

RELAY: Risk-Calibrated Human Neural Trajectory Distillation for Visual Reasoning

Supplementary Material

Planning-only numerical scenario — not empirical findings. All experimental outcomes, statistical tests, illustrations, audit values, and resource measurements in this edition are hypothetical design targets, not results from EEG recordings or trained models. Behavioral values come from a prespecified hierarchical probability model with item, family, initialization, and method interactions; they are neither empirical estimates nor forecasts. Dataset facts are attributed to the public release. Mathematical results are stated with their assumptions. Replace every experimental value and the associated interpretation before submission. This marked edition must not be submitted as an experimental paper.

Reading map. Sections [A–I](#) give the statistical and mathematical foundations. Sections [J–K](#) establish the revised path-compatible theorem, independent audits, and exact control identities. Sections [L–T](#) specify the complete experimental protocol, extended results, visual atlas, and record schemas. All entries below and PDF bookmarks are clickable.

Contents

A . Formal Setup and the Object Being Transferred	2
B . Unbiased Estimation and Exact Sampling Variance	3
C . Projection onto a Valid Trajectory-Kernel Set	4
D . Risk Calibration after Temporal Selection	4
E . Globally Coupled Monotone Alignment	6
F . What Joint Trajectory Geometry Identifies	7
G . Prediction Stability and Student Capacity	8
H . Information Budgets and Additional Sensors	9
I . Implementation and Numerical Verification	10
J . Path-Compatible Calibration and Independent Risk Audits	11
K . Mechanism Controls with Exact Invariances	12
L . Dataset, Events, and Evaluation Units	13
M . Signal Processing and Feature Construction	14
N . Models, Optimization, and Exact Loss Evaluation	15
O . Baselines as Explicit Objectives	16
P . Statistical Estimands, Intervals, and Multiplicity	17
Q . Sensor Budgets and Decision-State Interventions	18
R . Extended Results and Calibration Diagnostics	20
S . Visual Atlas: Kernels, Decisions, and Interventions	26
T . Computation, Records, and Reproduction	28

A. Formal Setup and the Object Being Transferred

A.1. Conditioning, notation, and units

Sections A–I establish the fixed-projector statistical core. Section J applies it to the selected path-compatible projector after conditioning on the independent selector. Throughout the analysis, condition on the distinct training images, their task labels, the fixed feature-extraction specification, the nuisance design, and the task-only warm start. None of these quantities is fitted using the donor EEG involved in a subsequent theorem. Donor-specific preprocessing is permitted when the same prespecified rule is fitted independently within each donor’s training records. A pooled preprocessing fit across the selector, calibrator, and teacher would invalidate this separation.

For n training problems and T neural windows, set $N = nT$. Time-major stacking means row $(t - 1)n + i$ refers to problem i in window t . The neural matrix for donor s is $Z_s \in \mathbb{R}^{N \times p}$. The feature dimension p is not the number of independent observations. The independent sampling units in the teacher analysis are donors. We use

$$\langle A, B \rangle_N = \frac{\text{tr}(A^\top B)}{N^2}, \quad \|A\|_N^2 = \langle A, A \rangle_N. \quad (\text{S1})$$

For an $n \times n$ time block, the analogous normalization is by n^2 . A global kernel loss is therefore the average of its T^2 block losses.

Table S1. Notation. Dimensions, operators, and target quantities.

Symbol	Meaning
n, T, N	Distinct paired problems, windows, and $N = nT$
p, d, L	Neural width, student-state width, reasoning steps
Z_s, M, E_s	Donor features, population mean, donor deviation
$P, \mathcal{P}$	Item nuisance projection and $I_T \otimes P$
K, B	Population trajectory kernel and unbiased estimate
m_j, ζ_j	Raw pair moments and variance components
π, G_π^0	Ordered path and warm-start trajectory kernel
a, v_b, ρ	Warm-start discrepancy, teacher risk, transfer weight
$\mathcal{A}, \mathcal{C}, \mathcal{B}$	Selector, calibrator, teacher cohorts
$\tilde{K}$	Frozen, projected distillation target

A.2. The observation model

Assumption 1 (Independent, identically processed donors). Conditional on the fixed training design, donor matrices $Z_1, \dots, Z_S$ are independent copies of a random

matrix Z , with $\mathbb{E}Z = M$. Write $E_s = Z_s - M$. Dependence between any two entries of a single E_s is allowed. The row-bounding transform ensures $\|Z_s[i, :]\|_2 \leq 1$ almost surely.

The assumption concerns processed measurements under the fixed response-censoring and invalid-feature map. Eligible donors define the observation population, and sensor coordinates remain in the common BioSemi feature frame.

Assumption 2 (Fixed common coordinates and projection). $P = P^\top = P^2 \in \mathbb{R}^{n \times n}$ is fixed before donor-dependent target selection, and $\mathcal{P} = I_T \otimes P$. All donors use the same feature coordinates and window construction. Each warm-start row and each student row before projection has norm at most one.

A shared orthogonal change of the neural feature coordinates leaves every kernel unchanged. A different rotation for each donor does not generally leave the population mean kernel unchanged. The estimator therefore retains one shared feature coordinate frame across donors.

Lemma S1 (Feasibility and norm bounds). With $U = \mathcal{P}M$, $K = UU^\top$, and $X_s = \mathcal{P}Z_s$, the matrices K and G_π^0 belong to

$$\mathcal{C} = \{A = A^\top \succeq 0 : \mathcal{P}A\mathcal{P} = A, \text{tr } A \leq N\}. \quad (\text{S2})$$

Moreover, $\|K\|_N \leq 1$, $\|G_\pi^0\|_N \leq 1$, and $\|\text{sym}(X_s X_r^\top)\|_N \leq 1$.

Proof. Jensen’s inequality gives $\|M[i, :]\|_2 \leq \mathbb{E}\|Z[i, :]\|_2 \leq 1$, hence $\|M\|_F^2 \leq N$. Orthogonal projection is a contraction, so $\text{tr } K = \|\mathcal{P}M\|_F^2 \leq N$. The range and PSD conditions follow from the factorization. The same reasoning applies to the time-stacked bounded warm-start states, including repetitions induced by π . For a PSD matrix, the Frobenius norm is at most the trace, proving the first two norm bounds. Finally,

$$\begin{aligned} \|\text{sym}(X_s X_r^\top)\|_F &\leq \|X_s X_r^\top\|_F \\ &\leq \|X_s\|_F \|X_r\|_F \leq N. \end{aligned}$$

Dividing by N proves the claim. $\square$

A.3. Why a joint kernel is the target

The population target is a block matrix $K_{tu} = P M_t M_u^\top P$. A diagonal block is a within-window geometry. An off-diagonal block compares states across both different problems and different windows. These are different statistical objects. Replacing the joint kernel by $\text{diag}(K_{11}, \dots, K_{TT})$ discards all cross-time coordinate information, even when every retained block is estimated perfectly. Section F gives an explicit counterexample and a positive identification result.

B. Unbiased Estimation and Exact Sampling Variance

B.1. The U-statistic identity

Define $h_{sr} = \frac{1}{2}(X_s X_r^\top + X_r X_s^\top)$ for $s \neq r$ and $B = \binom{S}{2}^{-1} \sum_{s < r} h_{sr}$.

Proposition S2 (Unbiasedness and the self-product bias). *For $S \geq 2$, $\mathbb{E}B = K$. If $\bar{Z} = S^{-1} \sum_s Z_s$, then*

$$\mathbb{E}[\mathcal{P} \bar{Z} \bar{Z}^\top \mathcal{P}] = K + \frac{1}{S} \mathcal{P} \mathbb{E}[E_s E_s^\top] \mathcal{P}. \quad (\text{S3})$$

The second term is PSD and need not be a multiple of the identity.

Proof. For distinct donors, independence gives $\mathbb{E}[Z_s Z_r^\top] = MM^\top$, because both linear noise terms and $\mathbb{E}[E_s E_r^\top]$ vanish. Symmetrization and projection therefore preserve expectation K . Every term in the average has that expectation. For the second assertion, write $\bar{Z} = M + \bar{E}$. The cross terms have zero expectation. Expanding $\mathbb{E}[\bar{E} \bar{E}^\top]$ produces S equal self-products divided by S^2 ; all $S(S-1)$ cross-products vanish. This proves Eq. (S3). A covariance-like matrix $\mathbb{E}[E_s E_s^\top]$ is PSD, and congruence by $\mathcal{P}$ preserves PSD. $\square$

The computational identity follows by expanding $(\sum_s X_s)(\sum_s X_s)^\top$. Its self-products cancel against $\sum_s X_s X_s^\top$, leaving $\sum_{s < r} (X_s X_r^\top + X_r X_s^\top)$. Division by $S(S-1)$ equals division of the symmetric pair sum by $\binom{S}{2}$.

