Supplementary analyses of the post-anoxic coma cohort

Bernadotte lab

Supplementary analyses of the post-anoxic coma cohort

To support the descriptor- and model-level coma results in the main text, we provide additional views of the same PhysioNet I-CARE analysis: the per-window feature construction, the burst-suppression substrate of the mode count, cohort- and patient-level dimensional trajectories, the marginal contribution of each feature block, the clinical operating points, and the geometric/topological descriptors. All are computed under the same patient-grouped cross-validation and carry no new claims beyond the main text. Time is hours since ROSC.

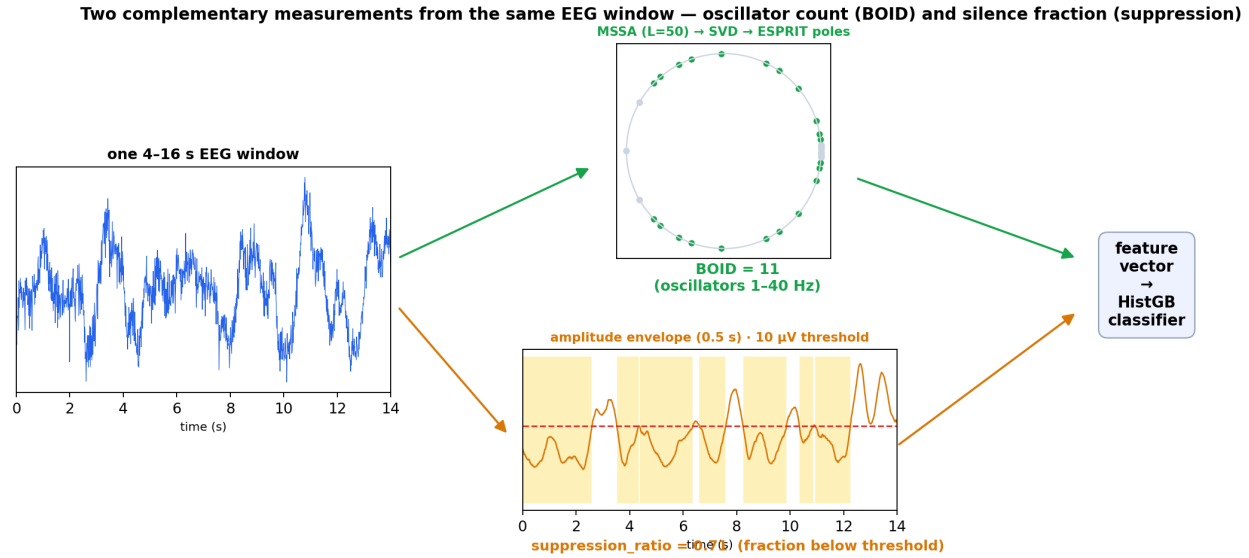

Figure 1: **Two complementary readouts from one EEG window.** From each 4-s window the MSSA/ESPRIT front-end returns the oscillator count (BOID; in-band poles) while an amplitude-envelope branch returns the suppression ratio (fraction below a 10 µV threshold). Oscillatory complexity and electrical-silence fraction are near-independent and both enter the feature vector.

