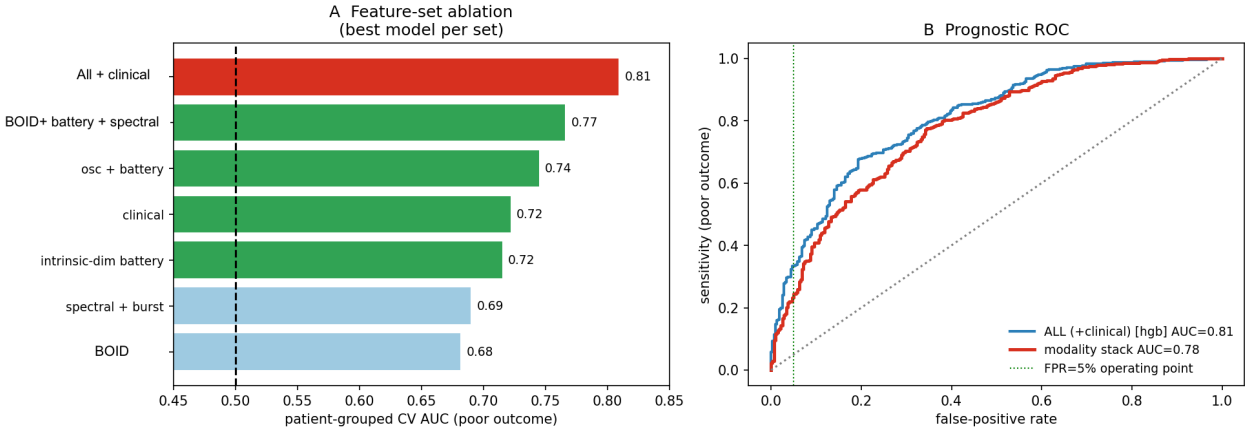

Figure 10: A fused dimensionality pipeline predicts coma outcome. Feature-set ablation shows monotone improvement from BOID descriptors to the full model with clinical variables. ROC curves show the fused model and stacked ensemble; the dotted line marks the conservative FPR ≤5% operating region.

The I-CARE database underlies the George B. Moody PhysioNet/Computing-in-Cardiology Challenge 2023 [23], whose official metric is exactly the sensitivity at a 5% false-positive rate for poor outcome at 72 h that we report, which lets us place our result against the field (Table 3). The leading challenge entries apply deep-network ensembles to many hours of full multichannel EEG together with ECG and clinical data, are tuned for this operating point, and are scored on a hidden test set; they reach 0.60–0.79 on the metric [23], and canonical deep-learning prognosticators report AUROC 0.88–0.90 [24]. Our pipeline is deliberately different: it reads only the first few minutes of resting EEG through a transparent, low-dimensional geometric feature set, yet attains AUROC 0.81—within the published range—at balanced accuracy 0.72. At the aggressive 5%-FPR operating point it reaches 0.33, below the deep-network leaders, reflecting both the minimal data it consumes and the absence of operating-point-specific tuning (these numbers are not strictly comparable: ours is patient-grouped cross-validation on the public training subset, theirs a hidden held-out test). The contribution here is therefore not a leaderboard-topping system but an interpretable, data-light marker: it shows that

the *intrinsic dimensionality* of the EEG carries genuine, independent prognostic signal, and it unifies the bedside coma prognosis with the consciousness read-out that the same geometry provides in the intact brain.

Method	I-CARE	n	Input	Evaluation	AUROC	Sens @5%	Key caveat
Reyna et al. 2023 [23] *(Challenge overview)*	yes	1020	19-ch EEG, ECG, clinical	hidden test; one hospital held out	—	0.79	overview; metric TPR@FPR≤5%, no AUROC
Zabihi et al. 2023 [25] *(hyperensemble: CatBoost + stacking)*	yes	607	EEG, ECG, clinical	hidden test + int. CV	0.92	0.79	reduced features under compute limits
Povinelli & Dupont 2024 [26] *(dynamical-systems features)*	yes	607	EEG, clinical	hidden test (1 run)	—	0.43	severe overfit (train 0.98 vs test 0.43)
Dionisio et al. 2024 [27] *(Transformer, attention over hourly EEG)*	yes + ext.	607	EEG only, PSD	single 80/20; ext. NTUH	0.82	—	patient-wise 0.58; per-hospital 0.33–0.85
Zheng et al. 2021 [28] *(recurrent Bi-LSTM)*	cons.†	1038	18-ch EEG, ∼58k h	5-fold CV (no hidden test)	0.78	—	consortium DB, not public set; CV only
Nesaragi et al. 2025 [29] *(Biaxial-former, dual-axis attention)*	yes	607	18-ch EEG	leave-one-hospital-out CV	0.77	0.52	self-made LOHO splits, not hidden test
This work (2026) *(subspace geometry + dimensionality battery)*	**yes**	**575**	**∼4 min, 19-ch EEG + clinical**	**leakage-free 5-fold patient-grouped CV**	**0.81**	**0.33**	**geometry adds verified +0.009; no epoch curation**
Sci. Rep. 2025 [30] *(gradient-boosted qEEG features)*	subset	277	qEEG, ECG, clinical	10-fold + 80/20, SMOTE	1.0	—	AUROC=1.0 implausible ⇒ likely leakage
Rohr et al. 2023 [31] *(Transformer with time prior)*	yes	607	EEG only, raw	CV only (no test score)	0.84‡	0.41	no comparable hidden-test result

Table 3: Reported prognostic performance of works using the I-CARE coma dataset, versus our result. Rows are ordered roughly by evaluation rigour (hidden-test Challenge entries, then external-validated, then internal cross-validation, then unscored). The official PhysioNet/CinC 2023 metric is the true-positive rate at false-positive rate ≤ 5% for poor outcome at 72 h on a sequestered hidden test set (one hospital held out); other rows report internal-CV or single-split AUROC, which is threshold-free, in-distribution and *not* comparable to the hidden-test sensitivity. Our pipeline reads only the first ∼4 min of 19-channel EEG with routine clinical variables, curates no epochs, and is scored by leakage-free patient-grouped cross-validation. †Zheng et al. used the I-CARE *consortium* database (the 7-hospital precursor to the frozen public challenge release), not the public set; the Explainable-ML row uses only a 277-patient EEG+ECG subset and reports an implausible AUROC of 1.0, so both must be read with care. ‡Value reported by Rohr et al. is AUPRC, not AUROC.

The reported numbers do not share a scale. Automated outcome prediction on the I-CARE corpus is an active benchmark, with reported areas under the ROC curve from 0.77 to above 0.9 and the official Challenge sensitivity (true-positive rate at a 5% false-positive rate, 72 h) reaching 0.79 in the best entries [23, 25]. These figures are easy to over-read, because they do not lie on a common scale. Only the hidden-test Challenge entries report the sensitivity at FPR$\leq$ 5% on a sequestered set with an entire hospital held out—the one number actually stress-tested for cross-site generalization; the Challenge winner reaches sensitivity 0.792 and AUROC 0.916 there [25]. Every other entry, including the transformer of Dionisio et al. (0.82) [27], the cross-hospital Biaxialformer (0.77) [32], and our own pipeline (0.81), reports an internal cross-validation or single-split AUROC—a threshold-free, in-distribution, intrinsically easier quantity. A CV-AUROC of 0.81 and a hidden-test sensitivity of 0.79 are not comparable; they measure different things on different data partitions.

