Research Square

Preprints are preliminary reports that have not undergone peer review.
They should not be considered conclusive, used to inform clinical practice,
or referenced by the media as validated information.

Manifold polarity as the geometric cause of cross patient failure in Riemannian EEG seizure detection

James Appenzeller

jamesappenzeller@gmail.com

National University

Research Article

Keywords: EEG seizure detection, manifold polarity, Riemannian geometry, quantum machine learning, cross-patient generalization

Posted Date: October 6th, 2026

DOI: https://doi.org/10.21203/rs.3.rs-10787059/v1

License: This work is licensed under a Creative Commons Attribution 4.0 International License.
Read Full License

Additional Declarations: No competing interests reported.

Abstract

Cross-patient EEG seizure detection using Riemannian geometry on the symmetric positive definite (SPD) covariance matrix manifold has consistently underperformed eigenvalue-based features, a finding attributed to inter-subject variability without a specific geometric mechanism. We identify patient-specific manifold polarity as the mechanism: seizure-related covariance shifts point in opposite directions for different patients on the SPD manifold. A previously unreported feature-regime interaction governs method comparisons: polarity prevalence scales with feature dimensionality, not method choice.

We validate this across 22 CHB-MIT patients (leave-one-subject-out). In the full-feature regime (336 features), AUC reaches 0.952 with zero polarity inversions. In the low-dimensional regime (8–36 features), Riemannian tangent space projection (0.659 calibrated AUC, 36 features), the A-Gate quantum circuit (0.637 calibrated, 8 features), and minimal eigenvalues (0.628 calibrated, 8 features) are statistically indistinguishable (Wilcoxon $p = 0.4358$). Polarity was discovered not through classical analysis but through quantum hardware experiments. The A-Gate circuit, a quantum geometric probe with a fixed measurement basis, makes polarity an explicit observable: per-patient AUC below 0.50 directly indicates manifold inversion, because the circuit's architectural rigidity prevents the adaptive absorption of polarity that flexible classifiers perform silently.

Statevector simulation of the same circuit performs worse than noisy hardware, a reversal of the standard assumption that simulation provides an optimistic upper bound. Simulation AUC degrades from 0.460 to 0.386 as input quality improves, because better-conditioned features sharpen polarity-induced amplitude cancellation; hardware shot noise partially breaks this cancellation stochastically, though the benefit diminishes as polarity projections strengthen at longer windows. The divergence is a property of the problem geometry, not measurement error. Manifold polarity, defined as the projection of a patient's covariance shift vector onto the circuit's measurement basis, is a continuous geometric quantity whose sign is encoding-dependent and whose magnitude determines measurement reliability. Whether the projection structure generalizes across encoding configurations, and what it reveals about the underlying manifold geometry, are the questions that motivate subsequent work.

1. Introduction

Seizure detection from electroencephalography (EEG) is a clinically critical task with significant unmet need: epilepsy affects approximately 50 million people worldwide, and roughly one third remain refractory to pharmacological treatment [WHO, 2019]. Automated systems must generalize across patients; individuals differ substantially in seizure morphology, electrode placement sensitivity, and baseline dynamics, and cross-patient generalization remains the field's central challenge [Yger et al., 2017]. Methods that perform well within a single subject frequently degrade when evaluated in leave-one-subject-out (LOSO) settings, a discrepancy that has been documented on the CHB-MIT dataset and

others but has not been explained at the level of specific geometric mechanism [Shoeb, 2009; Congedo et al., 2017].

The symmetric positive definite (SPD) manifold has emerged as a principled representation for EEG covariance matrices [Barachant et al., 2012; Pennec et al., 2006]. Riemannian geometry on the SPD manifold offers a theoretically rich framework: the geodesic distance between covariance matrices respects the manifold's curvature, and classification via tangent space projection provides a locally Euclidean approximation that is equivariant to electrode permutation and invariant to certain montage differences. In practice, however, Riemannian methods have consistently underperformed simpler eigenvalue-based representations in cross-patient evaluations [Zanini et al., 2018; Wu et al., 2020].

This paper identifies the geometric mechanism underlying this discrepancy. We show that seizure-related covariance shifts point in opposite directions on the SPD manifold for different patients, a phenomenon we term manifold polarity. A patient with standard polarity has ictal states whose covariance matrices lie on one side of the Riemannian reference point (Fréchet mean) [Fletcher & Joshi, 2004]; a patient with inverted polarity has ictal states on the opposite side. Riemannian tangent space projection with a single linear classifier cannot accommodate this, because the projection preserves the directional structure that makes patients incompatible [Barachant et al., 2012]. Sorted eigenvalue representations succeed precisely because sorting destroys directionality: an increasing and a decreasing eigenvalue spectrum both produce a sorted vector that points away from the interictal baseline.

This finding was not reached through classical analysis alone. The polarity structure was first made apparent through experiments with a quantum positive-negative (PN) neuron circuit (the A-Gate architecture) run on IBM Heron superconducting hardware [IBM Quantum, 2024]. The A-Gate circuit encodes EEG phase-locking value (PLV) features [Lachaux et al., 1999] as rotation angles in a quantum circuit and measures the expectation value of a fixed observable. The fixed measurement basis, motivated by the PN neuron's explicit separation of excitatory and inhibitory sub-units [Gupta et al., 2003], means the circuit cannot adaptively absorb polarity: patients with inverted polarity produce AUC values below 0.50, making the inversion an explicit, numerically observable event rather than a hidden variable compensated by classifier flexibility.

Beyond the polarity mechanism itself, we report a previously undescribed phenomenon in quantum machine learning: statevector simulation of the A-Gate circuit performs worse than noisy hardware at short windows, and both backends degrade as input data quality improves. The standard assumption in quantum ML is that noiseless simulation provides an optimistic upper bound on hardware performance [Preskill, 2018; Cerezo et al., 2021]. At 1.95-second windows, simulation achieves 0.460 AUC while hardware achieves 0.524 raw and 0.637 calibrated. At 20-second windows, simulation degrades to 0.386 (6/7 patients inverted) while hardware also degrades to 0.404 raw / 0.605 calibrated (6/7 inverted). We attribute this inversion to polarity-induced amplitude cancellation: when opposing-polarity patients

contribute with similar feature magnitudes, their exact-amplitude representations cancel in aggregate. Hardware shot noise breaks the cancellation stochastically, producing net positive discrimination.