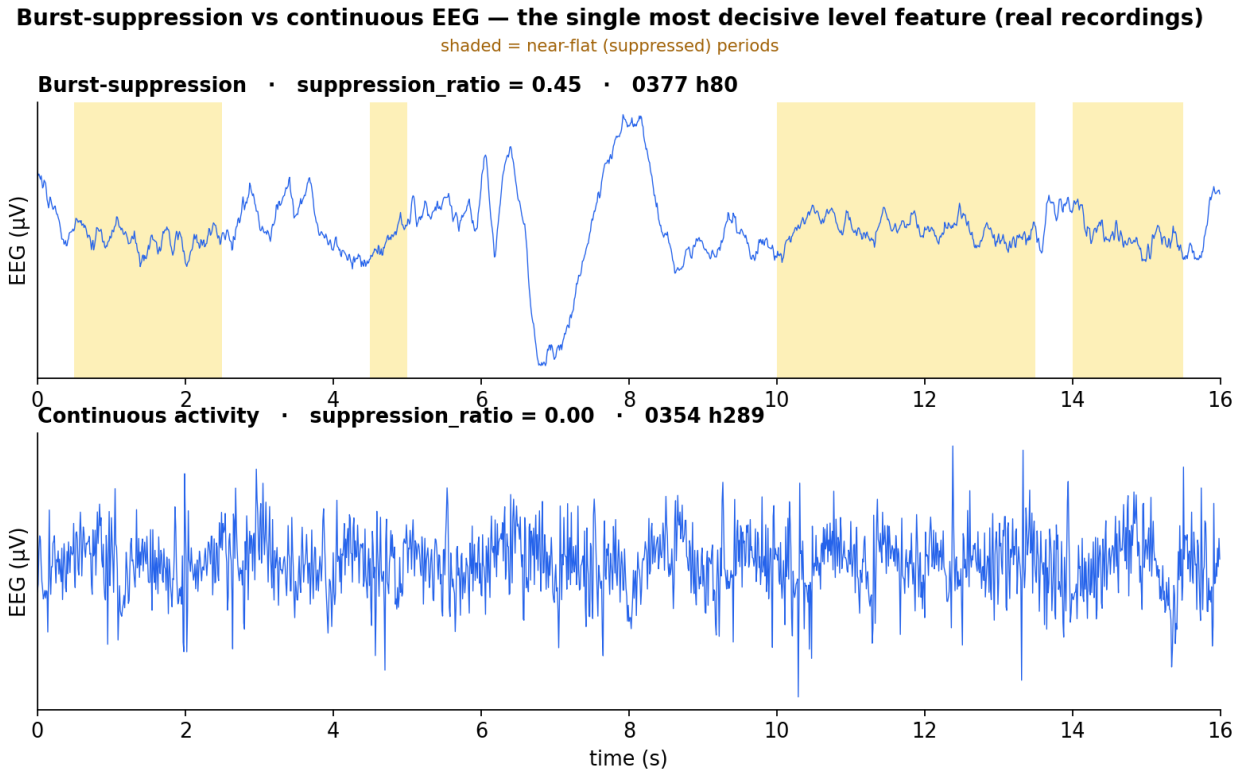

Figure 2: **Burst-suppression versus continuous EEG (real I-CARE recordings).** A burst-suppression recording (suppression ratio 0.45; suppressed stretches shaded) beside a continuous recording (0.00). The transient high-frequency bursts inflate the detected mode count without preserved consciousness, consistent with the pathological-proliferation interpretation in the main text.