B.2. A direct orthogonal decomposition

Write $V_s = \mathcal{P} E_s$ and $U = \mathcal{P} M$. Define

$$A_s = UV_s^\top + V_s U^\top, \quad C_{sr} = \text{sym}(V_s V_r^\top). \quad (\text{S4})$$

Then

$$h_{sr} = K + \frac{1}{2}(A_s + A_r) + C_{sr}, \quad (\text{S5})$$

and summing pairs gives the exact decomposition

$$B - K = \frac{1}{S} \sum_s A_s + \binom{S}{2}^{-1} \sum_{s < r} C_{sr}. \quad (\text{S6})$$

Both components have mean zero. More importantly, C_{sr} is conditionally mean zero given either of its donors:

$$\mathbb{E}[C_{sr} | V_s] = 0, \quad \mathbb{E}[C_{sr} | V_r] = 0. \quad (\text{S7})$$

Theorem S3 (Exact teacher mean-squared error). *Let $\alpha = \mathbb{E}\|A_1\|_N^2$ and $\beta = \mathbb{E}\|C_{12}\|_N^2$. Then*

$$\mathbb{E}\|B - K\|_N^2 = \frac{\alpha}{S} + \frac{2\beta}{S(S-1)}. \quad (\text{S8})$$

Equivalently, with $m_0 = \mathbb{E}\|h_{12}\|_N^2$, $m_1 = \mathbb{E}\langle h_{12}, h_{13} \rangle_N$, and $m_2 = \mathbb{E}\langle h_{12}, h_{34} \rangle_N$,

$$\begin{aligned} m_2 &= \|K\|_N^2, \\ \zeta_1 &= m_1 - m_2 = \alpha/4, \\ \zeta_2 &= m_0 - 2m_1 + m_2 = \beta. \end{aligned} \quad (\text{S9})$$

Thus $v_S = 4\zeta_1/S + 2\zeta_2/[S(S-1)]$.

Proof. For $s \neq r$, independence and zero means imply $\mathbb{E}\langle A_s, A_r \rangle_N = 0$. A term A_s is orthogonal in expectation to every C_{uv} . If $s \notin \{u, v\}$, independence suffices; if $s = u$ or $s = v$, condition on the shared donor and use Eq. (S7). Likewise, two distinct C terms are orthogonal in expectation. Disjoint pairs are independent. Pairs sharing one donor are conditionally independent given that donor, with conditional mean zero for each term. Expanding the squared norm of Eq. (S6) therefore leaves $S\alpha/S^2$ and $\binom{S}{2}\beta/\binom{S}{2}^2$, proving Eq. (S8).

For the moment identities, disjoint pairs h_{12} and h_{34} are independent and both have mean K , so $m_2 = \|K\|_N^2$. Expanding m_1 using the pair decomposition, only the common $A_1/2$ term contributes beyond $\|K\|_N^2$; hence $m_1 = m_2 + \alpha/4$. Expanding m_0 leaves two A terms with factor $1/4$ and one C term, so $m_0 = m_2 + \alpha/2 + \beta$. Substitution gives Eq. (S9). Nonnegativity follows from the definitions of α and β . $\square$

The proof permits dependence across time bins, image pairs, and EEG channels within a donor. Its independent sample size is the number of donors.

Corollary S4 (Independent donors monotonically reduce teacher variance). *For $S \geq 2$,*

$$v_S - v_{S+1} = \frac{4\zeta_1}{S(S+1)} + \frac{4\zeta_2}{S(S-1)(S+1)} \geq 0. \quad (\text{S10})$$

Proof. Subtract the two expressions from Theorem S3; both coefficients and variance components are nonnegative. $\square$

The result quantifies independent repetitions of the same measurement design. Additional sensors are analyzed through conditional innovations in Section H.

B.3. Exact finite-sample moment estimators

Let $(S)_r = S(S-1)\cdots(S-r+1)$. An explicit implementation is

$$\hat{m}_0 = \frac{1}{(S)_2} \sum_{a \neq b} \|h_{ab}\|_N^2, \quad (\text{S11})$$

$$\hat{m}_1 = \frac{1}{(S)_3} \sum_{a,b,c \text{ distinct}} \langle h_{ab}, h_{ac} \rangle_N, \quad (\text{S12})$$

$$\hat{m}_2 = \frac{1}{(S)_4} \sum_{a,b,c,d \text{ distinct}} \langle h_{ab}, h_{cd} \rangle_N. \quad (\text{S13})$$

Each summand has the appropriate population expectation. An implementation may instead average over unordered pairs, all edge pairs sharing exactly one donor, and all disjoint edge pairs. Those multiplicities cancel and produce exactly the same three estimates. Self-pairs must never appear in $\hat{m}_2$.

Lemma S1 implies $0 \leq m_0 \leq 1$, $0 \leq m_2 \leq 1$, and $0 \leq \zeta_1, \zeta_2 \leq 1$. For ζ_2 , use $m_1 \geq m_2$ to obtain $\zeta_2 \leq m_0 - m_2 \leq 1$. We clip estimated variance components to $[0, 1]$ when constructing a practical variance estimate. The *unclipped* quantities in Eq. (S13) are unbiased; clipping is a stabilization step and is logged separately.

C. Projection onto a Valid Trajectory-Kernel Set

C.1. Existence and a closed form

The set $\mathcal{C}$ is nonempty, closed, and convex. It is bounded because a PSD matrix with trace at most N has Frobenius norm at most N . Hence every matrix has a unique Frobenius projection onto $\mathcal{C}$.

Let $R \in \mathbb{R}^{N \times r}$ have orthonormal columns spanning the range of $\mathcal{P}$. For an arbitrary input A , symmetrize it and form $A_R = R^\top \text{sym}(A)R$. Let $A_R = V \text{diag}(\lambda)V^\top$.

Proposition S5 (Spectral trace projection). *The projection is*

$$\Pi_{\mathcal{C}}(A) = RV \text{diag}(w)V^\top R^\top, \quad w_j = (\lambda_j - \tau)_+, \quad (\text{S14})$$

where $\tau = 0$ if $\sum_j (\lambda_j)_+ \leq N$; otherwise $\tau > 0$ is the unique threshold satisfying $\sum_j (\lambda_j - \tau)_+ = N$.

Proof. Every feasible matrix has the form $RHR^\top$ with $H \succeq 0$ and $\text{tr} H \leq N$. Orthogonal decomposition of the Frobenius norm shows that the part of A outside this range and its skew-symmetric part add constants independent of H . We must therefore minimize $\|H - A_R\|_F^2$.

Write $D = V^\top HV$. Frobenius invariance gives

$$\|D - \text{diag}(\lambda)\|_F^2 = \sum_j (D_{jj} - \lambda_j)^2 + \sum_{i \neq j} D_{ij}^2. \quad (\text{S15})$$

If $D \succeq 0$, its diagonal entries are nonnegative. Replacing D by its diagonal preserves feasibility and cannot increase the objective. Thus solve

$$\min_{w_j \geq 0, \sum_j w_j \leq N} \frac{1}{2} \sum_j (w_j - \lambda_j)^2. \quad (\text{S16})$$

The KKT conditions give $w_j - \lambda_j + \tau - \mu_j = 0$, $\mu_j w_j = 0$, $\tau(\sum_j w_j - N) = 0$, and nonnegative multipliers. Therefore $w_j = (\lambda_j - \tau)_+$. The continuous, decreasing threshold equation determines τ whenever the trace constraint is active. Strict convexity gives uniqueness of the projected matrix. $\square$

C.2. Why projection cannot increase the target error

Lemma S6 (Projection inequality). *For $K \in \mathcal{C}$ and $\tilde{K} = \Pi_{\mathcal{C}}(A)$,*

$$\|\tilde{K} - K\|_N^2 \leq \|A - K\|_N^2 - \|A - \tilde{K}\|_N^2 \leq \|A - K\|_N^2. \quad (\text{S17})$$

Proof. For $0 \leq t \leq 1$, $\tilde{K} + t(K - \tilde{K})$ is feasible. The right derivative at $t = 0$ of its squared distance to A is nonnegative, yielding $\langle A - \tilde{K}, K - \tilde{K} \rangle_N \leq 0$. Expand

$$\begin{aligned} \|A - K\|_N^2 &= \|A - \tilde{K}\|_N^2 + \|\tilde{K} - K\|_N^2 \\ &\quad + 2\langle A - \tilde{K}, \tilde{K} - K \rangle_N. \end{aligned}$$

The last term is nonnegative. Rearrangement proves the result. $\square$

The PSD step is performed *after* estimating moments. Computing moments from already projected pair kernels would change their expectations and invalidate the preceding unbiasedness statements.

D. Risk Calibration after Temporal Selection

D.1. The exact bias–variance calculation

Let $\mathcal{F}_{\mathcal{A}, \mathcal{C}}$ denote information in the selector and calibrator cohorts. Conditional on this information, π , G_π^0 , and ρ are fixed, while $B^{\mathcal{B}}$ remains an independent unbiased estimator of K . Set $a = \|G_\pi^0 - K\|_N^2$ and $v = \mathbb{E}\|B^{\mathcal{B}} - K\|_N^2$.

Theorem S7 (Conditional risk identity and oracle comparison). *For $A_\rho = (1 - \rho)G_\pi^0 + \rho B^{\mathcal{B}}$,*

$$\mathbb{E}[\|A_\rho - K\|_N^2 \mid \mathcal{F}_{\mathcal{A}, \mathcal{C}}] = (1 - \rho)^2 a + \rho^2 v. \quad (\text{S18})$$

If $a + v > 0$, the unique minimizing weight is $\rho^* = a/(a + v)$ and

$$R(\rho) = \frac{av}{a + v} + (a + v)(\rho - \rho^*)^2. \quad (\text{S19})$$

Projection onto $\mathcal{C}$ can only reduce this risk. At ρ^* , it is at most $\min(a, v)$.