The apparent ceiling thins under scrutiny. Once protocol and provenance are matched, the apparent $\sim$0.9 ceiling is less solid than it appears. The highest deep-learning AUROCs are typically obtained on manually curated artifact-free epochs from many hours of recording, which inflates in-sample performance and degrades off-distribution [27]. Recording-wise training can manufacture much of the gain: the transformer's 0.82 collapses to 0.58 under patient-wise training and to 0.65 on an external cohort, and ranges from 0.33 to 0.85 across hospitals [27]. At least one recent multimodal model reports an AUROC of 1.0 on a small subset—a value that signals leakage rather than skill. And the often-cited 0.78$\rightarrow$0.88 of Zheng et al. was obtained on the I-CARE *consortium* database, the multi-hospital precursor rather than the frozen public release, under cross-validation with no held-out test [28]. Several entries with strong internal cross-validation never obtained any official hidden-test score at all.

A confound common to all, including us. Beneath these methodological differences lies a confound that no internal validation removes: in these cohorts a poor outcome is frequently death following withdrawal of life-sustaining treatment, a decision that is itself informed by the EEG. Every EEG-based prognosticator therefore partly predicts the treating team's behaviour rather than neurobiology alone, which inflates all reported discriminations. The single number that is both stress-tested and honestly reported—the Challenge winner's hidden-test sensitivity of 0.79 at 5% FPR—is best read as the realistic ceiling for current EEG-based prognosis, and it is still far from the near-certainty a withdrawal-of-care decision demands.

Where our result stands. Against this backdrop we make no claim to the leaderboard: our AUROC of 0.81 sits mid-pack among internal-CV results, and our sensitivity of 0.33 at 5% FPR trails the hidden-test winners. Our contribution is of a different kind. It reads only the first few minutes of standard 19-channel EEG with routine clinical variables, curates no epochs, and is scored by leakage-free, patient-grouped cross-validation, so its 0.81 is a deliberately conservative figure rather than a tuned one. More importantly, it is mechanistically interpretable—the prognostic signal is the brain's intrinsic oscillator dimensionality—and it unifies the bedside coma read-out with the same geometry that indexes the level of consciousness in sleep, anesthesia and meditation. The optimized discriminative-subspace classifier adds a small but rigorously verified increment (+0.009

AUROC, surviving nested re-selection and a shuffled-feature control). In a literature where high accuracy and low interpretability travel together, and where headline numbers seldom survive a change of evaluation protocol, a transparent, data-light, reproducible marker—reported with its confounds stated—is the more durable contribution.

3 Discussion

We introduce brain oscillator intrinsic dimension (BOID) as a geometric marker of brain-state regime. BOID estimates the number of active exposinusoidal oscillatory degrees of freedom required to generate a brain time series. Across sleep, anesthesia, meditation, and post-anoxic coma, the same observable exposed a common principle: the healthy brain alternates between high-dimensional conscious engagement and low-dimensional recovery or efficiency modes, whereas poor-prognosis coma reflects failure of this dimensional regulation.

The first result is direct. Conscious wakefulness is high-dimensional. In human Sleep-EDF recordings, BOID decreased monotonically from wakefulness through N1 and N2 to N3 slow-wave sleep, and the same low-dimensional non-REM state was reproduced in mouse sleep. Propofol sedation produced a pharmacological version of the same transition, reducing BOID relative to wakefulness. These convergent findings indicate that suppression of consciousness in the intact brain is not merely a reduction in signal amplitude or a shift in spectral power. It is a coordinated contraction of the active oscillator repertoire.

This distinction is central. BOID is not alpha power, delta power, or any other frequency-band amplitude. Spectral analysis asks which rhythms are strong. BOID asks how many oscillatory degrees of freedom are active. A sedated brain may show strong organized alpha activity while remaining low-dimensional, because the signal is dominated by a smaller number of synchronized modes. Conversely, active wakefulness or mental effort may have less dominant alpha power but higher BOID because the brain recruits a broader repertoire of independent oscillatory generators. BOID therefore provides a complementary geometric layer: it measures the dimensionality of the latent oscillatory system, not the energy of a single rhythm.

The second result is equally important: low BOID is not pathological by itself. In healthy physiology, low BOID appears to be a recovery mode. Deep sleep, propofol sedation, and expert meditative stabilization all reduce the number of active oscillatory degrees of freedom, but they do so in organized, reversible, and potentially restorative ways. In these states, conscious access narrows, mental load decreases, and the brain enters a slower, more synchronized, lower-resource regime. Healthy consciousness is therefore not maximal dimensionality at all times. It is the capacity to switch: to enter high-dimensional engagement when interaction, cognition, and reportability are required, and to return to low-dimensional recovery when stabilization and restoration are needed.

Meditation provides a critical dissociation between low dimensionality and loss of function. Expert meditators showed lower intrinsic dimension than novices, suggesting that trained mental activity

can be maintained with fewer active dynamical degrees of freedom. In this interpretation, expertise reduces the dimensional cost of control. Expert meditation is therefore not a collapse of brain dynamics, but a low-BOID waking efficiency mode: a state in which the brain remains awake and controlled while using a smaller, more synchronized oscillator repertoire. This finding extends the sleep/anesthesia result by showing that low-dimensional organization can also occur during trained waking cognition.

Coma reveals the pathological boundary of the same principle. In post-anoxic coma, high BOID did not indicate preserved consciousness. Instead, poor neurological outcome was associated with elevated BOID and a pathological mode-proliferation index combining high mode count, high-frequency poles, burst intermittency, and mode-frequency entropy. This reverses the naive interpretation that high dimensionality is always beneficial. In the injured brain, high BOID may reflect hyperexcitability, epileptiform activity, myoclonus, burst-like fragmentation, unstable high-frequency poles, and metabolically costly overactivation. The poor prognostic sign is therefore not unconsciousness itself, but failure of the injured brain to downshift into an organized low-dimensional unconscious recovery state.

The resulting theory is not "more dimension is better." It is regulated dimensional switching. In the intact brain, high-dimensional dynamics support conscious interaction, cognitive flexibility, and mental effort. In deep sleep and anesthesia, low-dimensional dynamics support recovery. In expert meditation, low-dimensional dynamics support efficient waking control. In poor-prognosis coma, high-dimensional dynamics reflect pathological overactivation and failed downshifting. The biological meaning of dimensionality therefore depends on tissue integrity and state context. BOID becomes informative precisely because it separates these regimes within a single geometric framework.

The clinical implication is that coma prognosis may depend not only on whether the EEG is active, suppressed, or burst-suppressed, but on whether the injured brain can regulate its dimensionality. In the I-CARE cohort, BOID and PMPI carried significant outcome information, and BOID contributed nontrivially to a multimodal prognostic pipeline with intrinsic-dimension, spectral/burst, and clinical features. The full model reached leakage-free AUROC 0.81 under patient-grouped validation. This performance should not be overinterpreted as a replacement for established clinical prognostication or hidden-test challenge systems. Rather, the key contribution is mechanistic: a transparent, data-light geometric marker extracted from only the first minutes of standard EEG carries independent prognostic information and connects coma outcome to the same dimensional-regulation principle observed in sleep, anesthesia, and meditation.