The anomalous per-patient AUC values (patients producing AUC below 0.50 on a circuit with no trained parameters) were first reported in the initial A-Gate hardware validation study [Appenzeller, 2026]. That work established the circuit architecture and identified the encoding bottleneck; the geometric explanation for the inversion was not available. The present paper takes that anomalous quantum observation as its starting point: we show that sub-chance AUC on a fixed-basis circuit is not a failure mode but a direct readout of manifold polarity, validate the mechanism across 22 patients using purely classical methods, and establish that the phenomenon is general to any direction-preserving classifier, quantum or classical. The quantum circuit's role is thus retrospectively diagnostic: it made a geometric structure visible that had been present in the data throughout.

Our contributions are: (1) identification of patient-specific manifold polarity as the geometric mechanism explaining Riemannian cross-patient failure; (2) demonstration that sorted eigenvalues are implicitly polarity-invariant, explaining their consistent advantage; (3) a quantum geometric probe that makes polarity an explicit observable; (4) characterization of a simulation-hardware divergence in quantum ML that is fundamental to the geometry of the problem, not a noise artifact; and (5) evidence that manifold polarity is a continuous projection quantity rather than a binary patient label, with encoding-dependent sign and measurable magnitude.

2. Background

2.1 EEG Seizure Detection

Seizure detection from scalp EEG involves identifying transitions from interictal (baseline) to ictal (seizure) states in multi-channel time series. The CHB-MIT scalp EEG dataset [Shoeb, 2009] is the standard benchmark, comprising recordings from 23 pediatric patients with refractory epilepsy. Cross-patient evaluation, defined as training on all patients except one and testing on the left-out patient, provides the most clinically realistic performance estimate, as deployed systems must generalize to new patients without retraining.

Published baselines on CHB-MIT vary considerably with evaluation protocol. Systems using within-patient validation routinely exceed 0.90 AUC; cross-patient LOSO evaluations typically yield 0.62–0.82 AUC depending on feature set and windowing [Wu et al., 2020; Usman et al., 2021]. XGBoost with spectral features has been reported at 0.625–0.820 AUC in LOSO settings. Competition-winning solutions on the related Kaggle seizure detection challenge used 20–30 second windows with maximum probability aggregation across sub-windows [Barachant et al., 2017], a design confirmed by our own experiments: Tier 2 eigenvalue features improve from 0.723 AUC at 1.95-second windows to 0.873 AUC at 20-second windows.

2.2 Riemannian Geometry for EEG

The space of symmetric positive definite (SPD) matrices forms a Riemannian manifold with a natural metric. For an n-channel EEG recording, the sample covariance matrix is an n×n SPD matrix. The affine-invariant metric on the SPD manifold [Pennec et al., 2006] defines geodesic distances that are invariant to linear transformations of the EEG signal, which provides robustness to electrode impedance variations and common-mode noise.

Tangent space projection (TSP) maps the SPD manifold to a Euclidean tangent space at a reference point P (typically the Fréchet mean of the training set):

$$S_{tan} = \mathbf{P}^{-1/2} \, log \left(\mathbf{P}^{-1/2} \mathbf{S} \mathbf{P}^{-1/2} \right) \mathbf{P}^{1/2}$$

The resulting symmetric matrix is vectorized (upper triangle) to produce a Euclidean feature vector suitable for standard classifiers. TSP has shown strong within-subject performance across BCI paradigms [Barachant et al., 2012], but cross-patient performance has been inconsistent.

Prior work attributes cross-patient TSP failure to inter-subject variability in covariance structure [Blankertz et al., 2008; Congedo et al., 2017; Barachant et al., 2012]. We show here that the failure has a more specific cause, polarity inversion on the manifold, and that this cause is visible across the full 22-patient CHB-MIT cohort (27% of patients inverted under Riemannian, 45% under minimal eigenvalue features). Domain adaptation methods address reference frame misalignment; they do not address the case where the class-separating direction itself points in opposite directions for different patients.

2.3 Positive-Negative (PN) Neuron Architecture

The positive-negative (PN) neuron is a biologically inspired neural architecture that explicitly separates excitatory (positive) and inhibitory (negative) sub-units within a single neuron [Gupta, Jin, and Homma, 2003]. In the classical formulation, the excitatory unit E and inhibitory unit I receive input through separate weighted pathways and interact through cross-connections w_{IE} and w_{EI} [see Appenzeller, 2026, Fig. 1]. Each sub-unit maintains its own state variable, decay constant, and activation function. The neuron's output reflects the balance between excitatory and inhibitory drives, a structure that maps naturally to the excitatory-inhibitory (E-I) balance that regulates cortical dynamics and is disrupted during seizures [van Hugte et al., 2023].

The A-Gate architecture implements this PN structure as a quantum circuit. RZ rotation gates encode the excitatory and inhibitory components of the EEG feature vector as rotation angles, and ECR (echoed cross-resonance) entangling gates implement the cross-connections w_{IE} and w_{EI}. Measurement in the Z basis at the output qubit produces a binary classification output whose expectation value corresponds to the net E-I balance of the encoded brain state.

2.4 Quantum Hardware for EEG Classification

The practical use of current quantum hardware for machine learning faces two fundamental limitations: circuit depth and the simulation-hardware gap. Near-term devices (NISQ era) [Preskill, 2018] support

circuits of limited depth before decoherence degrades fidelity. Statevector simulation, which provides exact quantum amplitude computation, is commonly used as a development proxy, but for circuits encoding domain-specific data with geometric structure, simulation and hardware can produce qualitatively different results. This paper reports a case where this divergence is not incidental but mechanistically explained.

3. Methods

3.1 Dataset and Preprocessing

We used the CHB-MIT Scalp EEG Database [Shoeb, 2009], a publicly available dataset of continuous EEG recordings from 23 pediatric patients with intractable epilepsy, recorded at Children's Hospital Boston. Recordings use the international 10–20 system at 256 Hz sampling rate. For classical experiments, we used all 22 available patients (chb12 excluded due to missing channel data), yielding a total of 163 seizures across patients. For quantum hardware experiments, a 7-patient subset (chb01, chb03, chb05, chb07, chb11, chb14, chb21) representing a range of seizure frequencies was used, consistent with prior quantum cost constraints; these 7 patients contributed 37 seizures. All classical results in this paper are reported on the 22-patient cohort unless otherwise noted.