Proof. Expand $A_\rho - K = (1 - \rho)(G_\pi^0 - K) + \rho(B^\mathcal{B} - K)$. The mixed inner product has conditional expectation zero by independence and unbiasedness. This proves Eq. (S18). Expanding the right side yields $a - 2a\rho + (a + v)\rho^2$. Completing the square gives Eq. (S19); the minimizer lies in $[0, 1]$. Since $av/(a + v) \leq a$ and $av/(a + v) \leq v$, the oracle comparison follows. Lemma S6 proves the projected bound. If $a = v = 0$, both sources equal K almost surely and every weight has zero risk. $\square$

This identity also explains the anchor control. At $\rho = 0$, the objective still regularizes the student's trajectory toward its task-only warm start. Removing EEG is not the same intervention as removing all regularization.

D.2. Why the calibrator must be independent of the selector

For a path fixed independently of cohort $\mathcal{C}$, define

$$\hat{a}_{\text{raw}} = \|G_\pi^0\|_N^2 - 2\langle G_\pi^0, B^\mathcal{C} \rangle_N + \hat{m}_2^\mathcal{C}. \quad (\text{S20})$$

Then $\mathbb{E}[\hat{a}_{\text{raw}} \mid \pi] = a$. Clipping to $[0, 4]$ cannot increase pointwise estimation error because a lies in that interval. The same argument applies to clipping each $\hat{\zeta}_j$ to $[0, 1]$.

If π is selected on $\mathcal{C}$ itself, G_π^0 becomes dependent on $B^\mathcal{C}$. The displayed expectation no longer follows: the selected cross term is data-adaptive. Separating the two cohorts lets us condition on the selected path and apply an ordinary fixed-function concentration argument. This is the statistical reason for three roles rather than a nominal train/test split within one noisy target.

D.3. A self-contained concentration lemma

Lemma S8 (Bounded U-statistic concentration). *Let U_S be a complete U-statistic of order r , formed from independent observations with a symmetric kernel taking values in $[a_0, b_0]$. Let $q = \lfloor S/r \rfloor \geq 1$. Then*

$$\mathbb{P}(|U_S - \mathbb{E}U_S| \geq \epsilon) \leq 2 \exp\left[-\frac{2q\epsilon^2}{(b_0 - a_0)^2}\right]. \quad (\text{S21})$$

Proof. For any permutation of the S observations, split its first qr entries into q disjoint blocks of size r and average the kernel over those blocks. Denote this block average by V_σ . Averaging V_σ over all permutations gives U_S : by symmetry, every unordered r -subset appears equally often, including when S is not divisible by r .

For completeness, if a bounded centered random variable W lies in an interval of length $b_0 - a_0$, its log moment generating function $g(t) = \log \mathbb{E}e^{tW}$ has $g(0) = g'(0) = 0$.

Its second derivative is the variance under an exponentially tilted distribution, which is at most $(b_0 - a_0)^2/4$. Integrating twice gives $g(t) \leq t^2(b_0 - a_0)^2/8$.

The q block-kernel values are independent. Applying this bound to their average yields

$$\mathbb{E}e^{t(V_\sigma - \mathbb{E}U_S)} \leq \exp\left[\frac{t^2(b_0 - a_0)^2}{8q}\right].$$

Convexity of the exponential and averaging over permutations give the same bound for U_S . Markov's inequality then bounds the upper tail by $\exp[-t\epsilon + t^2(b_0 - a_0)^2/(8q)]$. Minimize over $t > 0$ to obtain the exponent in Eq. (S21). Apply the same argument to $-U_S$ and add the two tails. $\square$

D.4. Finite-sample calibration error after path selection

Conditional on the selector, apply Lemma S8 to four fixed functions of the calibration cohort: $u = \langle G_\pi^0, B^\mathcal{C} \rangle_N$, and $\hat{m}_0, \hat{m}_1, \hat{m}_2$. Their orders are 2, 2, 3, 4, and their kernel ranges have lengths at most 2, 1, 2, 2, respectively. For a desired failure probability δ , define

$$\begin{aligned} e_u &= 2\sqrt{\frac{\log(8/\delta)}{2\lfloor c/2 \rfloor}}, & e_0 &= \sqrt{\frac{\log(8/\delta)}{2\lfloor c/2 \rfloor}}, \\ e_1 &= 2\sqrt{\frac{\log(8/\delta)}{2\lfloor c/3 \rfloor}}, & e_2 &= 2\sqrt{\frac{\log(8/\delta)}{2\lfloor c/4 \rfloor}}, \end{aligned} \quad (\text{S22})$$

where $c = |\mathcal{C}|$. Symmetrize the triple and quadruple kernels by averaging their possible pairings before applying the lemma; averaging preserves these ranges. A union bound gives simultaneous coverage at least $1 - \delta$.

Because clipping toward an interval containing the truth is nonexpansive,

$$|\hat{a} - a| \leq e_a := 2e_u + e_2, \quad (\text{S23})$$

$$|\hat{v}_b - v_b| \leq e_v := \frac{4}{b}(e_1 + e_2) + \frac{2}{b(b-1)}(e_0 + 2e_1 + e_2). \quad (\text{S24})$$

No union over the selected paths is needed: the entire calibration cohort is independent of the path-selection operation.

Theorem S9 (Post-selection plug-in risk bound). *On the event in Eq. (S24), put $s = a + v > 0$ and $e = e_a + e_v$. If $e \leq s/2$, then*

$$\begin{aligned} |\hat{\rho} - \rho^*| &\leq \frac{2e}{s}, \\ \mathbb{E}[\|\tilde{K} - K\|_N^2 \mid \mathcal{F}_{\mathcal{A}, \mathcal{C}}] &\leq \frac{av}{a+v} + \frac{4e^2}{a+v}. \end{aligned} \quad (\text{S25})$$

Proof. Let $\delta_a = \hat{a} - a$ and $\delta_v = \hat{v} - v$. Direct subtraction gives

$$\hat{\rho} - \rho^* = \frac{v\delta_a - a\delta_v}{(a+v)(\hat{a}+\hat{v})}. \quad (\text{S26})$$

The numerator has magnitude at most $ve_a + ae_v \leq se$. The estimated denominator sum is at least $s - e \geq s/2$, so the ratio is bounded by $2e/s$. Substitute this into Eq. (S19) and apply the projection inequality. $\square$

The plug-in algorithm uses the estimated moments. The strict certificate instead uses the distribution-free errors in Eq. (S22). For $c = 6$ and $\delta = 0.05$, $e_a = 6.86$, exceeding the admissible mismatch range $a \in [0, 4]$. Consequently, $a_L = 0$ for every record in this configuration. Zero transfer in every calibration record follows directly from this bound, while the plug-in estimator remains data-adaptive.

D.5. An optional confidence-certified transfer weight

Define $a_L = \max(0, \hat{a} - e_a)$ and $v_U = \hat{v} + e_v$. Let $\rho_{\text{cert}} = a_L / (a_L + v_U)$, with value zero when the denominator is zero.

Proposition S10 (Certified improvement relative to the anchor). *On the simultaneous calibration event,*

$$\mathbb{E} \|\tilde{K}_{\text{cert}} - K\|_N^2 \leq a - \frac{a_L^2}{a_L + v_U}. \quad (\text{S27})$$

For $a_L = 0$, the certificate chooses the anchor.

Proof. Subtract the anchor risk a from Eq. (S18). For $0 \leq \rho \leq 1$,

$$R(\rho) - a = (-2\rho + \rho^2)a + \rho^2 v. \quad (\text{S28})$$

The coefficient of a is nonpositive. Therefore $a \geq a_L$ and $v \leq v_U$ imply

$$R(\rho) - a \leq -2\rho a_L + \rho^2(a_L + v_U).$$

Substituting ρ_{cert} proves Eq. (S27); projection can only help. $\square$

The primary algorithm uses the plug-in weight. Under this six-donor configuration, the certified objective reduces algebraically to the EEG-selected-path anchor.

D.6. Rotation and independence

Three cyclic assignments reuse all donors in different roles, but each individual risk statement concerns a teacher cohort disjoint from its own selector and calibrator. The three final targets are not independent of one another. Linearity of expectation nevertheless bounds their average risk by the average of the three bounds above; independence is not required for that step. If a simultaneous confidence statement is desired, use $\delta/3$ per rotation. During training, the three

losses may have different selected paths, so replacing them by one averaged target is generally not algebraically equivalent.

E. Globally Coupled Monotone Alignment

E.1. The actual optimization problem

Let $\mathcal{P}_{T,L}$ contain paths with $\pi(1) = 1$, $\pi(T) = L$, and increments in $\{0, 1\}$. Choosing the $L - 1$ transition positions among $T - 1$ gaps gives

$$|\mathcal{P}_{T,L}| = \binom{T-1}{L-1}. \quad (\text{S29})$$

The selector minimizes

$$\hat{J}(\pi) = \|G_\pi^0\|_N^2 - 2\langle G_\pi^0, B^A \rangle_N. \quad (\text{S30})$$

Its expectation is $J(\pi) = \|G_\pi^0 - K\|_N^2 - \|K\|_N^2$. Negative empirical values are permitted; they are not negative squared distances, because a common constant was omitted. We do not clip the selection score, which would create artificial ties.

A separable dynamic-time-warping recurrence is not exact for Eq. (S30), because off-diagonal blocks couple choices at nonadjacent times. The default $T = 8$, $L = 4$ yields 35 paths, allowing an exact search with no approximation to the stated objective.