This is especially important because automated coma prognosis is a difficult and confounded benchmark. Published systems differ in input duration, artifact curation, evaluation protocol, hospital splits, access to ECG and clinical variables, and operating-point tuning. Moreover, poor outcome after cardiac arrest is often partly endogenous: death may follow withdrawal of life-sustaining treatment, and withdrawal decisions are themselves influenced by the EEG. Any EEG-based prognosticator therefore risks partly predicting clinical behavior rather than neurobiology alone. For this reason,

our result should be read conservatively. We do not claim a leaderboard-topping model. We claim that intrinsic oscillator dimensionality is an interpretable, reproducible, and biologically meaningful signal that survives leakage-controlled validation and adds information beyond conventional feature families.

The framework also connects several previously separate approaches to brain dynamics. Neural manifold analysis emphasizes low-dimensional latent structure in population activity. Delay embedding reconstructs dynamical state spaces from time series. MSSA and ESPRIT provide a way to recover damped oscillatory modes from Hankel trajectories. Spectral analysis describes where power lies in frequency space. Persistent homology describes the cyclic organization of reconstructed trajectories. BOID links these layers by asking a simple question: how many oscillatory degrees of freedom are needed to generate the observed brain state? This makes BOID a bridge between dynamical-systems theory, signal decomposition, topology, and clinically interpretable EEG analysis.

Several limitations remain. BOID estimates depend on preprocessing, montage, sampling rate, window length, signal-to-noise ratio, embedding parameters, and the criteria used to retain oscillatory poles. Absolute BOID values should therefore not be compared across datasets without calibration; the strongest evidence in this study comes from within-dataset state contrasts and convergent replication across independent estimators. The sleep and anesthesia analyses use open datasets with specific montages and sample sizes. The meditation dataset is modest, and expertise effects may be influenced by arousal, movement, breathing, or task strategy. The coma model was internally validated on the public I-CARE training partition and requires prospective, multicenter, hospital-wise validation against established quantitative EEG markers and challenge operating metrics. Future work should also test whether BOID trajectories predict within-patient transitions, whether successful recovery from brain injury is preceded by a low-dimensional stabilization phase, and whether the same regime-switching principle generalizes to intracranial EEG, MEG, fMRI, and closed-loop neurointerface recordings.

In summary, BOID identifies a dimensional regulation principle of brain dynamics. Conscious engagement is high-dimensional. Healthy recovery and expert control are low-dimensional. Poor-prognosis coma is not simply low consciousness, but failed downshifting from pathological high-dimensional overactivation into restorative low-dimensional organization. Brain oscillator dimensionality therefore provides a single interpretable read-out linking consciousness, sleep, anesthesia, meditation, and coma.

4 Conclusion

We propose brain oscillator intrinsic dimension (BOID) as an interpretable geometric read-out of brain-state regime. BOID does not measure the power of a frequency band; it estimates the number of active oscillatory degrees of freedom required to generate the observed brain signal. This distinction allows the same framework to unify apparently different states: conscious wakefulness,

deep sleep, anesthesia, expert meditation, and post-anoxic coma.

Across intact brain states, the results support a regime-switching principle. Conscious engagement is high-dimensional and resource-demanding. Deep sleep and propofol anesthesia are low-dimensional recovery states. Expert meditation is a low-dimensional waking efficiency mode, suggesting that trained control can reduce the dynamical cost of conscious activity. Thus, healthy brain function is not maximal dimensionality, but regulated dimensional switching between high-dimensional engagement and low-dimensional stabilization.

Coma exposes the failure mode of this regulation. In poor-prognosis post-anoxic coma, elevated BOID and pathological mode proliferation do not indicate preserved consciousness; they indicate fragmented, hyperexcitable, metabolically costly dynamics and failure to enter a protective low-dimensional recovery state. Incorporating BOID into a multimodal prognostic pipeline shows that this dimensional failure carries clinically relevant information while remaining mechanistically interpretable.

Together, these findings identify oscillator dimensionality as a bridge between dynamical-systems theory, brain-signal geometry, and clinical neurophysiology. BOID provides a single interpretable variable for tracking consciousness, recovery, efficiency, and failed downshifting. The central principle is therefore not that higher dimensionality is better, but that healthy brains can regulate dimensionality—and injured brains may lose this capacity.

5 Materials and Methods

5.1 Brain time-series datasets

We analyzed open brain time-series datasets that together span the range of conscious states, from alert wakefulness through graded loss of consciousness to the injured brain after cardiac arrest (Table 4). These comprised whole-night human polysomnography with expert sleep-stage annotations (Sleep-EDF Expanded), polysomnographic mouse sleep providing a cross-species replication (OpenNeuro ds006366), propofol anesthesia with repeated behavioral awakenings that dissociate the presence and absence of experience (OpenNeuro ds005620), expert-versus-novice focused-attention meditation (OpenNeuro ds001787), and continuous EEG from comatose survivors of cardiac arrest dichotomized by neurological outcome (PhysioNet I-CARE). The goal was not to build a single diagnostic model, but to test whether intrinsic dimension behaves as a general geometric marker of the level of consciousness across physiologically and pathologically distinct routes to its loss.

5.1.1 I-CARE Dataset and cohort selection

We analyzed the International Cardiac Arrest Research consortium (I-CARE) database, released for the George B. Moody PhysioNet/Computing-in-Cardiology Challenge 2023 [34, 23]: continuous EEG with routine clinical variables from 1020 comatose survivors of cardiac arrest across seven hospitals

Table 4: Open brain time-series datasets spanning the range of conscious states, from intact wakefulness through reversible loss of consciousness to post-anoxic coma.

Dataset	Modality	Conditions
Sleep-EDF Expanded [33]	Human PSG/EEG	Wake / N1 / N2 / N3 / REM (n = 10 subjects)
OpenNeuro ds006366 [19]	Mouse PSG/EEG	Wake / NREM / REM (cross-species replication)
OpenNeuro ds005620 [20]	EEG	Awake vs. propofol-sedated (repeated awakenings)
OpenNeuro ds001787 [21]	EEG	Expert meditators (n = 12) Novice meditators (n = 12)
PhysioNet I-CARE [22]	Continuous EEG	Comatose post-cardiac-arrest; Good vs. Poor outcome (n = 575, 24–72 h)

in the USA and Europe, dichotomized at 3–6 months into Good (Cerebral Performance Category, CPC, 1–2) versus Poor (CPC 3–5) neurological outcome. Only the public training partition ($\sim$ 60% of patients) is released; the validation (10%) and test (30%) partitions are permanently sequestered, with one hospital held out entirely to the test set. All analyses were therefore restricted to the public training partition and evaluated by leakage-free, patient-grouped cross-validation; we did not access the hidden test set, so our internal cross-validated AUROC is not directly comparable to the Challenge's hidden-test operating-point metric.

Inclusion criteria. (i) adult comatose cardiac-arrest patient in the public I-CARE training release; (ii) a defined outcome label (Good or Poor); (iii) at least one EEG recording containing a finite, non-flat segment within the first 72 h after return of spontaneous circulation (24 h for the early-window subspace analysis). No manual selection of artifact-free epochs was performed.

Exclusion criteria. (i) patients present only in the sequestered validation/test partitions (inaccessible); (ii) records that could not be retrieved or decoded, or that contained no usable EEG segment in the target window; (iii) *for the subspace-geometry classifier only*, patients lacking the complete 19-channel 10–20 montage, because the method compares signal subspaces within a common ambient space—patients recorded with additional channels were retained (the canonical 19 forming a subset), whereas those missing any canonical channel were excluded.