Preprocessing applied a 5th-order Butterworth bandpass filter (0.5–40 Hz, zero-phase via filtfilt), a 60 Hz notch filter (Q = 30) to remove power line interference, and per-channel z-score normalization. The 40 Hz upper cutoff follows standard clinical scalp EEG convention, removing muscle artifact while preserving delta through beta activity; this is more conservative than competition-winning pipelines that retained gamma content up to 128–180 Hz [Barachant et al., 2017; Temko et al., 2011], meaning our Tier 2B performance ceiling (0.952 AUC) is achieved without high-gamma features. EEG signals were segmented into non-overlapping windows of varying duration (1.95, 5, 10, 15, 20, and 30 seconds) to characterize the effect of window size on each method. For quantum circuit experiments, 1.95-second and 20-second windows were used; the shorter window matches our initial hardware experiments and the longer window aligns with the 20-second segmentation used in the winning solution to the Melbourne University seizure prediction challenge [Barachant et al., 2017].

Interictal segments were extracted from periods at least 30 minutes from any seizure onset or termination to reduce preictal contamination. Class balance was maintained at approximately 1:1 ictal-to-interictal by random sampling of interictal windows. Channel selection used a consistent 8-channel montage (FP1-F7, F7-T7, T7-P7, P7-O1, FP2-F8, F8-T8, T8-P8, P8-O2) providing bilateral temporal coverage relevant to seizure detection.

3.2 Evaluation Protocol

All methods were evaluated using leave-one-subject-out (LOSO) cross-validation. In each fold, the model is trained on all patients except one and evaluated on the held-out patient. Final performance is reported as mean area under the ROC curve (AUC) [Fawcett, 2006] across all 7 folds, with per-patient AUC

reported separately to characterize polarity-related variance. The same evaluation protocol, channel configuration, and class balance were maintained across all three methods (Tier 2, Tier 3, and quantum) to ensure valid cross-method comparisons. Statistical comparisons used the Wilcoxon signed-rank test, selected because per-patient AUC distributions are non-normal (Shapiro-Wilk $p < 0.05$ for all methods) and the sample sizes are small. Effect sizes are reported as rank-biserial correlation r alongside p-values.

3.3 Tier 2A and 2B: Eigenvalue Features with XGBoost

For each EEG window, we computed the 8×8 sample covariance matrix across the 8 EEG channels. The eigenvalue decomposition produces 8 non-negative eigenvalues sorted in descending order to form the feature vector. For a window containing N samples, the sample covariance matrix is:

$$\mathbf{C} = \frac{1}{N}\mathbf{X}\mathbf{X}^{\top}$$

where X is the 8×N channel data matrix with z-score normalized rows. The sorted eigenvalue vector is polarity-invariant: eigenvalues of C are invariant to sign changes in the columns of X, meaning patients whose seizures manifest as increased or decreased correlation both produce distinguishable eigenvalue patterns.

Tier 2A uses the 8 sorted eigenvalues as the complete feature vector. Tier 2B extends this to a 336-dimensional feature set inspired by the winning submission to the UPenn and Mayo Clinic seizure detection challenge [Hills, 2014]: for each 1-second sub-window, log10 FFT magnitudes across 40 frequency bins (1–40 Hz) are computed per channel (8 × 40 = 320 features), concatenated with eigenvalues of the inter-channel correlation matrix in both frequency domain (8 features) and time domain (8 features). The frequency-domain correlation matrix is computed from the column-normalized FFT magnitude matrix; the time-domain correlation matrix is computed directly from the raw sub-window signals. Tier 2B was included to establish the classical ceiling and to test whether polarity inversion persists when feature dimensionality provides redundant representation. The contrast between Tier 2A and Tier 2B isolates the effect of feature dimensionality from method choice.

For multi-window aggregation, we used MAX pooling: the maximum classification probability across sub-windows within a segment was taken as the segment-level prediction. This consistently outperformed MEAN pooling by 10–15 AUC points across all conditions, consistent with the transient signal hypothesis that ictal patterns appear in bursts rather than uniformly across a segment. Post-competition analysis of the Melbourne University challenge confirmed this pattern: MAX aggregation over 20-second sub-windows outperformed MEAN by 13 percentage points [Barachant et al., 2017].

Classification used XGBoost [Chen & Guestrin, 2016] with hyperparameters tuned on the training fold: n_estimators = 200, max_depth = 5, learning_rate = 0.05, subsample = 0.8, colsample_bytree = 0.8. These parameters were fixed across all window sizes and patient folds without per-patient tuning.

3.4 Tier 3: Riemannian Geometry with Tangent Space Projection

For each EEG window, the 8×8 sample covariance matrix was computed identically to Tier 2. The Riemannian Fréchet mean of all training covariance matrices was computed using the iterative algorithm of Fletcher and Joshi [2004], implemented via the pyRiemann library [Barachant et al., 2024]. Tangent space projection was then applied at the Fréchet mean, and the upper triangle of the resulting symmetric matrix was vectorized to produce a 36-dimensional feature vector. Classification used linear discriminant analysis (LDA) with automatic shrinkage (Ledoit-Wolf estimator [Ledoit & Wolf, 2004]).

Polarity calibration was applied post-hoc for analysis purposes: for each patient, we computed the sign of the median(AUC_ictal - AUC_interictal) across the training fold, and applied a sign flip to the test predictions when this value was negative. This calibration represents the minimum information required to resolve polarity: a single bit per patient, obtainable from a small calibration recording. We emphasize that this calibration is applied only to measure the polarity ceiling and is not proposed as a deployment strategy without modification.

3.5 A-Gate Quantum Circuit

The A-Gate architecture implements a quantum positive-negative neuron using RZ rotation gates and ECR entangling gates on IBM Heron superconducting hardware [IBM Quantum, 2024]. The circuit encodes EEG phase-locking value (PLV) features in the theta (4–8 Hz) and alpha (8–15 Hz) bands as rotation angles, exploiting the theoretical connection between PLV and quantum phase coherence.

Feature extraction for quantum encoding: for each 8-channel EEG window, we computed the PLV matrix across all 28 channel pairs in the theta and alpha bands [Lachaux et al., 1999]. For channel pair (i, j), instantaneous phases were extracted via the Hilbert transform and PLV was computed as:

$$PLV_{ij} = \left| \frac{1}{N} \sum_{t=1}^{N} e^{i\left(\varphi_i(t) - \varphi_j(t)\right)} \right|$$

where φ_i is the instantaneous phase of channel i. PLV values were mapped to rotation angles:

$$\theta_{ij} = \pi \cdot PLV_{ij}$$

producing a 56-dimensional encoding vector (28 channel pairs × 2 frequency bands).