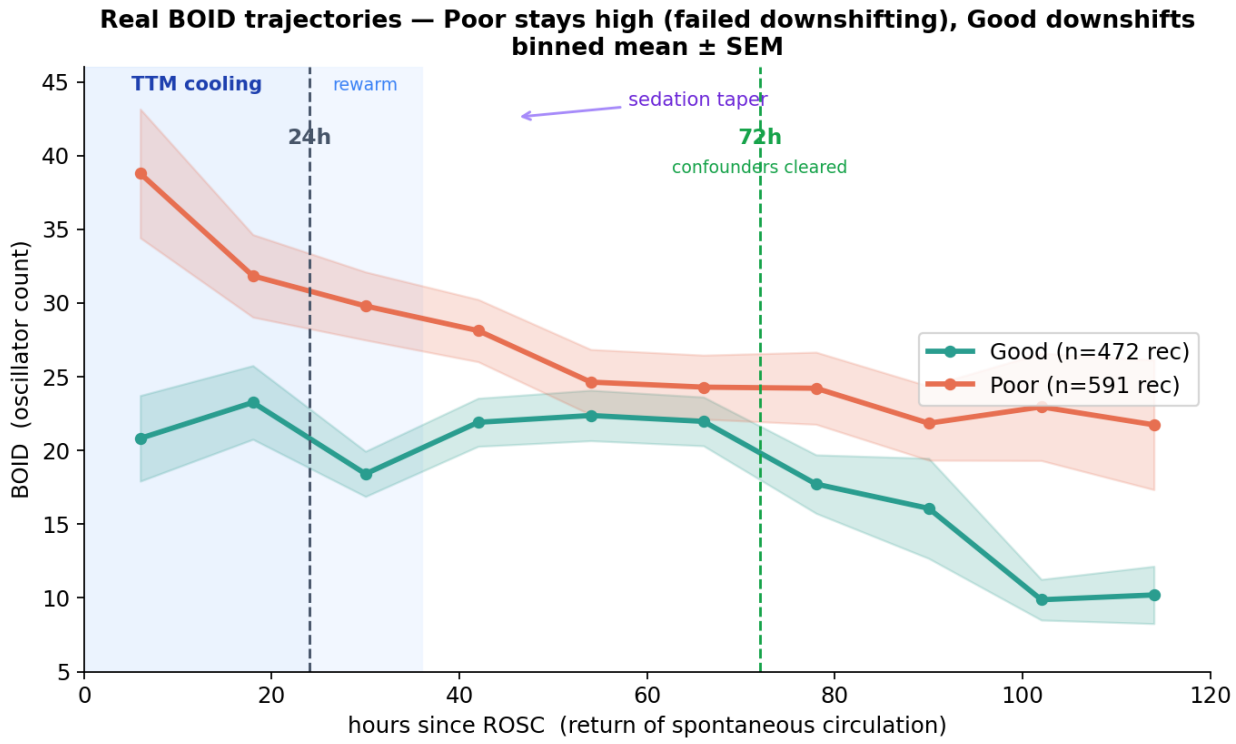

Figure 3: **Cohort dimensional trajectories.** Binned mean ± SEM BOID against time (472 Good / 591 Poor recordings). Poor-outcome BOID stays high (failed downshifting) while Good-outcome BOID declines toward a low-dimensional recovery state; the separation widens after 72 h.