Resulting cohorts. Applying these criteria to the retrieved records yielded 575 patients (217 Good / 358 Poor; 826 recordings at 24–72 h) for the oscillator-dimensionality battery and multimodal fusion, and 500 patients (195 Good / 305 Poor; one $\leq$ 24 h recording per patient) for the discriminative-subspace classifier. Both cohorts preserve the $\sim$ 60% majority of poor outcomes seen in the full dataset, and no patient contributed to more than one cross-validation fold.

5.2 Brain time series representation

Let

$$X(t) = (x_1(t), \ldots, x_m(t)) \in \mathbb{R}^m \tag{2}$$

be a multichannel brain time series, where m is the number of channels, sensors, electrodes, or regions of interest. For discrete recordings, we write

$$X_n = X(t_n), \qquad n = 1, \ldots, N. \tag{3}$$

Each recording is divided into windows

$$W_s = \{X_s, X_{s+1}, \ldots, X_{s+T-1}\}, \tag{4}$$

where T is the window length.

5.3 Hankel-delay embedding

For each multichannel signal $X(t) \in \mathbb{R}^m$, we constructed delay vectors

$$Z_t = [X(t), X(t-\tau), \ldots, X(t-(L-1)\tau)] \in \mathbb{R}^{mL}. \tag{5}$$

The set of all such vectors forms the reconstructed trajectory

$$\mathcal{T}_X = \{Z_t\}. \tag{6}$$

Equivalently, the embedded signal is represented by a block-Hankel matrix H_X, whose columns are delay vectors [7, 4, 3, 5, 6]. The intrinsic dimension of the brain time series is defined as the intrinsic dimension of this reconstructed trajectory [3, 5, 6].

5.4 Brain oscillator intrinsic dimension in the Bernadotte oscillator model

We estimate the intrinsic dimension of brain signals within the Bernadotte oscillator model. In this model, observed neural activity is represented as a finite superposition of latent exposinusoidal components introduced by Bernadotte [3, 5, 6]. Each component corresponds to an active neural mode with its own frequency, phase, damping or growth coefficient, and spatial projection. A brain time series is therefore not treated as an arbitrary stochastic process, but as a structured signal generated by a finite system of exposinusoidal oscillatory degrees of freedom.

A multichannel brain signal is modeled as

$$X(t) = \sum_{k=1}^{K} \mathbf{a}_k e^{\lambda_k t} \sin(\omega_k t + \varphi_k) + \varepsilon(t), \tag{7}$$

where $\mathbf{a}_k$ is the spatial pattern of the k-th mode, ω_k is its frequency, λ_k is a damping or growth parameter, φ_k is its phase, and $\varepsilon(t)$ denotes noise or unresolved activity.

The intrinsic dimension of the reconstructed brain trajectory is defined as the effective number of dynamically independent exposinusoidal components required to represent the signal [3, 5, 6]. This concept of brain-signal intrinsic dimension was introduced by Bernadotte and is operationalized here using Hankel-delay embeddings and algorithms for estimating the number of active components in brain-computer interface signals [4, 3]. The computational EEG/BCI implementation is further aligned with the lightweight Hankel-embedded pipeline by Menshikov, Elfimov, and Bernadotte [5, 6]. Thus, in the proposed framework, the intrinsic dimension of a brain signal is estimated as the number of active exposinusoidal degrees of freedom generating the observed neural dynamics.

5.5 Preprocessing and windowing

Across datasets we used a deliberately light, uniform preprocessing so that within-dataset contrasts reflect dynamics rather than analysis choices. Each recording was restricted to its EEG channels (the standard 10–20 montage where available; the single cortical EEG channel for the mouse data), band-pass filtered to 0.5–40 Hz, and resampled to 100 Hz (128 Hz for the mouse recordings, supplied at that rate). Recordings were segmented into non-overlapping epochs—4 s for the oscillator and frequency analyses, 30 s for the sleep singular-spectrum descriptors—and each epoch was treated as one window W and embedded independently. Flat or non-finite channels were dropped, and epochs with non-finite samples or near-zero variance were skipped. Every window-level descriptor was reduced to a per-recording or per-subject median before statistical comparison; only within-dataset gradients are interpreted, because absolute values depend on montage, window length and sampling rate.

5.6 MSSA front-end and ESPRIT oscillator extraction

The oscillator descriptors (the count and frequencies that carry most of the paper) are produced by a multichannel singular spectrum analysis (MSSA) front-end followed by ESPRIT pole extraction [3, 5, 6].

Each window is block-Hankel embedded with lag L and all C channels stacked, and the trajectory matrix H_X is decomposed by SVD; the signal-subspace rank r is selected by a 95% energy criterion on the singular values. ESPRIT exploits the lag shift-invariance of the block-Hankel row layout: dropping the last and the first lag-block of the leading r left singular vectors U yields two subspaces related by the diagonal of poles, $U_2 \approx U_1\Phi$ with $\mathrm{eig}(\Phi) = \{z_k\}$, recovered by total least squares. The complex poles $z_k = \exp((\sigma_k + \mathrm{i}\omega_k)\,\Delta t)$ give each mode's frequency $f_k = \omega_k/2\pi$, damping σ_k, and per-channel complex amplitude A_{ck}.

The BOID (*oscillator count*) n_{modes} is the number of recovered poles with positive frequency in 1–40 Hz; under conjugate pairing the implied number of real oscillatory modes is $\widehat{r} = n_{\mathrm{modes}}/2$.

From the modes we form the oscillator-family descriptors used throughout: the power-weighted ($|A_k|^2$) median and mean oscillator frequency; the fraction of oscillatory power in the delta, theta, alpha and beta bands and above 20 Hz; and the Shannon entropies of the mode-frequency and mode-power distributions. As a surrogate-thresholded alternative we also compute an oscillatory dimension d_{osc}, the size of the leading block of singular values of H_X that exceed a null built from independent per-channel temporal shuffles of the window; a temporal-shuffle null is required because phase-randomization preserves the autocovariance and hence the entire Hankel spectrum.

5.7 Hankel spectral rank

For each Hankel matrix H_X, we computed the singular value decomposition

$$H_X = U\Sigma V^\top. \tag{8}$$

From the singular values σ_i, we computed effective Hankel rank, energy rank, participation ratio, entropy rank, and stable rank:

$$d_{\mathrm{PR}} = \frac{\left(\sum_i \sigma_i^2\right)^2}{\sum_i \sigma_i^4}, \tag{9}$$

$$d_{\mathrm{ent}} = \exp\left(-\sum_i p_i \log p_i\right), \qquad p_i = \frac{\sigma_i}{\sum_j \sigma_j}. \tag{10}$$

These quantities provide linear estimates of the effective number of active Hankel components [4, 3, 5, 6].

5.8 Additional intrinsic dimension estimators

To avoid relying on a single estimator, we also computed nearest-neighbor and graph-based intrinsic dimension estimates, including the Levina–Bickel maximum likelihood estimator, TwoNN estimator, minimum-spanning-tree scaling estimators, and the Brito–Quiroz–Yukich MST-degree estimator [15, 16, 17, 18].

For each embedded point Z_i, let $T_j(Z_i)$ be the Euclidean distance to its j-th nearest neighbor. The local Levina–Bickel estimate for neighborhood size k is

$$\hat{d}_k(Z_i) = \left[\frac{1}{k-1}\sum_{j=1}^{k-1} \log \frac{T_k(Z_i)}{T_j(Z_i)}\right]^{-1}. \tag{11}$$

The global estimate is obtained by averaging over embedded points and, where applicable, over a range of neighborhood sizes.