E.2. A quadratic sufficient-statistic tensor

Set $U_k = PH_k^0$ and partition B^A into $n \times n$ blocks B_{tu} . Precompute

$$q_{tu}(k, l) = \frac{\|U_k U_l^\top\|_F^2 - 2\langle U_k U_l^\top, B_{tu} \rangle_F}{n^2 T^2}. \quad (\text{S31})$$

Then $q_{tu}(k, l) = q_{ut}(l, k)$, and

$$\hat{J}(\pi) = \sum_t q_{tt}(\pi(t), \pi(t)) + 2 \sum_{t < u} q_{tu}(\pi(t), \pi(u)). \quad (\text{S32})$$

This identity follows by expanding the Frobenius norm and inner product block by block. Candidate scoring is subsequently $O(T^2)$, independent of the number of images and the state width. The same tensor gives global min-marginals for the visualization:

$$C_{kt} = \min_{\pi: \pi(t)=k} \hat{J}(\pi) - \min_{\pi} \hat{J}(\pi). \quad (\text{S33})$$

An infeasible (k, t) receives an explicit mask, not an arbitrary large plotted number.

E.3. An admissible branch-and-bound extension

For longer sequences, fix a prefix $p_1, \dots, p_j$. At a future time $t > j$, a necessary label set is

$$\mathcal{L}_t = \{k : \max(p_j, L - (T - t)) \leq k \leq \min(L, p_j + t - j)\}. \quad (\text{S34})$$

Let C_{pre} be the exact contribution of pairs within the prefix. The following quantity is a lower bound on every completion:

$$\begin{aligned} \text{LB} = C_{\text{pre}} &+ \sum_{t>j} \min_{k \in \mathcal{L}_t} \left[q_{tt}(k, k) + 2 \sum_{u \leq j} q_{ut}(p_u, k) \right] \\ &+ 2 \sum_{j < t < u} \min_{k \in \mathcal{L}_t, l \in \mathcal{L}_u} q_{tu}(k, l). \end{aligned} \quad (\text{S35})$$

Proposition S11 (Exact search with admissible pruning). *Pruning a prefix when its lower bound exceeds the best complete-path score cannot remove a better solution. Exhaustive search and this branch-and-bound procedure return the same global optimum, with a common deterministic tie rule.*

Proof. Every feasible completion uses a label from each $\mathcal{L}_t$. Its future unary-plus-prefix contribution is at least the corresponding independent minimum in the second line of Eq. (S35). Each future–future pair is likewise at least its independent minimum. Summing these inequalities proves admissibility, even though the independently minimizing labels may be mutually inconsistent. If the bound exceeds an incumbent complete score, all completions are worse and can be removed. A finite search that prunes only such prefixes retains a globally optimal complete path. Resolving equal complete scores by the same lexicographic rule makes the returned optimizer deterministic. $\square$

The implementation uses a small floating-point margin before pruning and compares against exhaustive enumeration in unit tests. No speedup factor is reported before runtime is measured.

E.4. Selection stability

Theorem S12 (Global path regret). *Let $\hat{\pi}$ minimize Eq. (S30), and π^* minimize the population discrepancy $a_\pi = \|G_\pi^0 - K\|_N^2$. Then*

$$a_{\hat{\pi}} - a_{\pi^*} \leq 4\|B^A - K\|_N, \quad \mathbb{E}[a_{\hat{\pi}} - a_{\pi^*}] \leq 4\sqrt{v_{|A|}}. \quad (\text{S36})$$

Proof. For every candidate, $\hat{J}(\pi) - J(\pi) = -2\langle G_\pi^0, B^A - K \rangle_N$. Cauchy–Schwarz and Lemma S1 bound its magnitude by $2\|B^A - K\|_N$, uniformly over all paths. Since $\hat{J}(\hat{\pi}) \leq \hat{J}(\pi^*)$, add and subtract the two empirical scores to obtain a total deviation of at most twice that uniform bound. The constant $\|K\|_N^2$ cancels. Finally, Jensen’s inequality gives $\mathbb{E}\|B^A - K\|_N \leq \sqrt{v_{|A|}}$. $\square$

If the population optimum is unique with gap $g > 0$ to the next path, the selected path is correct whenever $4\|B^A - K\|_N < g$. This motivates reporting both the selected path and its cohort-to-cohort stability rather than presenting a single heatmap as a discovered biological clock.

F. What Joint Trajectory Geometry Identifies

F.1. One isometry for the observed trajectory

Theorem S13 (Gram identification on observed spans). *Let $F \in \mathbb{R}^{N \times d}$ and $U \in \mathbb{R}^{N \times p}$ satisfy $FF^\top = UU^\top$. There is an isometry from the span of the rows of U to the span of the rows of F mapping each observed row of U to its corresponding row of F .*

Proof. Let the common nonzero eigendecomposition be $V\Lambda V^\top$, where Λ is positive diagonal of rank r . Define $R_F = F^\top V\Lambda^{-1/2}$ and $R_U = U^\top V\Lambda^{-1/2}$. Then

$$R_F^\top R_F = I_r, \quad R_U^\top R_U = I_r. \quad (\text{S37})$$

The feature matrices factor as $F = V\Lambda^{1/2}R_F^\top$ and $U = V\Lambda^{1/2}R_U^\top$. Hence $F = UO$ with $O = R_U R_F^\top$. For any row-span vector $z = a^\top R_U^\top$, $zO = a^\top R_F^\top$, which has the same norm and preserves all inner products. This is the required isometry. After padding the smaller ambient dimension with zero coordinates, it can be extended to an orthogonal transformation on the common ambient space. $\square$

The theorem applies to the entire stacked trajectory. It therefore uses one coordinate correspondence across all windows, rather than a different correspondence at each window. It identifies the geometry of observed centered states; the physical meaning of an EEG channel remains that of a measurement coordinate.

F.2. A counterexample to independent time-wise matching

Let $n = 2$, $T = 2$, and $u = (1, -1)^\top/2$. Let the reference trajectory be $U_1 = u$, $U_2 = u$, while the student uses $F_1 = u$, $F_2 = -u$. At each time,

$$U_t U_t^\top = F_t F_t^\top = uu^\top. \quad (\text{S38})$$

Nevertheless, the reference change is $U_2 - U_1 = 0$, while the student change is $F_2 - F_1 = -2u$. The reference cross-time block is $uu^\top$ and the student’s is $-uu^\top$. Thus exact instantaneous matching can coexist with an incorrect transition geometry. The joint kernel distinguishes these two trajectories.

F.3. A quantitative bound for changes and contrasts

Let $D \in \mathbb{R}^{(T-1) \times T}$ be the first-difference matrix, with -1 and $+1$ in adjacent columns, and let $\Delta = D \otimes I_n$.

Proposition S14 (Increment transfer). *For any F, U ,*

$$\|\Delta(FF^\top - UU^\top)\Delta^\top\|_F \leq 4\|FF^\top - UU^\top\|_F. \quad (\text{S39})$$

For a general linear contrast operator C , the factor 4 is replaced by $\|C\|_{\text{op}}^2$.

Proof. For a sequence of vectors z_t ,

$$\sum_{t=1}^{T-1} \|z_{t+1} - z_t\|_2^2 \leq 2 \sum_{t=1}^{T-1} (\|z_{t+1}\|_2^2 + \|z_t\|_2^2) \leq 4 \sum_{t=1}^T \|z_t\|_2^2. \quad (\text{S40})$$

Therefore $\|\Delta\|_{\text{op}} \leq 2$. Apply $\|ABC\|_F \leq \|A\|_{\text{op}}\|B\|_F\|C\|_{\text{op}}$ to the kernel difference. The same argument yields the general contrast statement. $\square$

The result concerns inner products of observed state changes. Establishing equivalence of dynamical systems requires additional conditions beyond Gram matching [1, 2]. Accordingly, the intervention experiments ask whether transferred state directions participate in the *student's* decisions, a property that can be measured directly.

G. Prediction Stability and Student Capacity

G.1. An explicit readout certificate

Consider two PSD training kernels $K_1, K_2 \in \mathbb{R}^{m \times m}$, label matrix $Y \in \mathbb{R}^{m \times c}$, regularization $\gamma > 0$, and test-to-training kernel vectors κ_1, κ_2 . Define kernel-ridge scores

$$p_j = \kappa_j^\top (K_j + m\gamma I)^{-1} Y. \quad (\text{S41})$$

Theorem S15 (Kernel-ridge prediction stability). *Writing $\delta_K = \|K_1 - K_2\|_{\text{op}}$ and $\delta_\kappa = \|\kappa_1 - \kappa_2\|_2$,*

$$\|p_1 - p_2\|_2 \leq \|Y\|_{\text{op}} \left[\frac{\delta_\kappa}{m\gamma} + \frac{\|\kappa_2\|_2 \delta_K}{(m\gamma)^2} \right]. \quad (\text{S42})$$

If the second score vector has a winning-class margin larger than twice this bound, the two readouts select the same class.

Proof. Let $A_j = (K_j + m\gamma I)^{-1}$. PSD implies $\|A_j\|_{\text{op}} \leq 1/(m\gamma)$. Decompose

$$p_1 - p_2 = (\kappa_1 - \kappa_2)^\top A_1 Y + \kappa_2^\top (A_1 - A_2) Y. \quad (\text{S43})$$

The resolvent identity gives $A_1 - A_2 = A_1(K_2 - K_1)A_2$. Apply operator-norm submultiplicativity to both terms to obtain Eq. (S42). Each score coordinate changes by no more than the vector's ℓ_2 bound. The difference between a winner and any competitor therefore changes by at most twice that amount, proving the margin claim. $\square$

The test-to-training term is measured using held-out donor EEG only after model selection is frozen. It is a diagnostic requirement, not something inferred from a low training loss. This makes the certificate useful: it identifies precisely which missing generalization quantity an experiment must estimate.