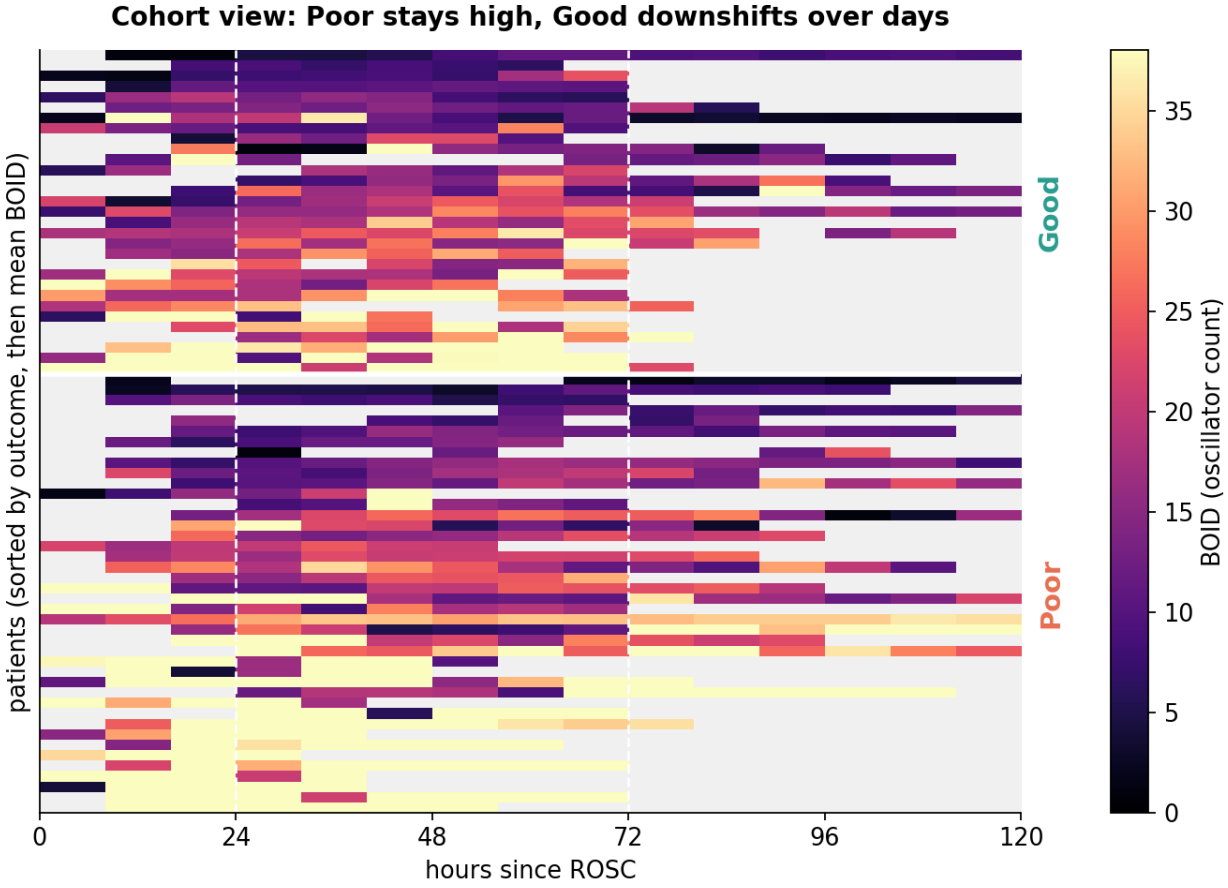

Figure 4: **Per-patient BOID across the cohort.** Each row is one patient's BOID trajectory, sorted by outcome. Poor-outcome patients remain high-dimensional; Good-outcome patients downshift toward low BOID, reproducing the failed-downshifting pattern at the level of individual patients rather than pooled windows.

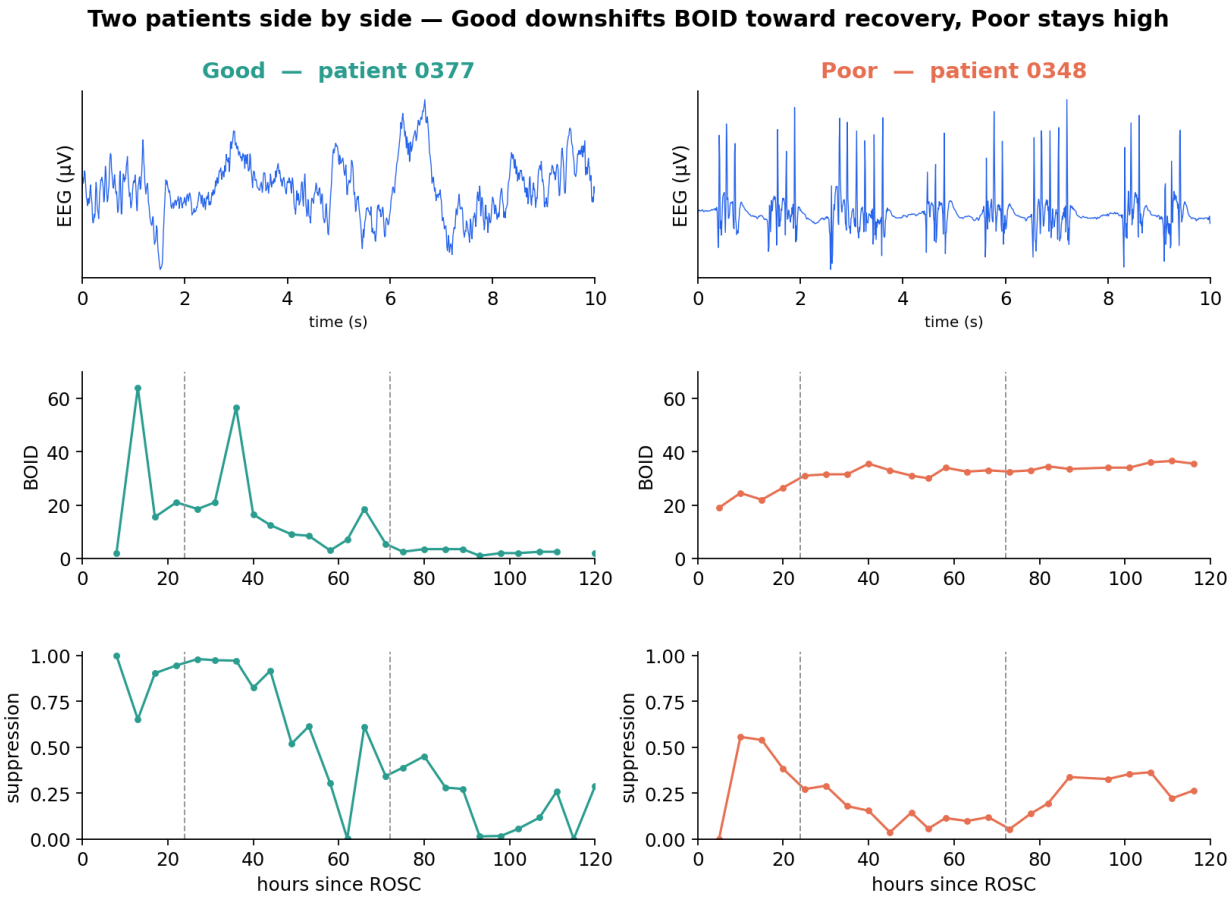

Figure 5: **Two further representative patients.** EEG, BOID, and suppression for one Good- and one Poor-outcome patient (distinct from the main-text example). The Good patient downshifts BOID and clears suppression toward recovery; the Poor patient remains high-dimensional.