Agreement among estimators was interpreted as evidence for stable manifold-like structure. Disagreement was treated as a potential marker of nonstationarity, fragmentation, or departure from

smooth low-dimensional dynamics.

5.9 Fractal estimators

As a check independent of the singular spectrum we additionally computed three established estimators. The Grassberger–Procaccia correlation dimension D_2 is the scaling exponent of the correlation integral $C(\epsilon) = \frac{2}{K(K-1)} \sum_{i<j} \mathbf{1}[\|Z_i - Z_j\| < \epsilon] \sim \epsilon^{D_2}$, estimated as the slope of $\log C$ versus $\log \epsilon$ over the lower-middle of the pairwise-distance distribution [35]. The false-nearest-neighbor embedding dimension is the smallest delay-embedding dimension at which the fraction of neighbors that separate on extending the embedding falls below a tolerance, computed per channel and averaged. The Higuchi fractal dimension is the slope of $\log L(k)$ versus $\log(1/k)$ for the curve length $L(k)$ of the series sub-sampled at lag k, again per channel and averaged. Because the correlation dimension acts on the multichannel delay embedding while the false-nearest-neighbor and Higuchi estimators act on the raw single-channel signals, their agreement with the spectral and graph estimators provides convergent evidence independent of the Hankel construction. The estimator-dispersion index Δ_{ID} is the inter-quartile range across the scalar estimators and is itself reported as a state descriptor.

5.10 Topological descriptors based on persistent homology

We computed persistent homology of reconstructed Hankel-delay trajectories using Vietoris–Rips filtrations. For each trajectory, we extracted H_0 and H_1 persistence diagrams, persistence entropy, the number of long-lived features, and vectorized topological summaries.

For a persistence diagram

$$D_q = \{(b_j, d_j)\}_j, \tag{12}$$

we define lifetimes

$$\ell_j = d_j - b_j \tag{13}$$

and persistence entropy

$$\mathrm{PersEntropy}(D_q) = -\sum_j p_j \log p_j, \qquad p_j = \frac{\ell_j}{\sum_k \ell_k}. \tag{14}$$

The topological interpretation is direct: dominant oscillations produce loop-like structures; multiple oscillatory modes produce more complex cyclic structures; fragmentation or switching produces altered connected components and persistence patterns.

5.11 Bernadotte discriminative-subspace classifier: topological separation of time series by class projectors

Each multichannel recording X was represented by the r-dimensional signal subspace $U_r(X)$ of its block-Hankel trajectory. For each outcome class, we computed the extrinsic Grassmann mean of the training members' subspace projectors,

$$P_c = \left\langle U_i U_i^\top \right\rangle_{i \in c}. \tag{15}$$

The class-separation operator

$$\Delta = P_{\text{Good}} - P_{\text{Poor}} \tag{16}$$

is symmetric, with eigenvalues in $[-1, 1]$. Empirically, most of its spectrum formed a near-zero shared bulk, whereas a small number of extreme eigenvectors captured directions that separated the two classes. We therefore retained the k eigenvectors with largest absolute eigenvalues, defining the discriminative subspace W_k, and discarded the shared bulk.

A query recording was scored by the signed energy of its own Hankel subspace within these discriminative directions:

$$s(X) = \sum_{j=1}^{k} \lambda_j \left\| w_j^\top U_r(X) \right\|_F^2 = \operatorname{Tr}\left(U_r(X)^\top \Delta_k U_r(X) \right), \tag{17}$$

where $\Delta_k = \sum_{j=1}^{k} \lambda_j w_j w_j^\top$. The sign and magnitude of $s(X)$ quantify whether the recording projects more strongly onto Good- or Poor-associated characteristic directions. Taking k to full rank recovers the plain subspace-distance classifier of [3, 5, 6]; finite $k \ll mL$ restricts the classifier to the discriminative directions and removes the large shared component of the two classes. (Algorithm 1).

Parameters and evaluation

The embedding window was fixed at $L = 12$ throughout. The signal-subspace rank r and discriminant dimension k were the only tuned hyperparameters. The classifier was evaluated on 500 comatose patients (195 Good / 305 Poor; one $\leq$ 24-h recording per patient) under leakage-free, patient-grouped cross-validation. To exclude model-selection and out-of-fold leakage, Δ was rebuilt using training patients only, and (r, k) were re-selected by inner cross-validation within each outer fold. A label-permuted shuffled-score control collapsed performance to chance.

On its own, the discriminative-subspace classifier reached AUROC 0.72 with best performance at $r = 12$ and $k \approx 3$–6. When added to the intrinsic-dimension battery and clinical-feature fusion, it contributed a verified +0.009 AUROC improvement in 30-seed repeated cross-validation.

Algorithm 1 Bernadotte discriminative-subspace classifier (topological separation of time series by class projectors). Specialises the subspace classifier of [3] to the low-dimensional subspace in which the classes separate.

Require: labelled multichannel series $\{(X_i, y_i)\}_{i=1}^{n}$, $y_i \in \{A, B\}$; embedding lag L; subspace rank r; discriminant dimension k

Signal subspace of one recording

function SUBSPACE(X, r) ▷ $X \in \mathbb{R}^{C\times N}$: C channels, N samples

  $H \leftarrow$ block-Hankel embedding of X, lag L ▷ $H \in \mathbb{R}^{CL\times(N-L+1)}$, channels stacked

  $\Sigma \leftarrow \frac{1}{N-L+1} HH^{\top}$ ▷ lagged covariance, $CL \times CL$

  return $U_r(\Sigma)$ ▷ top-r eigenvectors of Σ

end function

Train: build the discriminant subspace

for $i = 1, \ldots, n$ **do** $U_i \leftarrow$ SUBSPACE(X_i, r)

end for

for $c \in \{A, B\}$ **do** $P_c \leftarrow \dfrac{1}{|I_c|} \displaystyle\sum_{i\in I_c} U_i U_i^{\top}$ ▷ extrinsic Grassmann mean (mean projector) of class c

end for

$\Delta \leftarrow P_A - P_B$ ▷ symmetric, eig $\in [-1, 1]$

$(\lambda_j, w_j) \leftarrow$ eig(Δ); keep the k pairs of largest $|\lambda_j|$ ▷ $W_k = [w_1 \ldots w_k]$, $\Delta_k = \sum_{j\le k} \lambda_j w_j w_j^{\top}$

Classify a query $X_\star$

$U_\star \leftarrow$ SUBSPACE($X_\star, r$)

$s(X_\star) \leftarrow \displaystyle\sum_{j=1}^{k} \lambda_j \, \|w_j^{\top} U_\star\|_F^2 = \mathrm{Tr}\big(U_\star^{\top} \Delta_k U_\star\big)$

return A if $s(X_\star) > \tau$ else B ▷ $\lambda_j > 0$: A-characteristic; $\lambda_j < 0$: B-characteristic