Circuit specification (CH8 configuration): the logical circuit uses 8 qubit pairs (16 qubits total plus 1 ancilla = 17 qubits), where each pair corresponds to one EEG channel pair's PLV encoding. The excitatory sub-circuit applies RZ(θ_E) rotations to the first qubit of each pair; the inhibitory sub-circuit applies RZ(θ_I) to the second qubit. ECR (echoed cross-resonance) gates between excitatory and inhibitory qubits implement the cross-connections w_{IE} and w_{EI}.

Hardware execution: circuits were transpiled for IBM Heron r1 (ibm_torino, 133 qubits) and executed with 1024 shots per circuit; no error mitigation was applied [Qiskit Contributors, 2023]. The raw classification score for each window was computed from the Z-basis expectation value:

$$\langle Z \rangle = \frac{N_0 - N_1}{1024}$$

where N_0 and N_1 are the counts of $|0\rangle$ and $|1\rangle$ outcomes respectively. Polarity calibration was applied identically to the Riemannian method.

Statevector simulation used Qiskit's statevector simulator [Qiskit Contributors, 2023] with identical circuit parameters, producing exact quantum amplitudes without shot noise. Both 1.95-second and 20-second window experiments were run in simulation; hardware experiments at 20-second windows were executed on ibm_torino in March 2026.

4. Results

4.1 Feature Regime and Classification Performance

Table 1 presents the primary comparison across 22 patients at 20-second windows. The dominant finding is a feature-regime interaction: performance differences between methods are driven by feature dimensionality, not by the choice of geometric representation.

Table 1

Classification AUC across methods (22-patient cohort, 20-second windows, LOSO cross-validation).
Quantum results from 7-patient hardware subset only. Tier 2B uses 336 features; Tier 2A and quantum use 8 features.

Method	Features	AUC (raw)	AUC (calibrated)	N inverted / total	% inverted
22-patient cohort (CHB-MIT, all except chb12), 20-second windows, LOSO cross-validation					
Tier 2B: Full eigenvalues + XGBoost	336	0.952	0.952 (+ 0.000)	0 / 22	0%
Tier 3: Riemannian + LDA	36	0.536	0.659 (+ 0.123)	6 / 22	27%
Tier 2A: Sorted eigenvalues + XGBoost	8	0.521	0.628 (+ 0.107)	10 / 22	45%
7-patient subset (quantum hardware experiments), 1.95s and 20s windows					
A-Gate quantum (hardware, 1.95 s)	8	0.524	0.637 (+ 0.113)	3 / 7	43%
A-Gate quantum (simulation, 1.95 s)	8	0.460	0.540 (+ 0.080)	~ 3 / 7	~ 43%
A-Gate quantum (simulation, 20 s)	8	0.386	0.630 (+ 0.244)	6 / 7	86%
A-Gate quantum (hardware, 20 s)	8	0.404	0.605 (+ 0.201)	6 / 7	86%

Tier 2B (336 eigenvalue features) achieves 0.952 raw AUC with no polarity inversion across any of the 22 patients. Feature redundancy (computing eigenvalues separately per frequency band and aggregating across bands) provides enough representation capacity that the per-patient polarity structure becomes irrelevant to classifier performance. Polarity calibration produces zero lift (0/22 inversions).

In the low-dimensional regime, all three methods are statistically indistinguishable after calibration: Tier 3 Riemannian 0.659 (36 features), A-Gate quantum 0.637 (8 features), Tier 2A eigenvalue 0.628 (8 features). Wilcoxon signed-rank tests confirm this structure: Tier 3 vs. Tier 2A cannot be distinguished ($p = 0.4358$, $r = 0.12$), while either low-dimensional method is significantly below Tier 2B ($p < 0.001$, $r > 0.80$). Table 2 reports the complete statistical comparisons.

Table 2
Wilcoxon signed-rank tests (22-patient cohort, 20-second windows, calibrated AUC).

Comparison	Wilcoxon p-value	Interpretation
Tier 3 vs. Tier 2A (36 vs. 8 features)	0.4358	Not significant — methods are statistically indistinguishable at matched dimensionality
Tier 3 vs. Tier 2B (36 vs. 336 features)	0.0001	Significant — full eigenvalue features dominate
Tier 2A vs. Tier 2B (8 vs. 336 features)	0.0000	Significant — dimensionality is the primary driver

The practical implication is that the comparison of Riemannian vs. eigenvalue methods in prior literature is confounded by feature dimensionality. In the same-dimensionality regime, the two approaches are equivalent in classification performance. The Riemannian approach's calibrated advantage (0.659 vs. 0.628, $\Delta = 0.031$) is non-significant; the uncalibrated gap is smaller still (0.536 vs. 0.521, $\Delta = 0.015$), because both methods suffer polarity-induced suppression without calibration, with eigenvalues affected more given their higher inversion rate (10/22 vs. 6/22).

4.2 Window Size Effects (7-patient subset)

Table 3 shows the window size sweep for Tier 2A and Tier 3 on the 7-patient subset. These results guided the choice of 20-second windows for the full 22-patient comparison.

Table 3
Window size sweep (7-patient subset). Covariance condition number decreases with window size, indicating better estimation quality. *Tier 2A uses MAX pooling across sub-windows. Tier 3 calibrated figures use per-patient polarity sign bit.*

Window	Tier 2A (MAX)	Tier 3 (raw)	Tier 3 (calibrated)	Covariance cond. #
1 s	0.803	0.516	0.696	67.3
2 s	0.785	0.520	0.695	60.1
5 s	0.795	0.510	0.702	55.7
10 s	0.773	0.525	0.716	51.5
15 s	0.778	0.551	0.722	50.8
20 s	0.873	0.570	0.690	49.0
30 s	0.834	0.543	0.724	47.4

Tier 2A shows a non-monotonic profile: AUC peaks at 20 seconds (0.873) and declines slightly at 30 seconds (0.834), suggesting diminishing returns from longer windows beyond 20 seconds. Tier 3 raw AUC is flat across all seven window sizes (0.51–0.57), confirming that covariance estimation quality is not the bottleneck. Tier 3 calibrated AUC is stable around 0.690–0.724. Notably, Tier 3 LDA polarity is

100% consistent per patient across all seven classical window sizes (chb03, chb14, chb21 always inverted; chb01, chb05, chb07, chb11 always standard), contrasting sharply with the A-Gate hardware results where 4/7 patients show encoding-dependent polarity flips (see § 4.3). This indicates that polarity instability in the quantum results arises from the PLV-based encoding projection, not from the underlying covariance geometry.