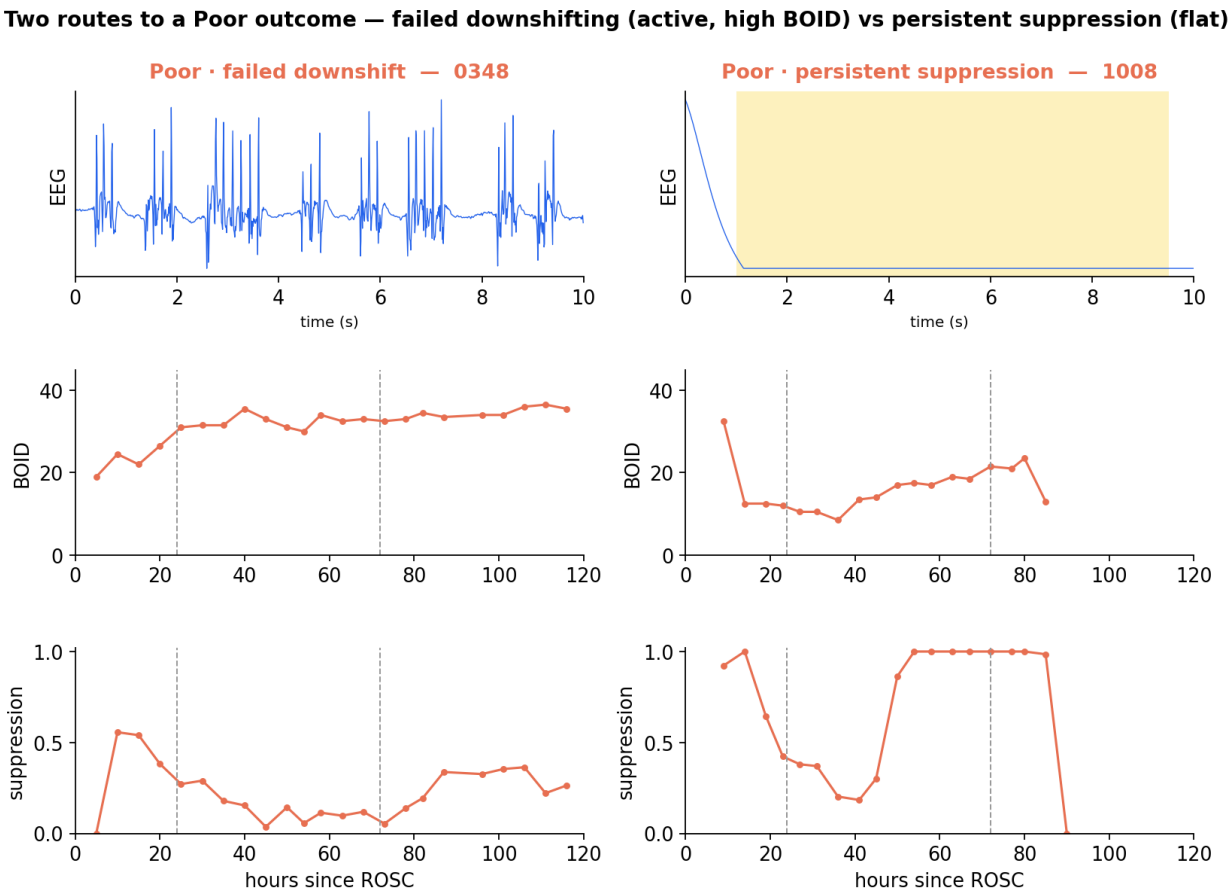

Figure 6: **Two routes to a poor outcome.** Failed downshifting (active EEG, persistently high BOID) versus persistent suppression (near-flat EEG, high suppression ratio). Both are poor-prognosis but occupy opposite ends of the amplitude axis, which is why BOID and suppression contribute complementary information.