5.12 Topology-driven classifier

For each window W_s, we constructed a feature vector

$$\Psi(W_s) = [\Psi_{\mathrm{ID}}, \Psi_{\mathrm{Hankel}}, \Psi_{\mathrm{PH}}, \Psi_{\mathrm{geom}}], \tag{18}$$

where Ψ_{ID} collects the nearest-neighbor, graph and fractal intrinsic-dimension estimates, Ψ_{Hankel} the singular-spectrum descriptors, Ψ_{PH} the persistent-homology summaries, and Ψ_{geom} additional trajectory-geometry features (no raw band power or amplitudes enter Ψ). The classifier F is a histogram gradient-boosting model with balanced class weights, evaluated under K-fold cross-validation *grouped by subject* (`GroupKFold`), so that no subject contributes windows to both the training and the test fold. We report balanced accuracy, macro-F_1, and (for binary contrasts) the ROC-AUC, with chance level 1/(number of classes), together with permutation feature importance. For the I-CARE coma data the outcome label is per patient, and all 575 patients with a usable 24 and/or 72 h recording were analyzed (the first 240 s of each, partially streamed). We tested whether

higher oscillator dimensionality marks poor outcome in two ways: (i) for each marker—the oscillator count, the mean oscillator frequency, and a pre-specified pathological mode-proliferation index (PMPI = z(n_modes) + z(hf-pole fraction) + z(burst intermittency) + z(mode-frequency entropy), z-scored within each timepoint) (we report the discrimination AUC for predicting Poor and a one-sided Mann–Whitney test; and (ii) a leakage-free, patient-grouped cross-validated logistic classifier (standardized features, `GroupKFold` so no patient is split across folds) over the oscillator count, PMPI and frequency descriptors, whose significance was assessed by a patient-level label-permutation test (3000 permutations reassigning Good/Poor across patients, two-sided on $|\mathrm{AUC} - 0.5|$). For the fused prognostic pipeline we grouped the per-recording features into four families—oscillator, the full intrinsic-dimension battery (Levina–Bickel MLE, TwoNN, correlation dimension, false-nearest-neighbors, MST-length, BQY-degree, persistent-homology H_1, effective/stable rank, Higuchi, estimator dispersion), spectral/burst-suppression, and clinical—and trained, for each family and each union, both a histogram gradient-boosting classifier and an ℓ_2-regularized logistic model under *outcome-stratified* patient-grouped five-fold cross-validation (`StratifiedGroupKFold`); we report the better model per feature set, scored by ROC-AUC, balanced accuracy, sensitivity at a false-positive rate $\leq 5\%$, and the Brier score, with a 2000-fold patient-level permutation test on the winner and on a three-modality out-of-fold stacked ensemble.

Statistical analysis and classification

All descriptor comparisons were performed on subject-level (or patient-level) medians, to avoid inflated significance from many correlated windows. Differences across more than two physiological states used the Kruskal–Wallis test across groups (the Friedman test for the within-subject sleep-stage design); within-subject two-condition contrasts (rest vs. task, awake vs. sedated) used the paired Wilcoxon signed-rank test; two-group contrasts (Good vs. Poor outcome, expert vs. novice) used the Mann–Whitney U test; and ordinal associations (with the CPC score) used the Spearman correlation. P-values are reported per descriptor; where a family of descriptors was screened this is stated, and only the strongest effects survive an informal multiplicity correction. All cross-validation is grouped by subject or patient (no leakage), and significance of cross-validated decoding was, where stated, assessed by label permutation respecting the grouping.

Algorithmic summary

Algorithm 2 Intrinsic dimension and topology-based analysis of brain time series

Require: Multichannel brain time series $X(t)$, labels y, embedding length L, delay τ
Ensure: Intrinsic dimension estimates, topological descriptors, brain-state classifier

```
for each subject do
    Preprocess and normalize multichannel recording
    Segment recording into windows W_s
    for each window W_s do
        Construct Hankel-delay trajectory T_X
        Compute SVD of the Hankel matrix H_X
        Estimate d_rank, d_PR, d_ent
        Estimate nearest-neighbor and graph-based intrinsic dimension
        Compute persistent homology descriptors D_0, D_1
        Build feature vector Ψ(W_s)
    end for
end for
Aggregate descriptors at subject and condition levels
Test group differences in intrinsic dimension and topology
Train topology-driven classifier with subject-wise cross-validation
Report statistical effects, estimator agreement, and classification performance
```


Data and materials availability

All datasets analyzed in this study are publicly available through PhysioNet or OpenNeuro.

Code for Hankel-delay embedding, oscillator extraction, intrinsic-dimension estimation, MSSA, and topological analysis will be made available in a public repository upon publication: `https://data.aicumene.com`.

Funding information

The research was funded by the authors.

Competing interests

The authors declare no competing interests.

Supplement

Within-subject reproducibility of sleep-related dimensional downshifting.

To verify that the sleep-depth effect was not driven by a single subject or by pooled-window statistics, we visualized whole-night BOID trajectories for each Sleep-EDF recording in the analyzed sample. In nearly all recordings, BOID was highest during wakefulness and lowest during N3 slow-wave sleep, with N1, N2, and REM occupying intermediate ranges. The running median followed the hypnogram over the night: dimensionality decreased during transitions into deeper NREM sleep and increased again during REM or wake transitions. This subject-level montage confirms that the reduction of BOID with sleep depth is reproducible across individuals and reflects a within-brain state transition from high-dimensional wakefulness to low-dimensional sleep recovery mode.(Fig. S11).

Individual-animal reproducibility of mouse sleep dimensional downshifting.

To verify that the mouse sleep effect was not driven by pooled epochs or a small number of animals, we visualized BOID trajectories separately for individual mice. Across recordings, non-REM sleep was associated with lower BOID than wakefulness, while REM generally occupied an intermediate or higher-dimensional regime. The within-animal boxplots reproduced the group-level ordering, and the time-resolved traces showed BOID decreasing during sustained non-REM periods and increasing during transitions back toward wakefulness or REM. These individual-animal trajectories support the conclusion that low-dimensional non-REM sleep is a conserved recovery-mode state across mammalian sleep.(Fig. S12).