G.2. Teacher risk propagates through a convex ideal student

A tractable kernel-space learner provides a second bridge. For $G \in \mathcal{C}$, define

$$\Phi(G) = \mu \text{tr}[Y^\top (G + \eta I)^{-1} Y], \quad \mu, \eta > 0, \quad (\text{S44})$$

and

$$\mathcal{S}(T) = \arg \min_{G \in \mathcal{C}} \{\Phi(G) + \lambda \|G - T\|_N^2\}, \quad \lambda > 0. \quad (\text{S45})$$

This is an analysis model and a possible diagnostic baseline, not an assertion that a recurrent network solves this convex optimization problem.

Proposition S16 (Nonexpansive teacher-to-student map). *The optimizer exists, is unique, and satisfies*

$$\|\mathcal{S}(T_1) - \mathcal{S}(T_2)\|_N \leq \|T_1 - T_2\|_N. \quad (\text{S46})$$

Consequently, the teacher-risk bound also bounds $\mathbb{E}\|\mathcal{S}(\tilde{K}) - \mathcal{S}(K)\|_N^2$.

Proof. Along a symmetric direction D , the second derivative of $\text{tr}[Y^\top (G + \eta I)^{-1} Y]$ equals

$$2\|(G + \eta I)^{-1/2} D(G + \eta I)^{-1} Y\|_F^2 \geq 0. \quad (\text{S47})$$

Thus Φ is convex. The quadratic term is strictly convex, and the feasible set is compact, so there is one optimizer. Let $\psi = \Phi + \iota_{\mathcal{C}}$ and $G_j = \mathcal{S}(T_j)$. Optimality gives subgradients $u_j \in \partial\psi(G_j)$ with $u_j = -2\lambda(G_j - T_j)/N^2$. Convex subgradients are monotone: adding their two defining inequalities yields $\langle u_1 - u_2, G_1 - G_2 \rangle_F \geq 0$. Substitute the expressions for u_j , rearrange, and apply Cauchy–Schwarz to obtain

$$\begin{aligned} \|G_1 - G_2\|_F^2 &\leq \langle T_1 - T_2, G_1 - G_2 \rangle_F \\ &\leq \|T_1 - T_2\|_F \|G_1 - G_2\|_F. \end{aligned} \quad (\text{S48})$$

Divide, with the zero-difference case immediate. The expectation bound follows pointwise by setting $T_1 = \tilde{K}, T_2 = K$. $\square$

G.3. A path and rank approximation floor

Let $R_\pi \in \mathbb{R}^{T \times L}$ encode the selected path, with row t equal to the one-hot vector for $\pi(t)$. Since each step is visited, $R_\pi^\top R_\pi$ is invertible. Define the orthogonal projector

$$W_\pi = [R_\pi (R_\pi^\top R_\pi)^{-1} R_\pi^\top] \otimes P. \quad (\text{S49})$$

Every realizable student trajectory kernel lies in the range of this projector and has rank at most d .

Proposition S17 (Unavoidable path–capacity error). *For PSD K , let λ_j be the decreasing eigenvalues of $W_\pi K W_\pi$. Then*

$$\min_{G \succeq 0, W_\pi G W_\pi = G, \text{rank } G \leq d} \|G - K\|_F^2 = \|K - W_\pi K W_\pi\|_F^2 + \sum_{j>d} \lambda_j^2. \quad (\text{S50})$$

It is therefore a lower bound for a finite neural architecture with the same path and state width.

Proof. The map $A \mapsto W_\pi A W_\pi$ is an orthogonal projection in the Frobenius inner product. Thus $K - W_\pi K W_\pi$ is orthogonal to every feasible $G - W_\pi K W_\pi$, and the squared error decomposes into the two corresponding squared norms. It remains to approximate $C = W_\pi K W_\pi \succeq 0$ by a rank- d matrix in the same range. Let V span the column space of any rank- d candidate. Projection onto that column space gives a residual at least $\|(I - VV^\top)C\|_F^2 = \text{tr}(C^2) - \text{tr}(V^\top C^2 V)$. Expanding in the eigenbasis of C^2 shows the captured trace is at most the sum of its largest d eigenvalues: the coefficients lie in $[0, 1]$ and sum to at most d . Hence the error is at least $\sum_{j>d} \lambda_j^2$. Retaining the largest d eigencomponents of C is PSD, lies in the required range, and attains the bound. $\square$

This decomposition distinguishes too few reasoning steps from too narrow a state. A sensor-scaling plateau can therefore be examined jointly with teacher reliability and the student’s approximation floor, rather than assigned a single explanation from its shape alone.

H. Information Budgets and Additional Sensors

H.1. A linear-Gaussian measurement model

To reason about sensor budgets, let a task-relevant latent vector obey $z \sim \mathcal{N}(0, \Sigma_0)$ with $\Sigma_0 \succ 0$. For an existing sensor set S , observe

$$y_S = A_S z + \epsilon_S, \quad \epsilon_S \sim \mathcal{N}(0, R_S), \quad R_S \succ 0, \quad (\text{S51})$$

independently of z . This model supplies an interpretable measurement analysis; its parameters are not presumed to be known for human EEG.

Proposition S18 (Posterior covariance). *The covariance of $z \mid y_S$ is*

$$\Sigma_S = (\Sigma_0^{-1} + A_S^\top R_S^{-1} A_S)^{-1}. \quad (\text{S52})$$

Proof. The negative log posterior, up to an additive constant, is

$$\frac{1}{2} z^\top \Sigma_0^{-1} z + \frac{1}{2} (y_S - A_S z)^\top R_S^{-1} (y_S - A_S z).$$

Collect the quadratic terms in z and complete the square. The resulting Gaussian precision is $\Sigma_0^{-1} + A_S^\top R_S^{-1} A_S$, proving Eq. (S52). $\square$

H.2. What a new independent measurement contributes

Add $y_+ = a^\top z + \epsilon_+$ with independent noise variance $\sigma^2 > 0$.

Theorem S19 (Conditional innovation, estimation gain, and information). *The updated covariance is*

$$\Sigma_{S+} = \Sigma_S - \frac{\Sigma_S a a^\top \Sigma_S}{\sigma^2 + a^\top \Sigma_S a}. \quad (\text{S53})$$

The reduction in Bayes squared-error risk for z and the conditional mutual information are

$$\Delta_z = \frac{a^\top \Sigma_S^2 a}{\sigma^2 + a^\top \Sigma_S a} \geq 0, \quad (\text{S54})$$

$$I(z; y_+ \mid y_S) = \frac{1}{2} \log \left(1 + \frac{a^\top \Sigma_S a}{\sigma^2} \right). \quad (\text{S55})$$

For a linear task score $c^\top z$, the Bayes risk reduction is

$$\Delta_c = \frac{(c^\top \Sigma_S a)^2}{\sigma^2 + a^\top \Sigma_S a}. \quad (\text{S56})$$

Proof. The new precision is $\Sigma_S^{-1} + a a^\top / \sigma^2$. Multiply the right side of Eq. (S53) by this precision; the rank-one terms cancel, leaving the identity. Thus Eq. (S53) is its inverse. Under squared error, the conditional mean is Bayes optimal: expanding $\|z - g(y)\|^2$ around $\mathbb{E}[z \mid y]$ leaves a zero cross term and a nonnegative excess term. Its Bayes risk is the trace of the posterior covariance. Taking the trace of the rank-one decrement gives Δ_z ; multiplying on both sides by c gives Δ_c .

For the information identity, $y_+ \mid y_S$ is Gaussian with variance $\sigma^2 + a^\top \Sigma_S a$, while $y_+ \mid z, y_S$ has variance σ^2 . The entropy of a scalar Gaussian of variance v is $\frac{1}{2} \log(2\pi e v)$, obtained by integrating its log density and using its variance. Subtract these two conditional entropies to obtain Eq. (S55). $\square$

A measurement can reduce uncertainty about z while contributing nothing to a particular task if $c^\top \Sigma_S a = 0$. This motivates evaluating task-relevant innovation, not just electrode count or effective rank.

H.3. Correlated sensor noise

Suppose $\text{Cov}(\epsilon_S, \epsilon_+) = r$. Form

$$\tilde{y}_+ = y_+ - r^\top R_S^{-1} y_S, \quad (\text{S57})$$

$$\tilde{a} = a - A_S^\top R_S^{-1} r, \quad (\text{S58})$$

$$\tilde{\sigma}^2 = \sigma^2 - r^\top R_S^{-1} r. \quad (\text{S59})$$

The residual noise is uncorrelated with ϵ_S and, under the joint Gaussian model, independent of it. It is also independent of z . Therefore Theorem S19 applies with a, σ^2 replaced by these residual quantities. The Schur complement $\tilde{\sigma}^2$ must be positive for the displayed inverse formulation. An exact digital copy of an existing observation contributes no conditional information; handle that singular limit directly rather than dividing 0 by 0.

This derivation suggests three distinct empirical comparisons: independent physical channels, deterministic copies of existing channels, and actually recorded additional modalities. EEG and eye tracking permit the third comparison in the chosen archive. fMRI and ECG would require their own paired measurements and measurement models; a fitted EEG channel curve supplies neither.