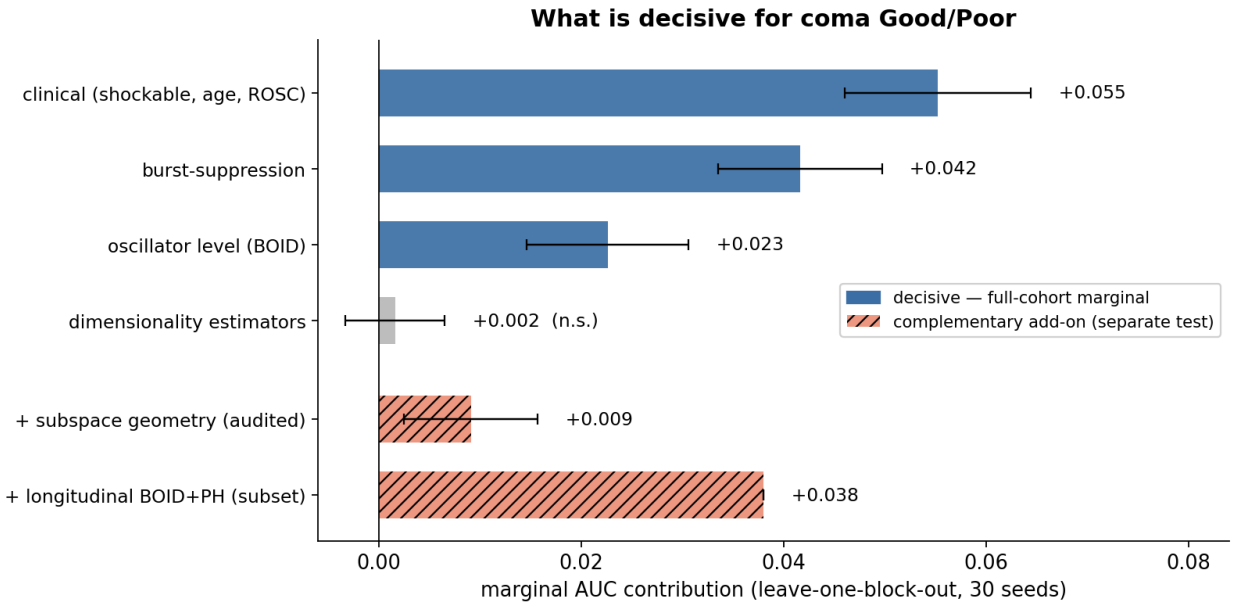

Figure 7: **Marginal contribution of each feature block.** Leave-one-block-out AUC change over 30 patient-grouped cross-validation seeds. Clinical variables and burst-suppression carry the largest unique signal, the oscillator count a smaller but consistent one, and the intrinsic-dimension battery little marginal value once the others are present.