4.3 Manifold Polarity Prevalence and Structure

Table 4 shows per-patient polarity assignments for the 7-patient quantum subset. Polarity prevalence scales with feature dimensionality: 45% of the 22-patient cohort shows inversion with Tier 2A (10/22), 27% with Riemannian Tier 3 (6/22), and 0% with the full Tier 2B representation. The quantum hardware rate (43%, 3/7) matches the Tier 2A rate in the same subset, consistent with both methods operating in the 8-feature regime.

Table 4

Per-patient polarity assignments across Riemannian, quantum hardware, and statevector simulation (7-patient subset). *Discordant patients (chb11, chb14) are highlighted. QHW = quantum hardware, QSim = statevector simulation.*

Patient	Riemannian polarity	QHW polarity	QSim polarity	QHW concordant?
chb01	Standard (+)	Standard (+)	Inverted (−)	Yes
chb03	Inverted (−)	Inverted (−)	Inverted (−)	Yes
chb05	Standard (+)	Standard (+)	Standard (+)	Yes
chb07	Standard (+)	Standard (+)	Standard (+)	Yes
chb11	Standard (+)	Inverted (−)	Inverted (−)	**No**
chb14	Inverted (−)	Standard (+)	Standard (+)	**No**
chb21	Inverted (−)	Inverted (−)	Inverted (−)	Yes

Within a given method, polarity consistency across encoding configurations varies by patient. Two patients (chb03, chb21) show stable inverted polarity across all signal-bearing hardware configurations, while four patients (chb01, chb05, chb07, chb14) show encoding-dependent sign flips. This indicates that polarity is not a fixed binary label but a continuous projection: strongly polar patients have large covariance shift magnitudes that project consistently regardless of encoding, while weakly polar patients sit near the sign boundary and resolve differently depending on channel count and window size. Cross-method concordance for quantum hardware vs. Riemannian is 71% (5/7), and across all three methods (including simulation) is 40.9%. The low all-method concordance is driven entirely by Tier 2B having zero inversions, a consequence of its high-dimensional representation rather than polarity absence. Per-patient AUC values for all 22 CHB-MIT patients under each method are available in the project repository at https://github.com/jappenzeller/QDNU.

Encoding-dependent polarity is the norm, not the exception. Of the 7 hardware patients, only chb03 and chb21 maintain consistent polarity signs across all configurations (both always inverted). The remaining patients show sign flips between channel counts or window sizes, confirming that polarity is a projection of the patient's covariance shift vector onto the specific measurement basis defined by the encoding. Different encodings probe different axes of the same underlying manifold, and the sign of the projection depends on the angle between the shift vector and the measurement axis.

4.4 Quantum Hardware vs. Statevector Simulation

At 1.95-second windows, statevector simulation produces 0.460 aggregate AUC while hardware produces 0.524 raw and 0.637 calibrated. At 20-second windows, simulation degrades to 0.386 (6/7 inverted) and hardware degrades to 0.404 raw / 0.605 calibrated (6/7 inverted). Both backends show increased polarity inversion at 20 seconds, consistent with longer windows producing more stable PLV estimates that resolve borderline patients away from chance. The simulation-hardware gap narrows at 20 seconds (0.386 vs. 0.404 raw) compared to 1.95 seconds (0.460 vs. 0.524 raw), suggesting that stronger polarity projections partially overcome the stochastic advantage of shot noise.

After oracle calibration, simulation at 20 seconds (0.630) approximately matches hardware at 1.95 seconds (0.637), but hardware at 20 seconds drops to 0.605. The degradation of hardware calibrated AUC from 1.95s to 20s indicates that longer windows do not simply improve feature quality: they also shift borderline patients into stronger polarity projections, reducing the number of patients contributing positive discrimination.

5. Discussion

5.1 Polarity as a Missing Mechanism in Riemannian BCI

The cross-patient failure of Riemannian methods has been documented repeatedly in the BCI literature [Barachant et al., 2012; Congedo et al., 2017; Yger et al., 2017]. Proposed solutions include recentering individual covariance matrices to a shared reference [Zanini et al., 2018], Procrustes analysis combining recentering with stretching and rotation [Rodrigues et al., 2019], and parallel transport on the SPD cone [Yair et al., 2019]. These methods address reference frame misalignment in a principled way and do improve cross-subject performance, but none identifies polarity inversion as the specific geometric mechanism at work. Recentering aligns distribution means; it does not resolve the case where the class-separating direction points opposite ways for different patients after recentering. Our 22-patient results generalize the 7-patient polarity finding: 45% of CHB-MIT patients show polarity inversion under minimal eigenvalue features, and 27% under Riemannian tangent space projection. This mechanism explains a previously reported binary detection failure: in a matched-regime pilot comparing QML and classical seizure detection, 7 of 17 QML patients produced zero detections, a pattern consistent with polarity inversion causing systematic misclassification rather than gradual performance degradation.

The feature-regime finding clarifies an important confound in prior comparisons: Riemannian and eigenvalue methods have typically been compared at different feature dimensionalities, with the Riemannian tangent space vectorization producing far fewer features than full spectral eigenvalue pipelines. In the matched low-dimensional regime (36 vs. 8 features), the two methods are statistically indistinguishable (Wilcoxon $p = 0.4358$). The practical advantage of high-dimensional eigenvalue methods is not geometric superiority. It is that feature redundancy across frequency bands effectively averages out per-patient polarity structure, incidentally achieving polarity invariance through dimensionality rather than design. Polarity invariance is a function of feature dimensionality, not method choice.

This framing also resolves an apparent paradox: why does the Riemannian method, which preserves more geometric information, perform worse than the eigenvalue method? The answer is that it preserves too much, specifically the directional information that produces polarity-induced cancellation across patients. A single sign bit of calibration per patient closes the gap between Riemannian and eigenvalue performance (0.659 vs. 0.628 calibrated AUC, $\Delta = 0.031$, $p = 0.4358$); the remaining difference is non-significant. The bottleneck is not information content but polarity resolution.