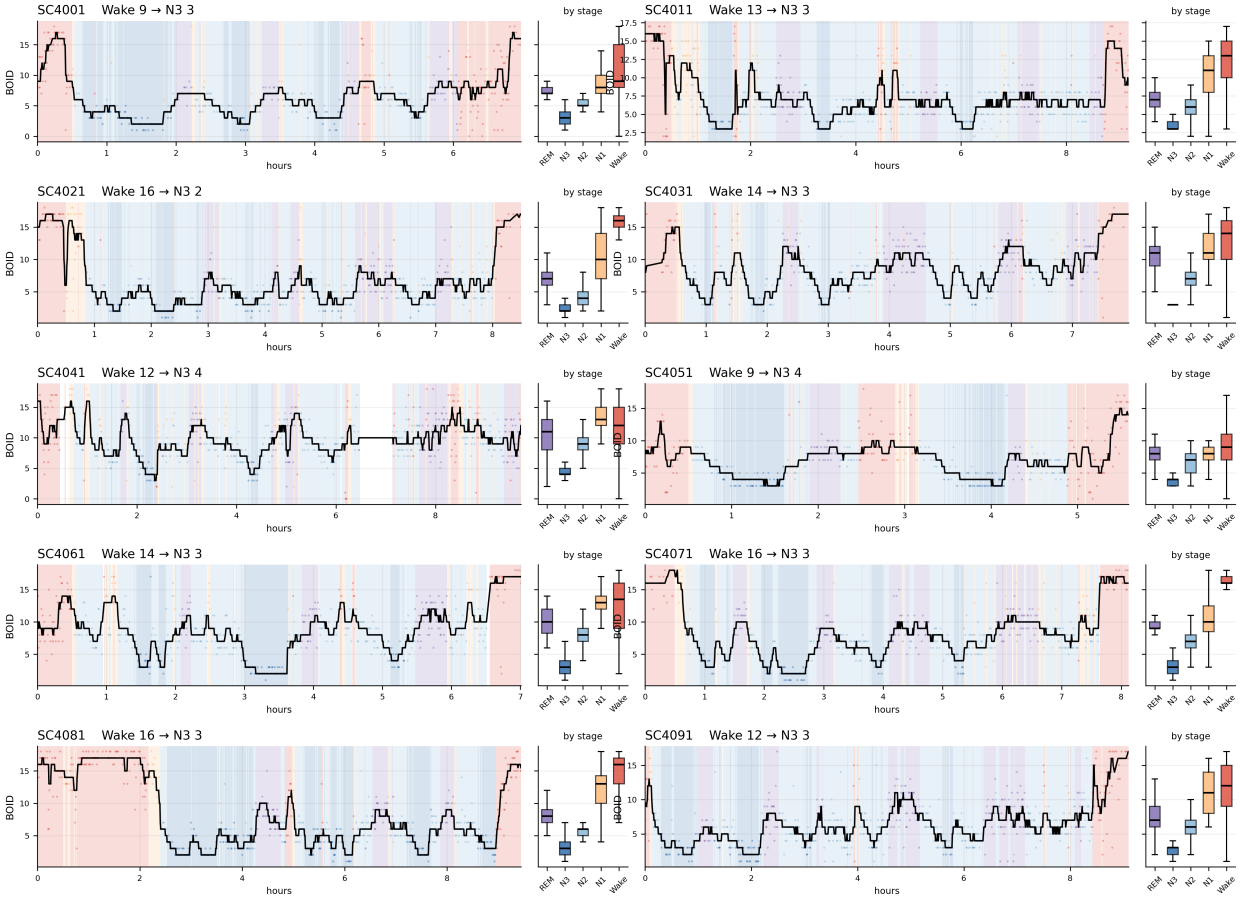

Figure 11: Whole-night oscillator dimensionality trajectories across the Sleep-EDF human sleep sample. Each panel shows one Sleep-EDF recording from the analyzed human sleep subset. The main trace in each panel is Brain oscillator intrinsic dimension (BOID), computed as oscillator count per 30-s epoch, across the night. Colored points indicate individual epochs; the black curve shows a running median. Background shading denotes manually scored sleep phase: Wake in red, N1 in orange, N2 in light blue, N3 in dark blue, and REM in purple. The small boxplot to the right of each trajectory summarizes BOID by sleep stage within the same recording. Across subjects, BOID is consistently highest during wakefulness, decreases during NREM sleep, and reaches its lowest values during N3 slow-wave sleep, with REM generally occupying an intermediate regime.

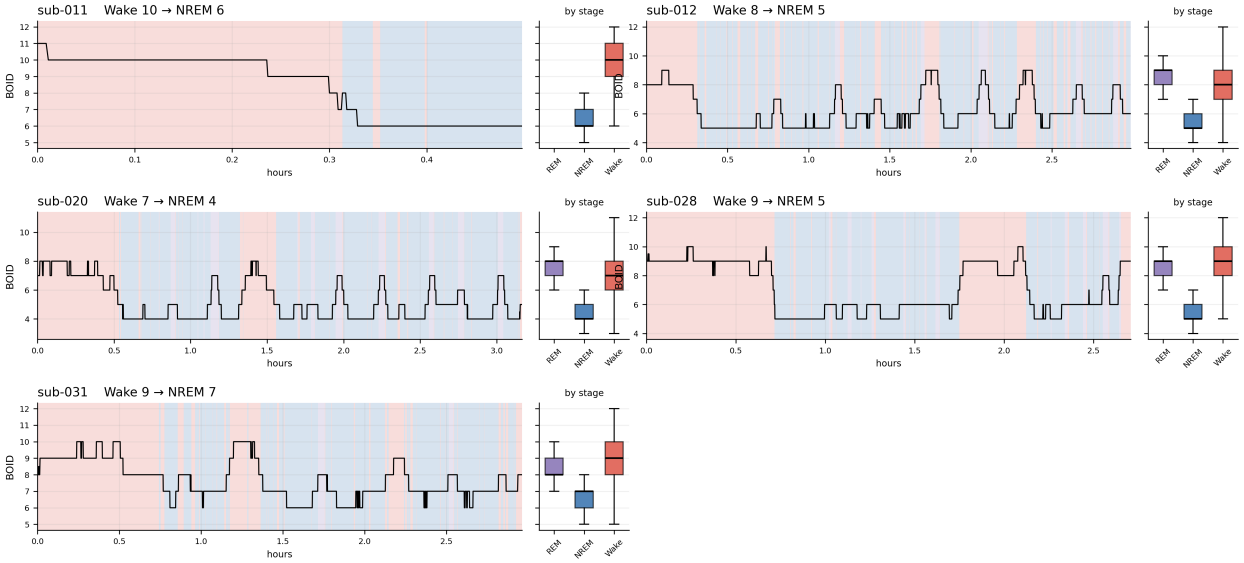

Figure 12: Mouse sleep reproduces the low-dimensional non-REM state across individual animals. Whole-recording BOID trajectories for representative mice from OpenNeuro ds006366. Each panel shows brain oscillator intrinsic dimension (BOID), computed as oscillator count across time, together with sleep-stage background shading. Wake is shown in red, non-REM sleep in dark blue, and REM in purple. The black step curve shows the BOID trajectory for each animal, and the adjacent boxplot summarizes BOID by sleep stage within the same recording. Across animals, BOID is consistently lower during non-REM sleep than during wakefulness, with REM generally occupying an intermediate or higher-dimensional regime. Panel titles report the within-animal Wake and non-REM medians, illustrating the repeated dimensional downshift from wakefulness into non-REM sleep. This montage shows that the mouse sleep effect is reproducible at the individual-animal level and is not driven by pooling epochs across animals.

References

[1] E. N. Brown, R. Lydic, and N. D. Schiff, General anesthesia, sleep, and coma. *New England Journal of Medicine* **363**, 2638–2650 (2010). doi:10.1056/NEJMra0808281.

[2] E. Adam et al., Modulatory dynamics mark the transition between anesthetic states of unconsciousness. *Proceedings of the National Academy of Sciences* **120**, e2300058120 (2023). doi:10.1073/pnas.2300058120.

[3] A. Bernadotte, Topology-driven classification of time series. *bioRxiv* (2026). doi:10.64898/2026.04.25.720787.

[4] A. Bernadotte, V. Buchstaber, Method for evaluating the number of signal sources and application to non-invasive brain-computer interface. arXiv:2410.11844 (2024). doi:10.48550/arXiv.2410.11844.

[5] I. Menshikov, N. Elfimov, A. Bernadotte, A lightweight Hankel-embedded pipeline for real-time EEG filtering and classification. *Proc. ICARA*, 589–593 (2026). doi:10.1109/ICARA69401.2026.11480292.

[6] A. Bernadotte, Estimating the number of sources in EEG with Hankel embedding for brain-computer interface. *Proc. ICARA*, 621–626 (2026). doi:10.1109/ICARA69401.2026.11480398.

[7] V. M. Buchstaber, Time series analysis and Grassmannians. *Amer. Math. Soc. Transl.* **162**, 1–17 (1994).