I. Implementation and Numerical Verification

I.1. Exact computations without a waveform generator

The released reference core takes bounded arrays with shape $[S, Tn, p]$. It implements the U-statistic, distinct-index moments, PSD trace projection, exact quadratic path search, admissible branch-and-bound, one cohort rotation, and the kernel gradient. Dataset parsing and student fitting are separate pipeline stages.

For large N , avoid materializing every pair kernel when estimating moments. The identity

$$\begin{aligned} & \langle \text{sym}(AB^\top), \text{sym}(CD^\top) \rangle_F \\ &= \frac{1}{2} \{ \text{tr}(A^\top CD^\top B) + \text{tr}(A^\top DC^\top B) \} \end{aligned} \quad (\text{S60})$$

uses feature-space cross-products. It follows by expanding both symmetrizations, observing that transpose-paired terms have equal traces, and cyclically permuting each trace. This is algebraically identical to the dense reference, not a new estimator.

If a PSD target has decomposition $\tilde{K} = V\Lambda V^\top$, the student loss can be computed as

$$\begin{aligned} \|FF^\top - \tilde{K}\|_F^2 &= \|F^\top F\|_F^2 - 2\text{tr}(F^\top V\Lambda V^\top F) \\ &\quad + \|\tilde{K}\|_F^2. \end{aligned} \quad (\text{S61})$$

Retaining all positive eigenvalues makes this exact. A truncated approximation must report omitted spectral energy and be compared with the exact target; it is not silently substituted into the theory.

I.2. Gradient and aggregation

For a fixed symmetric target T ,

$$\nabla_F \|FF^\top - T\|_N^2 = \frac{4}{N^2} (FF^\top - T)F. \quad (\text{S62})$$

To verify it, write the differential as $2\langle FF^\top - T, dFF^\top + FdF^\top \rangle_F / N^2$; symmetry makes the two terms equal. Move factors cyclically inside the trace to read off Eq. (S62). Back-propagate through the fixed projection and through repeated occurrences of each state along the path. In particular, add, rather than average away, all appearances of the same h_i^k before applying the loss normalization.

Pairwise losses from a small independently normalized mini-batch are not the full centered-kernel objective. The primary implementation caches frozen visual tokens and computes the neural loss on the complete training-item index. Memory-limited implementations can accumulate exact matrix products over item blocks; they must reproduce Eq. (S61) and its gradient.

I.3. The decision-carrying visual workspace

Input images are letterboxed to the frozen encoder’s specified resolution while preserving their board layout. DINOv2 patch tokens are projected to width $d = 256$. A workspace with eight learned tokens performs four recurrent updates using shared cross-attention, self-attention, and a feed-forward block. Pre-layer normalization and residual connections are used in each sub-block. There are four attention heads and MLP width 1024. The first workspace token is the decision state. Bound its norm to one after the recurrent residual update, and use that same state both for the kernel and for the final four-way classifier.

The classifier matrix is spectrally normalized to operator norm at most one in the forward readout; a fixed logit temperature of 0.2 controls their scale. There is no bypass from an image-conditioned answer head around the supervised state. All architecture controls retain these choices. A fixed DINO feature plus linear classifier and a candidate-only visual baseline are included to quantify how much reasoning the frozen representation already supports.

The VLM replication uses the locked Qwen2.5-VL checkpoint and a fixed prompt asking for one of four option indices. Extract states at four predeclared residual-depth positions from the final *prompt* position before any answer token is appended. The supervised state is extracted before answer conditioning. LoRA settings and trainable modules are specified in Section N. Depth alignment supplies the ordered representation axis for this replication.

I.4. Unit tests that precede any experiment

The numerical test suite checks algebraic identities and edge cases on fixed test arrays. It compares the fast U-statistic with an explicit pair average, verifies moment indexing,

tests a Gaussian special-case variance formula, checks PSD/range/trace constraints and the projection inequality, compares exhaustive path search with branch-and-bound, and verifies gradients by finite differences. It also checks that the counterexample in Section F has identical diagonal blocks and different cross-time blocks.

The test log records the software environment, random seed, and identity checks separately from the experiment manifest.

J. Path-Compatible Calibration and Independent Risk Audits

This section derives the target used by RELAY. The preceding U-statistic and projection results apply to any fixed orthogonal observation map. After path selection, the calibrator and teacher remain independent of that map; replacing $\mathcal{P}$ by W_π in those results gives the estimators below.

J.1. The temporal repetition operator

For an ordered path π , let $R_\pi \in \mathbb{R}^{T \times L}$ have $R_\pi[t, k] = 1\{\pi(t) = k\}$. Set

$$Q_\pi = R_\pi(R_\pi^\top R_\pi)^{-1}R_\pi^\top, \quad W_\pi = Q_\pi \otimes P. \quad (\text{S63})$$

Every step is visited, so the diagonal matrix $R_\pi^\top R_\pi$ has strictly positive entries. Direct multiplication gives $Q_\pi^\top = Q_\pi$ and $Q_\pi^2 = Q_\pi$. The same identities hold for W_π . For the stacked student $F_\pi = (R_\pi \otimes P)H$, $W_\pi F_\pi = F_\pi$, and hence $W_\pi G_\pi W_\pi = G_\pi$.

Lemma S20 (Orthogonal decomposition of the distillation target). *The operator $\mathcal{W}_\pi(A) = W_\pi A W_\pi$ is an orthogonal projection in matrix Frobenius space. For $K_\pi = W_\pi K W_\pi$ and every G satisfying $W_\pi G W_\pi = G$,*

$$\|G - K\|_N^2 = \|G - K_\pi\|_N^2 + \|K - K_\pi\|_N^2. \quad (\text{S64})$$

Proof. Idempotence follows from $W_\pi^2 = W_\pi$. Cyclic invariance of the trace gives $\langle A, W_\pi B W_\pi \rangle_N = \langle W_\pi A W_\pi, B \rangle_N$, so the operator is self-adjoint. An idempotent self-adjoint operator is an orthogonal projector. In particular, $K - K_\pi$ is orthogonal to every matrix in its range, including $G - K_\pi$. Expanding the squared norm of $(G - K_\pi) - (K - K_\pi)$ gives the result. $\square$

The second term is a temporal-repetition floor. It is constant during student fitting at a selected path. The first term is the component used to choose the teaching strength. Rank and row-norm constraints impose additional representational requirements; the convex teacher set enforces temporal compatibility rather than claiming that every target is exactly expressible by a width- d network.

J.2. Conditional moment estimation in the compatible space

Condition on the selector cohort and its selected path. For calibration donors define $X_s^\pi = W_\pi Z_s$, $h_{s,r}^\pi = \text{sym}(X_s^\pi X_r^{\pi\top})$, and the corresponding raw moments $m_{j,\pi}$. The common mean is $W_\pi M$ and the population kernel is K_π . Applying Theorem S3 yields

$$v_{b,\pi} = \frac{4(m_{1,\pi} - m_{2,\pi})}{b} + \frac{2(m_{0,\pi} - 2m_{1,\pi} + m_{2,\pi})}{b(b-1)}. \quad (\text{S65})$$

The unclipped estimate

$$\hat{a}_{\pi,\text{raw}} = \|G_\pi^0\|_N^2 - 2\langle G_\pi^0, B_\pi^c \rangle_N + \hat{m}_{2,\pi}^c \quad (\text{S66})$$

is conditionally unbiased for $a_\pi = \|G_\pi^0 - K_\pi\|_N^2$. Its expectation follows by linearity and independence from the selector. Clipping to $[0, 4]$ produces the operational mismatch estimate; the raw estimate remains in the audit ledger. Variance components are clipped individually to $[0, 1]$ before inserting the actual teacher-cohort size b .

Theorem S21 (Risk calibration at the student's temporal resolution). *Let the selector, calibrator, and teacher cohorts be disjoint. Conditional on the selector and calibrator, let $\rho \in [0, 1]$ be fixed and let B_π be an unbiased independent teacher with mean K_π and risk v_π . Define*

$$A_\rho = (1 - \rho)G_\pi^0 + \rho B_\pi, \quad (\text{S67})$$

$$\tilde{K}_\pi = \Pi_{\mathcal{C}_\pi}(A_\rho), \quad (\text{S68})$$

$$\mathcal{C}_\pi = \{A \succeq 0 : W_\pi A W_\pi = A, \text{tr } A \leq N\}. \quad (\text{S69})$$

For $c_\pi = \|K - K_\pi\|_N^2$,

$$\mathbb{E}\|A_\rho - K\|_N^2 = c_\pi + (1 - \rho)^2 a_\pi + \rho^2 v_\pi, \quad (\text{S70})$$

$$\mathbb{E}\|\tilde{K}_\pi - K\|_N^2 \leq c_\pi + \frac{a_\pi v_\pi}{a_\pi + v_\pi} + (a_\pi + v_\pi)(\rho - \rho_\pi^*)^2, \quad (\text{S71})$$

$$\rho_\pi^* = \frac{a_\pi}{a_\pi + v_\pi}. \quad (\text{S72})$$

When $a_\pi + v_\pi = 0$, set $\rho_\pi^* = 0$ and interpret the fraction as zero.