5.2 Quantum Circuits as Geometric Diagnostic Tools

The standard justification for quantum machine learning approaches is computational: quantum circuits can implement certain operations with fewer gates than classical circuits, potentially enabling speedup on specific problem structures [Biamonte et al., 2017]. Our results do not support a computational speedup claim, and we do not make one. The quantum circuit operates in the low-dimensional regime (8 features) where its calibrated AUC (0.637) is statistically indistinguishable from other low-dimensional classical methods (Riemannian 0.659 at 36 features, Tier 2A 0.628 at 8 features). Encoding the 336 features that achieve 0.952 AUC would require circuit depths beyond near-term hardware capability.

What the quantum circuit provides is a diagnostic capability not available from flexible classical classifiers. A classifier with learned parameters (XGBoost, LDA with sufficient feature dimension, or any model with adaptive decision boundaries) has enough degrees of freedom to accommodate polarity without exposing it: it simply learns different boundary orientations for different subgroups of the training data, absorbing the geometric asymmetry into its weights. The quantum circuit's fixed measurement basis removes those degrees of freedom entirely. The circuit functions as a geometric probe with a fixed orientation: it can only measure the component of the brain state signal that projects onto its measurement axis. When a patient's polarity is inverted relative to that axis, the circuit produces AUC < 0.50, making the inversion a numerically explicit, quantifiable event rather than a hidden variable silently corrected by model flexibility. This is not a limitation to be engineered around; it is the property that makes the circuit useful as an instrument. The architectural rigidity that prevents competitive classification performance is precisely what enables diagnostic sensitivity to manifold polarity.

The basis-dependence of polarity (71% concordance between quantum hardware and Riemannian, not 100%) indicates that PLV-based encoding and broadband covariance encoding probe different

projections of the same underlying manifold. Polarity is not a fixed patient label but a relationship between a patient's brain dynamics and a specific measurement projection, analogous to how the sign of a dot product depends on both the vector and the chosen basis.

Hardware results across multiple encoding configurations reveal that polarity is not a binary geometric invariant but a continuous projection quantity. Each patient's covariance shift vector has both a direction and a magnitude on the SPD manifold; the A-Gate circuit projects this vector onto its fixed measurement basis and reports the sign (AUC above or below 0.5) and magnitude (distance from 0.5) of the projection. Strongly polar patients (chb03, chb21) have shift vectors nearly aligned or anti-aligned with the measurement basis across all tested encodings. Weakly polar patients have shift vectors nearly orthogonal to the basis, producing small projections whose sign is sensitive to encoding parameters. This continuous interpretation explains both the encoding-dependent sign flips observed in 4/7 patients and the stable inversions in 2/7. The A-Gate circuit's fixed measurement basis makes the projection directly observable, functioning as a geometric probe rather than a classifier. Whether the projection structure is preserved across different measurement bases, and what the distribution of shift vector directions reveals about seizure pathophysiology, are questions for subsequent investigation.

5.3 The Simulation-Hardware Divergence

The degradation of statevector simulation with improved input data (0.460 at 1.95 s → 0.386 at 20 s) is, to our knowledge, an unreported phenomenon in the quantum ML literature. The standard assumption is that simulation is an optimistic upper bound on hardware performance: hardware shot noise and decoherence should degrade performance relative to ideal simulation [Cerezo et al., 2022]. Our results invert this assumption for a specific class of problem.

The mechanism is polarity-induced amplitude cancellation. When the aggregate quantum state is computed exactly across patients with opposing polarities, the contributions cancel to near-zero amplitude in the output qubit. This is a feature of the problem geometry, not a limitation of the circuit: the circuit correctly encodes each patient's data, but the aggregate measurement collapses across patients whose states interfere destructively. At 1.95-second windows, hardware shot noise breaks this cancellation stochastically, producing a substantial gap (0.524 raw vs. 0.460 simulation). At 20-second windows, however, stronger PLV estimates produce larger polarity projections that overwhelm the stochastic benefit: hardware also degrades (0.404 raw, 6/7 inverted), and the simulation-hardware gap narrows to 0.018. Shot noise attenuates polarity-induced cancellation but does not eliminate it when the projection magnitudes are large.

This finding has implications for QML evaluation practice. Simulation-based benchmarking is the norm for developing quantum ML circuits, both because hardware access is expensive and because simulation produces deterministic, reproducible results [Cerezo et al., 2022]. For problem domains where the training data contains geometric subgroups with opposing orientations, a pattern likely present in any cross-subject biological dataset, simulation may understate hardware performance when polarity projections are weak (short windows, noisy features). When projections are strong (long windows, well-

conditioned features), the gap narrows as polarity-induced cancellation overwhelms shot noise in both backends. The direction and magnitude of the simulation-hardware bias depends on the interaction between problem geometry and encoding parameters, and is not predictable without analyzing the manifold structure of the specific dataset.

5.4 Connection to Excitatory-Inhibitory Balance

The A-Gate's PN structure, with excitatory and inhibitory sub-circuits interconnected through ECR entangling gates, was motivated by the excitatory-inhibitory (E-I) balance model of seizure generation [van Hugte et al., 2023]. The ictal state is characterized by a shift in E-I balance toward excitation, disrupting the homeostatic regulation that maintains normal cortical dynamics [Turrigiano, 2008]. The polarity finding is consistent with this model: patients differ not only in whether their seizures shift E-I balance, but in the direction of the shift as measured by the specific features encoded.

The PLV-based encoding captures phase coupling between channels. Some patients may show increased phase synchrony during seizure (standard polarity: ictal states cluster toward high PLV, rotating the quantum state in one direction); others may show decreased inter-regional phase coupling as the seizure localizes (inverted polarity) [Wendling et al., 2003; Jiruska et al., 2013]. The Riemannian tangent space captures broadband covariance structure, which may reflect different aspects of the same E-I dynamics. The encoding-dependent polarity observed in 4/7 hardware patients (sign flips between window sizes or channel counts) is consistent with this model: different projections of the covariance shift vector onto different measurement bases capture different aspects of the E-I disruption. The strongly polar patients (chb03, chb21, always inverted across all configurations) may represent cases where the E-I shift is large and aligned with multiple measurement axes simultaneously; weakly polar patients may have smaller or more obliquely oriented shifts.