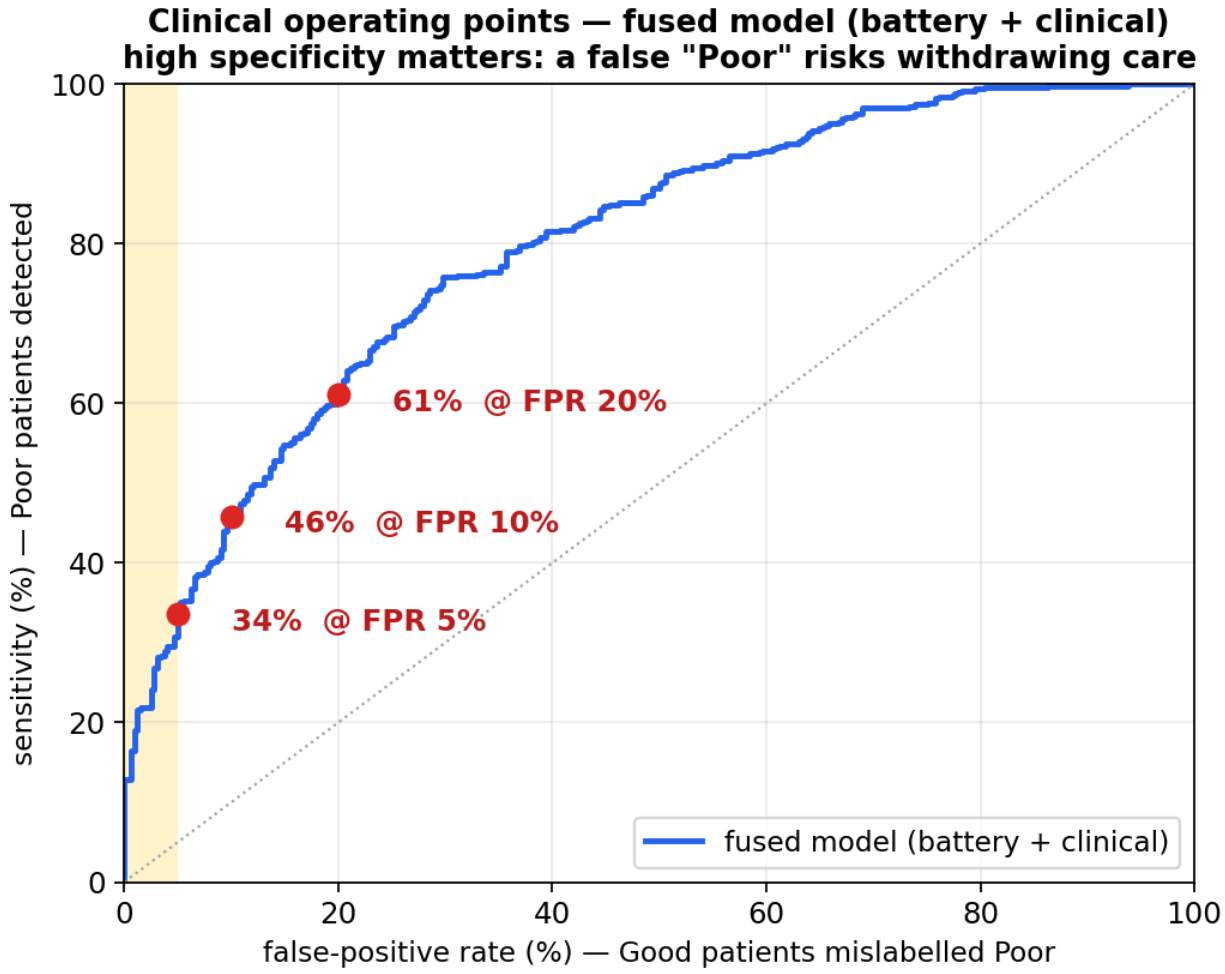

Figure 8: **Clinical operating points.** Sensitivity for poor outcome at fixed false-positive rates for the full fused model (intrinsic-dimension battery + clinical): 0.34 at 5 % FPR, 0.46 at 10 %, and 0.61 at 20 %. Because a false "poor" risks inappropriate withdrawal of care, the high-specificity end is the clinically relevant operating region; the AUC advantage is realized only at more permissive thresholds.