Proof. Write $B_\pi = K_\pi + \varepsilon$ with conditional mean-zero ε . Both $G_\pi^0 - K_\pi$ and ε lie in the range of $\mathcal{W}_\pi$. Lemma S20 separates the orthogonal residual $K - K_\pi$, giving

$$\|A_\rho - K\|_N^2 = c_\pi + \|(1 - \rho)(G_\pi^0 - K_\pi) + \rho\varepsilon\|_N^2.$$

Taking conditional expectation removes the cross term and gives the first identity. The bounded-row assumptions imply $K_\pi \in \mathcal{C}_\pi$, because projection preserves PSD and cannot increase the trace of $\mathcal{P} M M^\top \mathcal{P}$. The projection inequality

in Lemma S6, applied with C_π , bounds the compatible-space error by that of A_ρ . Both projected and unprojected matrices are in the same temporal range, so their full-space errors add the same c_π . Completing the square in $(1 - \rho)^2 a_\pi + \rho^2 v_\pi$ gives Eq. (S71). If $a_\pi = v_\pi = 0$, the compatible error vanishes almost surely and the full-space error is c_π . $\square$

The oracle minimizes the pre-projection conditional risk and its upper bound is no greater than $c_\pi + a_\pi$ or $c_\pi + v_\pi$. The implemented method uses the independently estimated ratio $\hat{a}_\pi / (\hat{a}_\pi + \hat{v}_\pi)$. Its calibration error appears explicitly in the final square term.

Corollary S22 (An irrelevant floor can inflate a transfer weight). *Holding the compatible variance $v > 0$ fixed, adding an orthogonal mismatch $c \geq 0$ changes the oracle ratio by*

$$\frac{a + c}{a + c + v} - \frac{a}{a + v} = \frac{cv}{(a + v)(a + c + v)} \geq 0. \quad (\text{S73})$$

Proof. Subtract the fractions over their common positive denominator. The numerator is $(a + c)(a + v) - a(a + c + v) = cv$. $\square$

This isolates the mismatch effect. A full-space estimator also changes the variance; its empirical comparison therefore recomputes both moments instead of assuming that the above monotonic statement applies unchanged.

J.3. Two-reference-donor risk differences

Let T_1, T_0 be fixed after training-cohort processing. Two new donors give $B_H = h_{12}$ with $\mathbb{E}[B_H | T_1, T_0] = K$. This reference is not PSD projected.

Proposition S23 (Observable, unbiased risk difference). *The statistic*

$$\hat{\Delta}(T_1, T_0) = \|T_1\|_N^2 - \|T_0\|_N^2 - 2\langle T_1 - T_0, B_H \rangle_N \quad (\text{S74})$$

has conditional expectation $\|T_1 - K\|_N^2 - \|T_0 - K\|_N^2$.

Proof. Expand both population squared errors. Their common term $\|K\|_N^2$ cancels. The remaining expression is linear in K , so substituting an independent unbiased B_H preserves its conditional expectation. $\square$

No estimate of reference variance is needed for this difference. In contrast, $\|T - B_H\|_N^2$ adds the reference's own mean-squared error to the target risk. Replacing B_H by its PSD projection changes its mean and no longer gives this unbiased audit.

J.4. Absolute risk with an independent eight-donor panel

The absolute audit reserves eight donors and uses three four-donor training roles. For the same fixed W_π , compute the raw reference $B_{H,\pi}$ and a distinct-four-donor moment $\hat{m}_{2,H,\pi}$.

Proposition S24 (Observable, unbiased absolute risk). *For $|H| \geq 4$ independent reference donors,*

$$\hat{R}_\pi(T) = \|T\|_N^2 - 2\langle T, B_{H,\pi} \rangle_N + \hat{m}_{2,H,\pi} \quad (\text{S75})$$

has conditional expectation $\|T - K_\pi\|_N^2$.

Proof. The reference kernel has conditional expectation K_π . Distinct reference pairs are independent, so $\mathbb{E}\hat{m}_{2,H,\pi} = \|K_\pi\|_N^2$. The result follows by expansion. The reference remains independent after conditioning on the chosen path and all training quantities because neither reference donor data nor their outcomes were used for selection. $\square$

Absolute estimates are not clipped. For uncertainty, target records are first averaged, then each reference donor is deleted from all records together. For $h = 8$ and delete-one values $\bar{R}_{(-s)}$,

$$\widehat{\text{SE}}_J^2 = \frac{h-1}{h} \sum_s (\bar{R}_{(-s)} - h^{-1} \sum_r \bar{R}_{(-r)})^2. \quad (\text{S76})$$

The reported interval is $\bar{R} \pm t_{7,975} \widehat{\text{SE}}_J$. This is a reference-donor jackknife interval, not a 240-independent-record interval. The student configurations for the two- and eight-reference audits are reported separately.

K. Mechanism Controls with Exact Invariances

K.1. Deleting a loss block versus setting its target to zero

The true time-wise baseline is

$$\mathcal{L}_{\text{diag}} = \frac{1}{T} \sum_t \frac{\|F_t F_t^\top - T_{tt}\|_F^2}{n^2}. \quad (\text{S77})$$

Its gradient is $4(F_t F_t^\top - T_{tt})F_t / (Tn^2)$. There are no off-diagonal residuals in that gradient. A different control uses the full loss against a block-diagonal target, introducing $\sum_{t \neq u} \|F_t F_u^\top\|_F^2 / N^2$. Both controls are included under distinct names. Separate inner validation selects their regularization weight from the same three-value grid, so a missing normalization factor cannot explain the comparison.

Proposition S25 (A time-wise geometry-preserving intervention). *For orthogonal matrices O_t shared across all donors in window t , define $Z'_{s,t} = Z_{s,t} O_t$. Every*

within-window cross-donor kernel is unchanged. A cross-time product changes to $Z_{s,t}O_tO_u^\top Z_{r,u}^\top$.

Proof. At $t = u$, $O_tO_t^\top = I$, so the original product is recovered, before and after symmetrization, averaging, and common item projection. At $t \neq u$, the displayed factor remains and is generally not the identity. This preserves each instantaneous geometry while modifying the temporal coordinate correspondence. $\square$

The rotations are QR factors of seeded Gaussian matrices. The same O_t is used in selector, calibrator, teacher, and audit donors. Cross-time information is changed before selection; the entire teacher pipeline is then rebuilt.

K.2. Readout-sensitive and null directions

Let the four-choice linear readout have weight $V \in \mathbb{R}^{4 \times d}$. Define $C_4 = I_4 - \mathbf{1}\mathbf{1}^\top/4$ and $J = C_4V$. The projector onto probability-sensitive directions is $P_{\text{read}} = J^+J$.

Lemma S26 (Final-state nullspace deletion preserves probabilities). $\text{rank}(P_{\text{read}}) \leq 3$. For every state h ,

$$\text{softmax}(Vh) = \text{softmax}(VP_{\text{read}}h). \quad (\text{S78})$$

Proof. Since C_4 has rank three, J has rank at most three. The vector $h - P_{\text{read}}h$ is in the nullspace of J , implying that $V(h - P_{\text{read}}h)$ is a constant multiple of $\mathbf{1}$. Adding a constant to all logits leaves softmax unchanged. This proves the identity. $\square$

Thus supervising the state does not alone establish that every supervised direction affects decisions. The experiments delete readout-sensitive directions, delete final-readout null directions, and compare neural-aligned versus random subspaces using the same samples. Early-state perturbations can propagate through later nonlinear updates; the exact nullspace identity applies to the final readout only.

K.3. Paired difference-in-differences

For method m , let $C_{m,a}(i, s)$ and $C_{m,r}(i, s)$ be correctness after aligned and random removal. The item-level interaction is

$$d_i = \frac{1}{S} \sum_s [C_{\text{RELAY},r}(i, s) - C_{\text{RELAY},a}(i, s) - C_{\text{ERM},r}(i, s) + C_{\text{ERM},a}(i, s)]. \quad (\text{S79})$$

Its average is the difference of the two extra impairments. All four conditions are resampled with the same original-problem indices. Subtracting marginal interval endpoints would discard this covariance and is not used.

L. Dataset, Events, and Evaluation Units

L.1. Released material and the paired panel

We use version 5 of *Human Neurocognition During Abstract Reasoning: Multimodal Recordings of EEG, Eye-Tracking, and Behavioural Responses* [3]. The release contains 25 participants, including 20 who completed all five sessions. Each session contains 80 problems, giving 400 distinct problems in eight pattern families. The analysis panel comprises the 20 complete-session donors. Repetitions across people are measurements of the same problem, not additional independent training examples. Donor aliases D01–D20 denote the lexicographically ordered complete-session identifiers; the alias map retains the original release identifiers.

The primary outer fold has 350 training problems and 50 test problems. The total nominal panel contains $20 \times 400 = 8,000$ donor–problem trials and 64,000 donor–problem–window cells before feature validity coding. These counts describe the complete-session design. The data-ingestion manifest records the actual source path, original identifier, session, problem identifier, pattern, answer mapping, event indices, and validity bits for each cell.

The experiment presents an encoding sequence followed by a decision phase with four candidate answers. Encoding items are presented for 600 ms; the decision phase permits up to 12 s [3, 4]. The released continuous EEG is acquired with 64 BioSemi channels at 2,048 Hz, with four EOG channels. EyeLink recordings are sampled at 2,000 Hz. We start from the minimally processed continuous EEG, keeping event and behavior files attached to their original recording. Dataset and stimulus use follow the release’s CC BY 4.0 license.