6. Limitations

Several limitations constrain the scope of our conclusions. First, quantum hardware experiments cover only 7 of 22 CHB-MIT patients due to computational cost constraints. The feature-regime finding rests on the full 22-patient cohort for classical methods only; quantum hardware coverage of the full cohort is infeasible on near-term hardware. Preliminary channel-scaling experiments (12 and 16 channels) show degrading hardware AUC as circuit depth increases, with 16-channel circuits (246 CZ gates, depth 192) returning chance-level performance, indicating that the coherence budget of current hardware sets a practical ceiling on encoding complexity.

Second, the feature-regime argument could be interpreted as 'quantum just needs more features to compete.' This is partially correct. Encoding 336 features would require a quantum circuit whose depth exceeds current hardware coherence limits by a substantial margin. The low-dimensional regime (8 features for the quantum circuit) is not a compromise; it is the only regime accessible to near-term quantum hardware, and in that regime the methods are equivalent. This is not a dismissal of quantum potential; it is an honest characterization of current capability [Preskill, 2018].

Third, we tested a single quantum encoding (PLV theta-alpha, CH8 configuration). The V2 band-power encoding produced degenerate hardware results (AUC = 0.500), indicating that not all encodings expose the polarity structure. Whether other encodings provide better performance or cleaner polarity resolution is unknown. Fourth, no error mitigation was applied to quantum circuits; dynamical decoupling or zero-noise extrapolation might improve raw AUC but would not address the fundamental polarity problem. Fifth, the polarity calibration procedure requires a labeled seizure sample per patient; for zero-shot cross-patient deployment, Tier 2B eigenvalue features with MAX aggregation remain the practical recommendation. Sixth, the oracle-calibrated AUC figures reported throughout this paper assume known per-patient polarity, equivalent to one labeled seizure sample per patient. Label-free polarity determination using population-level covariance geometry (Approach B) fails for direction-preserving classifiers because the population eigenvalue direction is always “standard”; polarity is a per-patient deviation from this population direction. Only polarity-invariant methods (Tier 2A sorted eigenvalues) produce clean-validated calibration (Approach B AUC 0.683 vs. oracle 0.689). For direction-preserving methods, the oracle-calibrated figures represent upper bounds on achievable performance. Finally, the Mandelbrot boundary distance analysis initially proposed as a convergent E-I measure was inconclusive due to dataset mismatches and is reserved for future work.

7. Conclusion

We have identified patient-specific manifold polarity as the geometric mechanism underlying the well-documented failure of Riemannian methods in cross-patient EEG classification, and we have validated this mechanism across 22 CHB-MIT patients. Specifically, we show that seizure-related covariance shifts point in opposite directions on the SPD manifold for different patients, a structure invisible to any classifier that normalizes or discards directional information. A previously unreported feature-regime interaction governs method comparisons: in the low-dimensional regime (8–36 features), geometry-preserving methods (Riemannian tangent space projection, quantum circuit) are statistically equivalent to geometry-discarding methods (sorted eigenvalues) after polarity calibration. In the full-feature regime (336 features), eigenvalue-based methods achieve 0.952 AUC with zero polarity inversion, not because the polarity problem is solved, but because feature dimensionality renders it irrelevant. Prior comparisons in the literature that favor eigenvalue over Riemannian methods are likely confounded by this dimensionality asymmetry.

The A-Gate quantum circuit serves as a geometric diagnostic tool that makes polarity an explicit, measurable observable. Its fixed measurement basis, which prevents the adaptive absorption of polarity that flexible classical classifiers perform silently, is precisely what revealed the phenomenon; the circuit’s architectural rigidity functions as a quantum geometric probe rather than a limitation. The simulation-hardware divergence is a novel finding with implications for quantum ML evaluation methodology. At short windows, hardware substantially outperforms simulation (0.524 vs. 0.460 raw AUC) because shot noise attenuates polarity-induced cancellation. At longer windows, both backends degrade as stronger polarity projections overwhelm the stochastic benefit, but the gap persists (0.404 vs. 0.386). Simulation-based benchmarking may systematically understate hardware performance for cross-subject biological

datasets with opposing geometric subgroups, particularly in the weak-projection regime where shot noise provides the greatest relative benefit.

Hardware experiments across multiple encoding configurations reveal that manifold polarity is not a binary invariant but a continuous projection: the sign and magnitude of a patient's covariance shift vector projected onto the circuit's measurement basis. Strongly polar patients show consistent sign across all encodings; weakly polar patients show encoding-dependent flips, providing direct evidence of the vector nature of the covariance shift. Whether the distribution of shift vector directions carries clinical information, and whether fixed-basis quantum circuits with different measurement axes can map the full shift geometry, are the questions this result opens.

More broadly, this work suggests that the inter-subject variability invoked to explain cross-patient failures in BCI is not a monolithic phenomenon. At least one component, manifold polarity, has a specific geometric cause, a specific geometric fix (a single bit of calibration per patient), and a quantum diagnostic that makes it numerically explicit. Identifying similar structured causes for other components of inter-subject variability, and building classifiers that are invariant to them by design rather than by accident of dimensionality, is a productive direction for cross-patient generalization research.

Declarations

Conflict of interest:

The author declares no competing interests.

Ethics and Consent to Participate

Not applicable.

Consent to Publish:

Not applicable.

Funding:

No funding was received for this research.

Author Contribution

J.A. conceived the study, designed and executed all experiments, performed all quantum hardware runs on IBM ibm_torino, conducted the classical and statistical analyses, and wrote the manuscript.

Data Availability

The CHB-MIT Scalp EEG Database used in this study is publicly available at PhysioNet (https://physionet.org/content/chbmit/1.0.0/). Quantum hardware results and per-patient analysis outputs are available in the project repository at https://github.com/jappenzeller/QDNU.