[8] R. Roy, T. Kailath, ESPRIT: Estimation of signal parameters via rotational invariance techniques. *IEEE Trans. Acoust. Speech Signal Process.* **37**, 984–995 (1989). doi:10.1109/29.32276.

[9] Y. Kuramoto, *Chemical Oscillations, Waves, and Turbulence* (Springer, 1984).

[10] J. A. Acebrón, L. L. Bonilla, C. J. Pérez Vicente, F. Ritort, and R. Spigler, The Kuramoto model: A simple paradigm for synchronization phenomena. *Reviews of Modern Physics* **77**, 137–185 (2005). doi:10.1103/RevModPhys.77.137.

[11] E. M. Izhikevich, Simple model of spiking neurons. *IEEE Transactions on Neural Networks* **14**, 1569–1572 (2003). doi:10.1109/TNN.2003.820440.

[12] E. M. Izhikevich,

Dynamical Systems in Neuroscience: The Geometry of Excitability and Bursting (MIT Press, 2007).

[13] J. Rissanen, Modeling by shortest data description. *Automatica* **14**, 465–471 (1978). doi:10.1016/0005-1098(78)90005-5.

[14] P. D. Grünwald, *The Minimum Description Length Principle* (MIT Press, 2007).

[15] E. Levina and P. J. Bickel, Maximum likelihood estimation of intrinsic dimension. In *Advances in Neural Information Processing Systems* **17**, 777–784 (2005).

[16] E. Facco, M. d'Errico, A. Rodriguez, and A. Laio, Estimating the intrinsic dimension of datasets by a minimal neighborhood information. *Scientific Reports* **7**, 12140 (2017). doi:10.1038/s41598-017-11873-y.

[17] M. R. Brito, A. J. Quiroz, and J. E. Yukich, Intrinsic dimension identification via graph-theoretic methods. *Journal of Multivariate Analysis* **116**, 263–277 (2013). doi:10.1016/j.jmva.2012.12.007.

[18] F. Camastra and A. Staiano, Intrinsic dimension estimation: Advances and open problems. *Information Sciences* **328**, 26–41 (2016). doi:10.1016/j.ins.2015.08.029.

[19] L. Rose et al., Mouse Sleep Staging Validation dataset and automated sleep staging in mice. *SLEEP Advances* **6**, zpaf025 (2025). doi:10.1093/sleepadvances/zpaf025.

[20] I. J. Bajwa et al., A repeated awakening study exploring the capacity of complexity measures to capture dreaming during propofol sedation. *Sci. Rep.* **15**, s41598-025-12695-z (2025). doi:10.1038/s41598-025-12695-z.

[21] A. Delorme and T. Brandmeyer, EEG meditation study. *OpenNeuro*, dataset ds001787. doi:10.18112/openneuro.ds001787.v1.0.2.

[22] E. Amorim et al., The International Cardiac Arrest Research (I-CARE) Consortium electroencephalography database. *Crit. Care Med.* (2023). Dataset: PhysioNet, doi:10.13026/m33r-bj81.

[23] M. A. Reyna, E. Amorim, R. Sameni, J. Weigle, A. Elola, A. Bahrami Rad, S. Seyedi, H. Kwon, W.-L. Zheng, M. M. Ghassemi, M. J. A. M. van Putten, J. Hofmeijer, N. Gaspard, A. Sivaraju, S. T. Herman, J. W. Lee, M. B. Westover, and G. D. Clifford, "Predicting neurological recovery from coma after cardiac arrest: the George B. Moody PhysioNet Challenge 2023," *Computing in Cardiology*, vol. 50, 2023. doi:10.22489/CinC.2023.205.

[24] M. C. Tjepkema-Cloostermans, C. da Silva Lourenço, B. J. Ruijter, S. C. Tromp, G. Drost, F. H. M. Kornips, A. Beishuizen, F. H. Bosch, J. Hofmeijer, and M. J. A. M. van Putten, "Outcome prediction in postanoxic coma with deep learning," *Critical Care Medicine* **47**, 1424–1432 (2019). doi:10.1097/CCM.0000000000003854.

[25] M. Zabihi, A. Chaman Zar, P. Grover, and E. S. Rosenthal, "Hyperensemble learning from multimodal biosignals to robustly predict functional outcome after cardiac arrest," *Computing in Cardiology*, vol. 50, 2023. doi:10.22489/CinC.2023.142.

[26] R. J. Povinelli and R. Dupont, "A dynamical systems approach to predicting patient outcome after cardiac arrest," *arXiv preprint* arXiv:2405.08827, 2024.

[27] J. Dionisio, M. Rizwan, H. Suzuki, K. M. Hassan, I. Chanpornpakdi, L.-Y. Lin, T. Tanaka, and C. Lin, Is attention all you need for EEG to predict neurological outcomes in cardiac arrest? In *Proc. IEEE EMBS International Conference on Biomedical and Health Informatics* (2024). doi:10.1109/BHI62660.2024.10913569.

[28] W.-L. Zheng, E. Amorim, J. Jing, O. Wu, M. Ghassemi, et al., "Predicting neurological outcome from electroencephalogram dynamics in comatose patients after cardiac arrest with deep learning," *IEEE Trans. Biomed. Eng.*, vol. 69, no. 5, pp. 1813–1825, 2021. doi:10.1109/TBME.2021.3139007.

[29] N. Nesaragi, A. Qadir, P. Halvorsen, and I. Balasingham, "Biaxialformer: leveraging channel independence and inter-channel correlations in EEG signal decoding for predicting neurological outcomes," *arXiv preprint* arXiv:2507.02879, 2025.

[30] Y. Niu, X. Chen, J. Fan, C. Liu, M. Fang, Z. Liu, X. Meng, Y. Liu, L. Lu, and H. Fan, "Explainable machine learning model based on EEG, ECG, and clinical features for predicting neurological outcomes in cardiac arrest patients," *Scientific Reports* **15**, 11498 (2025). doi:10.1038/s41598-025-93579-0.

[31] M. Rohr, T. Schilke, L. Willems, C. Reich, S. Dill, G. Güney, and C. Hoog Antink, "Transformer network with time prior for predicting clinical outcome from EEG of cardiac arrest patients," *Computing in Cardiology*, vol. 50, 2023. doi:10.22489/CinC.2023.173.

[32] S. Nesaragi, H. A. Qadir, P. Halvorsen, and I. Balasingham, "Biaxialformer: leveraging channel independence and inter-channel correlations in EEG signal decoding for predicting neurological outcomes," *arXiv preprint* arXiv:2507.02879, 2025.

[33] B. Kemp et al., Analysis of a sleep-dependent neuronal feedback loop: the slow-wave microcontinuity of the EEG. *IEEE Trans. Biomed. Eng.* **47**, 1185–1194 (2000). doi:10.1109/10.867928.

[34] E. Amorim, W.-L. Zheng, J. W. Lee, S. T. Herman, M. M. Ghassemi, A. Sivaraju, N. Gaspard, J. Hofmeijer, M. J. A. M. van Putten, M. A. Reyna, G. D. Clifford, and M. B. Westover, "I-CARE: International Cardiac Arrest REsearch consortium database (version 2.1)," *PhysioNet*, 2023. doi:10.13026/m33r-bj81.

[35] P. Grassberger, I. Procaccia, Characterization of strange attractors. *Phys. Rev. Lett.* **50**, 346–349 (1983). doi:10.1103/PhysRevLett.50.346.