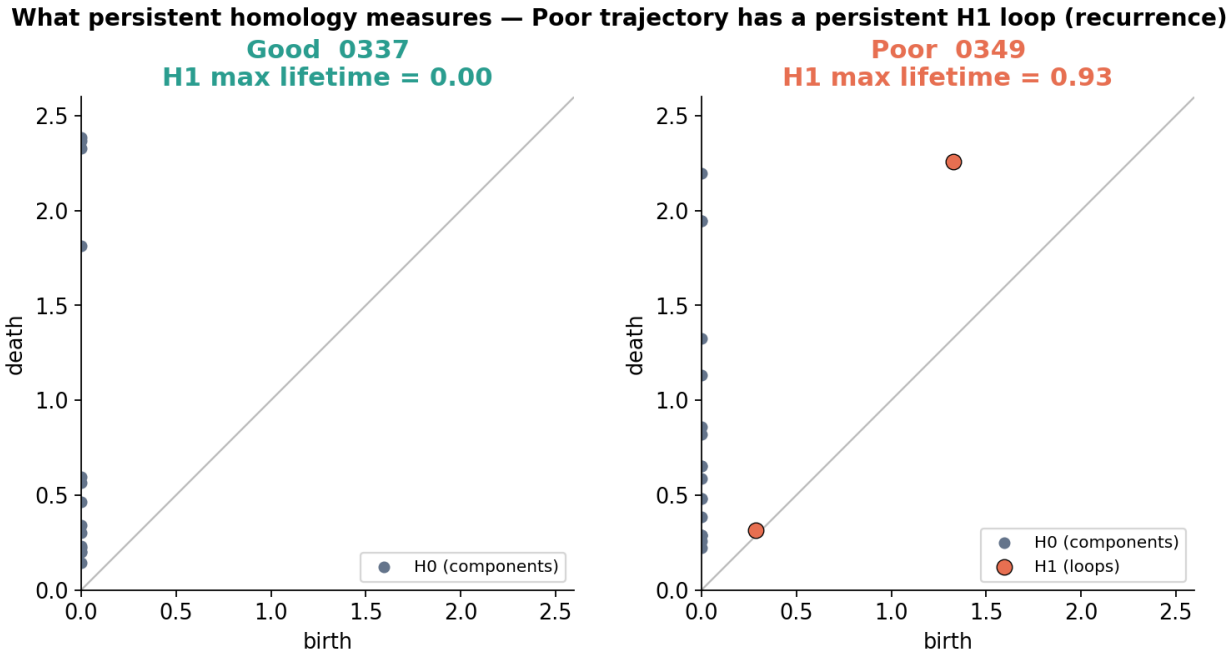

Figure 9: **Persistent homology of the dimensional trajectory.** Persistence diagrams of the time-delay (Takens) embedding of two patients' BOID trajectories: the poor-outcome trajectory carries a persistent H1 loop (recurrent, non-settling dynamics) absent in the good-outcome trajectory.

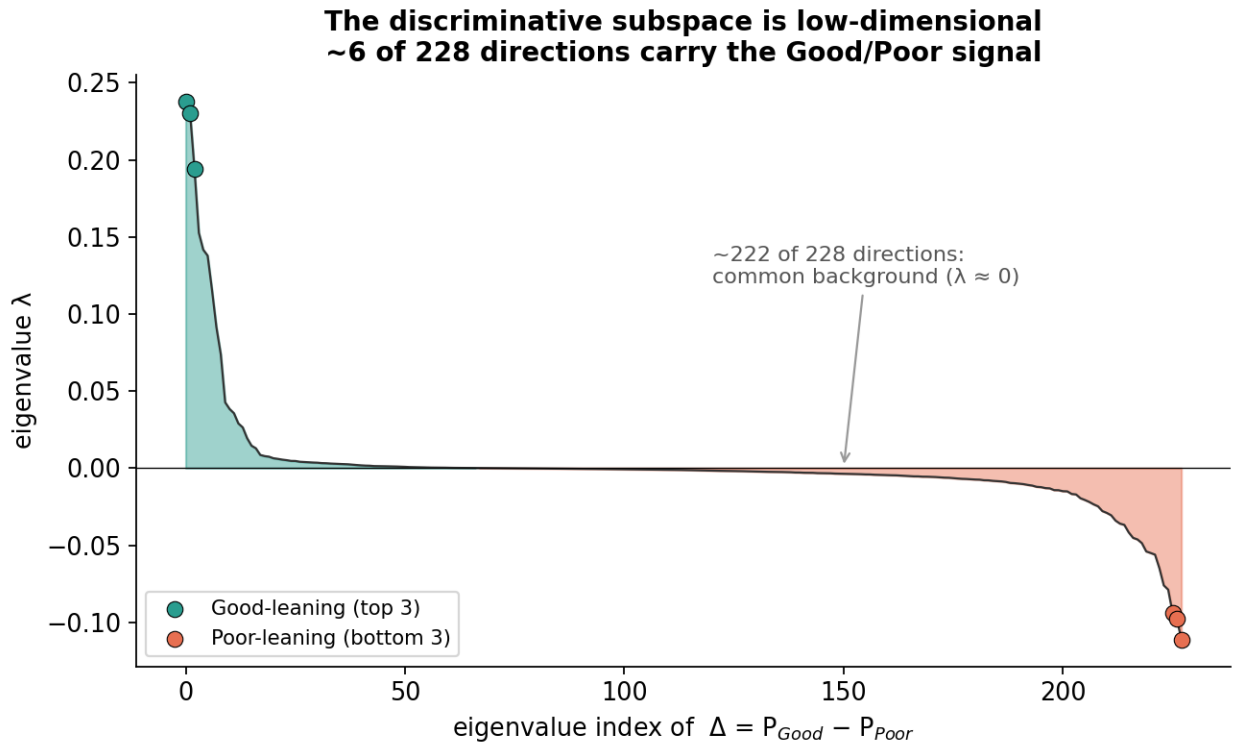

Figure 10: **The discriminative subspace is low-dimensional.** Eigenvalue spectrum of the class-projector difference Δ = P_Good − P_Poor (ambient dimension 228): only a handful of directions carry non-zero eigenvalue, the remainder being common background. This motivates the top-k truncation in the discriminative-subspace classifier.