References

1. Appenzeller J. Hardware-Validated Quantum Positive-Negative Neuron Architecture for Multi-Channel EEG Analysis. Discover Quantum Sci. 2026;2:30. 10.1007/s44464-026-00032-w.
2. Barachant A, Temko A, Li F, Titericz G. (2017). Melbourne University AES/MathWorks/NIH Seizure Prediction Challenge: Winning solution documentation. https://github.com/alexandrebarachant/kaggle-seizure-prediction-challenge-2016 [20-second non-overlapping windows, MAX aggregation, no filtering; ensemble of XGBoost models with spectral, correlation, and Riemannian tangent space features].
3. Barachant A, Bonnet S, Congedo M, Jutten C. Multiclass brain-computer interface classification by Riemannian geometry. IEEE Trans Biomed Eng. 2012;59(4):920–8. 10.1109/TBME.2011.2172210.
4. Barachant A, Barthélemy Q, King J-R, Gramfort A, Chevallier S, Rodrigues PLC et al. (2024). pyRiemann (v0.7). Zenodo. 10.5281/zenodo.593816
5. Biamonte J, Wittek P, Pancotti N, Rebentrost P, Wiebe N, Lloyd S. Quantum machine learning. Nature. 2017;549:195–202. 10.1038/nature23474.
6. Blankertz B, Tomioka R, Lemm S, Kawanabe M, Müller K-R. Optimizing spatial filters for robust EEG single-trial analysis. IEEE Signal Process Mag. 2008;25(1):41–56. 10.1109/MSP.2008.4408441.
7. Cerezo M, Arrasmith A, Babbush R, Benjamin SC, Endo S, Fujii K, McClean JR, Mitarai K, Yuan X, Cincio L, Coles PJ. Variational quantum algorithms. Nat Reviews Phys. 2021;3(9):625–44. 10.1038/s42254-021-00348-9.
8. Cerezo M, Verdon G, Huang H-Y, Cincio L, Coles PJ. Challenges and opportunities in quantum machine learning. Nat Comput Sci. 2022;2(9):567–76. 10.1038/s43588-022-00311-3.
9. Chen T, Guestrin C. (2016). XGBoost: A scalable tree boosting system. Proc. KDD, 785–794. 10.1145/2939672.2939785
10. Congedo M, Barachant A, Bhatia R. Riemannian geometry for EEG-based brain-computer interfaces; a primer and a review. Brain-Computer Interfaces. 2017;4(3):155–74. 10.1080/2326263X.2017.1297192.
11. Fawcett T. An introduction to ROC analysis. Pattern Recognit Lett. 2006;27(8):861–74. 10.1016/j.patrec.2005.10.010.
12. Fletcher PT, Joshi S. (2004). Principal geodesic analysis on symmetric spaces: Statistics of diffusion tensors. ECCV Workshop CVAMIA, LNCS 3117, 87–98.

13. Gupta MM, Jin L, Homma N. Static and Dynamic Neural Networks: From Fundamentals to Advanced Theory. Wiley-IEEE; 2003. [PN neuron architecture, Section 8.3, Fig. 8.22].
14. Hills M. (2014). UPenn and Mayo Clinic's Seizure Detection Challenge: Winning submission. 1st place, Kaggle. https://github.com/MichaelHills/seizure-detection [log10 FFT magnitudes 1–47 Hz, inter-channel correlation eigenvalues in frequency and time domains, Random Forest classifier].
15. IBM Quantum. (2024). IBM Heron r1 technical specifications. IBM Res. https://quantum.ibm.com
16. Lachaux J-P, Rodriguez E, Martinerie J, Varela FJ. Measuring phase synchrony in brain signals. Hum Brain Mapp. 1999;8(4):194–208. 10.1002/(SICI)1097-0193(1999)8:4%3C194::AID-HBM4%3E3.0.CO;2-C.
17. Ledoit O, Wolf M. A well-conditioned estimator for large-dimensional covariance matrices. J Multivar Anal. 2004;88(2):365–411. 10.1016/S0047-259X(03)00096-4.
18. Wendling F, Bartolomei F, Bellanger JJ, Chauvel P. Epileptic fast intracerebral EEG activity: evidence for spatial decorrelation at seizure onset. Brain. 2003;126(6):1449–59. 10.1093/brain/awg144.
19. Jiruska P, de Curtis M, Jefferys JGR, Schevon CA, Schiff SJ, Schindler K. Synchronization and desynchronization in epilepsy: controversies and hypotheses. J Physiol. 2013;591(4):787–97. 10.1113/jphysiol.2012.239590.
20. Pennec X, Fillard P, Ayache N. A Riemannian framework for tensor computing. Int J Comput Vis. 2006;66(1):41–66. 10.1007/s11263-005-3222-z.
21. Preskill J. Quantum computing in the NISQ era and beyond. Quantum. 2018;2:79. 10.22331/q-2018-08-06-79.
22. Qiskit Contributors. (2023). Qiskit: An open-source framework for quantum computing. Zenodo. 10.5281/zenodo.2573505
23. Rodrigues PLC, Jutten C, Congedo M. Riemannian Procrustes analysis: transfer learning for brain-computer interfaces. IEEE Trans Biomed Eng. 2019;66(8):2390–401. 10.1109/TBME.2018.2889705.
24. Shoeb AH. (2009). Application of machine learning to epileptic seizure onset detection and treatment. PhD thesis, MIT.
25. Temko A, Thomas E, Marnane W, Lightbody G, Boylan G. EEG-based neonatal seizure detection with support vector machines. Clin Neurophysiol. 2011;122(3):464–73. 10.1016/j.clinph.2010.06.034.
26. Turrigiano GG. The self-tuning neuron: synaptic scaling of excitatory synapses. Cell. 2008;135(3):422–35. 10.1016/j.cell.2008.10.008.
27. Usman SM, Usman M, Fong S. (2021). Epileptic seizures prediction using machine learning methods. Comput. Math. Methods Med., 2021, 1–10. 10.1155/2021/4132649
28. van Hugte EJH, Schubert D, Nadif Kasri N. Excitatory/inhibitory balance in epilepsies and neurodevelopmental disorders: depolarizing γ-aminobutyric acid as a common mechanism. Epilepsia. 2023;64(8):1975–90. 10.1111/epi.17651.
29. WHO. Epilepsy: A public health imperative. World Health Organization; 2019.

30. Wu J, Zhou T, Li T. Detecting epileptic seizures in EEG signals with complementary ensemble empirical mode decomposition and extreme gradient boosting. Entropy. 2020;22(2):140. 10.3390/e22020140.
31. Yair O, Ben-Chen M, Talmon R. Parallel transport on the cone manifold of SPD matrices for domain adaptation. IEEE Trans Signal Process. 2019;67(7):1797–811. 10.1109/TSP.2019.2894801.
32. Yger F, Berar M, Lotte F. Riemannian approaches in brain-computer interfaces: A review. IEEE Trans Neural Syst Rehabil Eng. 2017;25(10):1753–62. 10.1109/TNSRE.2016.2627016.
33. Zanini P, Congedo M, Jutten C, Said S, Berthoumieu Y. Transfer learning: a Riemannian geometry framework with applications to brain-computer interfaces. IEEE Trans Biomed Eng. 2018;65(5):1107–16. 10.1109/TBME.2017.2742541.