L.2. The original problem is the grouping key

The stable identifier is the original problem ID, not a rendered image filename, donor–trial index, or window index. Every donor recording, icon-renamed image, and option-permuted image derived from an original problem inherits this identifier. A hash-based assertion checks that no parent ID occurs in both sides of an outer or inner split. The held family’s name and all its associated original problems are excluded before fitting normalization, EOG correction, paths, teacher moments, or model hyperparameters.

Answers are represented by the semantic candidate identity and by the displayed option index. An option permutation updates the displayed index but not the semantic answer. Four-class probabilities are saved in displayed A–D order with this mapping. The complete-sequence family string is used by the split generator and analysis, never by the visual model. The model’s only instruction is: *Complete the visual sequence. Answer with A, B, C, or D.*

L.3. Outer, inner, and joint donor splits

The eight outer folds leave one entire pattern family out. Inner selection uses the seven remaining families, leaving one of those seven out in turn. This produces 300 inner-training and 50 inner-validation problems. No outer test result selects a hyperparameter, a montage, a donor group, or a visualization case.

In the primary donor split, D01 and D02 are independent references. D03–D20 are partitioned into three consecutive groups of six; their selector, calibrator, and teacher roles rotate cyclically. This gives three fixed targets per outer-fold/seed pair and $8 \times 10 \times 3 = 240$ calibration records. The three rotations share donors and are not independent replications.

Joint family–donor evaluation uses ten disjoint reference pairs: (D01,D02), (D03,D04), through (D19,D20). For each pair, the remaining 18 donors are partitioned into groups of six by the same ordered rule. All eight family folds are refitted. The reference pair’s EEG contributes neither teacher selection nor model fitting, and CKA is evaluated only on the 50 problems in the held family. Thus task accuracy and neural alignment refer to the same held-family–held-donor unit. Pair zero reuses the primary fits; the remaining nine pairs add $9 \times 2 \times 8 \times 10 = 1,440$ ERM/RELAY fits. Table S11 reports every pair.

The separate absolute-risk audit reserves D13–D20 as eight references and uses D01–D12 in three groups of four. It is a distinct 12-donor teacher experiment, rather than an absolute-error estimate attributed to the 18-donor primary teacher. Target construction and reference uncertainty are specified in Section J.

M. Signal Processing and Feature Construction

M.1. A causal, donor-local EEG pipeline

Channels are converted to microvolts using the recording metadata. Within each continuous session, a fourth-order Butterworth high-pass at 0.5 Hz and fourth-order Butterworth low-pass at 40 Hz are applied with second-order-section causal filtering. Filter states are carried through the continuous session rather than reset at each event. Following low-pass filtering, the signal is decimated by a factor of eight to 256 Hz. The primary pipeline does not use forward–backward filtering, future-window interpolation, or a separate notch after the 40-Hz low-pass.

Average referencing uses only the retained physical EEG channels. Sensor removal therefore precedes re-referencing and every subsequent preprocessing stage. EOG correction is fitted separately within each donor, using only training-problem decision prefixes. Horizontal and vertical EOG differences, their first differences, and an intercept form the regression design. Training-standardized ridge regression with penalty 10^{-3} estimates each EEG channel’s ocular con-

tribution. The corresponding transform is frozen for that donor’s remaining trials. No regression coefficient is pooled across selector, calibrator, teacher, and reference cohorts.

A channel–window is invalid when its raw absolute amplitude exceeds $150 \mu\text{V}$, its one-sample increment exceeds $50 \mu\text{V}$ after decimation, or more than 20% of samples are absent. Flat segments have a training-scaled standard deviation below 10^{-3} for at least 250 ms. Invalid channels are not interpolated from channels removed in a sensor-budget experiment. A trial–window with more than 20% invalid retained channels receives the all-zero feature code; otherwise invalid channel coefficients are set to zero. The validity mask is saved separately and is the complete input to the mask-only teacher.

M.2. Decision-prefix timing

The primary prefix begins at the event marking the decision-board onset. It ends 300 ms before the recorded response. For a timeout, the maximum decision duration supplies the response-censoring endpoint. Prefixes shorter than 400 ms receive the invalid-feature code. No post-response sample enters a feature, even through a filtering or interpolation boundary.

The valid prefix is divided into eight equal progress bins. Each bin must contain at least four retained samples. At 256 Hz the shortest admitted prefix contains about 102 samples, leaving more than four samples per bin. Each channel is expressed by the first four orthonormal type-II DCT coefficients within a bin. With 64 channels, $p = 64 \times 4 = 256$. The transform captures a local signal level and its low-order within-bin variation without requiring a raw-waveform decoder.

Per-channel coefficient centering and scale use the donor’s training problems only: the median and 1.4826 times the median absolute deviation, with a 10^{-6} scale floor. Values are clipped to $[-8, 8]$, concatenated in the common channel/coefficient order, and divided by $\max(1, \|z\|_2)$. The full time-major matrix is then projected by the fixed item-centering operator. The empirical estimand is the mean of these processed, response-censored measurements.

M.3. Time, features, and nuisance alternatives

The timing variants use response guards of 150, 300, and 450 ms, and a stimulus-clock alternative with eight 100-ms windows after decision onset. Feature sweeps use 4, 6, 8, and 12 progress bins and 2, 4, or 8 DCT coefficients. Every variant recomputes the path set and teacher from its own training features. The primary $T = 8, L = 4$ setting has 35 paths; for a general pair, the count is $\binom{T-1}{L-1}$.

The primary nuisance design contains only the intercept. A visual-nuisance variant adds training-image mean intensity, contrast, edge density, and the first eight frozen-DINO principal scores. A label-nuisance variant adds three indepen-

dent correct-option indicators. These designs define different residualized targets and are reported as sensitivity analyses, not silently substituted preprocessing. Their regressors are fixed before any donor-dependent teacher stage.

M.4. Gaze, EOG, and behavioral teachers

Gaze positions are converted to board-relative coordinates using the display metadata. Pupil values are standardized within donor using training trials. Invalid tracking samples and blink events are masked; gaps longer than 75 ms remain invalid, and short gaps use causal last-valid-sample filling. For each of the same eight decision bins, the gaze vector contains mean horizontal and vertical position, their standard deviations, mean pupil size, fixation count, saccade count, and mean saccade amplitude. Every count is divided by the observed bin duration. The feature vector is training-standardized and norm bounded using the EEG recipe.

The EOG-only teacher uses the four recorded EOG channels and the same DCT construction. Reaction-time supervision uses log response time together with its square and a timeout indicator; the mask-only teacher uses the saved validity pattern. No behavior-derived teacher receives test-family labels. EEG+gaze concatenates separately bounded EEG and gaze vectors, each scaled by $1/\sqrt{2}$, before forming the joint donor kernel. This preserves an explicit modality scale and a unit row bound.

N. Models, Optimization, and Exact Loss Evaluation

N.1. Visual input and the recurrent workspace

Boards are rendered from the released stimulus specification. Images are letterboxed to 448×448 with a neutral background, preserving the complete sequence and all four candidates. The frozen DINOv2 ViT-B/14 encoder supplies 1,024 patch tokens of width 768; no task labels update the encoder. The published `dinov2_vitb14_pretrain.pth` checkpoint defines the backbone. Its acquired bytes are locked once, before all experiment splits, and the run manifest records their SHA-256 digest.

A linear layer maps each patch to width 256. Eight learned workspace tokens undergo four shared updates, each containing pre-normalized cross-attention to image tokens, self-attention, and a 256–1,024–256 GELU feed-forward block. Attention has four heads. Attention and residual dropout are 0.1. The first workspace token is norm bounded and supplies the decision state at each step. A bias-free four-way linear readout is spectrally normalized to operator norm at most one and uses temperature 0.2. The supplied executable module has exactly 1,253,376 trainable parameters. Width and depth sweeps change only the named dimensions and keep the backbone frozen.

N.2. Warm start and continuation

The task-only warm start runs for 500 updates. Each continuation runs for 1,000 updates from the identical warm-start weights, optimizer initialization policy, and random seed. AdamW uses learning rate 3×10^{-4} , weight decay 0.01, $\beta_1 = 0.9$, $\beta_2 = 0.999$, and $\epsilon = 10^{-8}$. The learning rate has 50 linear warmup updates followed by cosine decay to 3×10^{-6} within each phase. Global gradient clipping uses norm 1.0. The primary RELAY loss weight is fixed at $\lambda = 0.3$.

An update uses the complete original training panel: 350 examples in an outer fold, or 300 in an inner fold. Computation is chunked into microbatches of 25; the cross-entropy is normalized by the full panel size. The neural Gram loss also uses that same complete panel, not independently centered minibatches. Thus the configured microbatch size changes memory use rather than the objective. No label-preserving image augmentation is given a copied EEG target. The primary numerical comparison uses the original paired boards only; optional icon and option transformations are evaluated as separate visual checks.

Validation runs every 50 updates. Checkpoint selection minimizes inner-family validation NLL, with earlier update count breaking ties. The continuation budget is identical for compared objectives. Outer test metrics are computed once from the selected checkpoint. All ten initialization identifiers are retained in their original, unsorted order: 11, 29, 47, 71, 101, 137, 173, 211, 257, and 307.

N.3. Teacher arithmetic and memory

Teacher features, distinct-index moments, path scores, and eigenvalue thresholding use float64. Student activations use bfloat16, with Gram products and the loss accumulated in float32. Target matrices are stored with their exact index order, selected path, role membership, and pre/post-projection spectra. For $n = 350$, $T = 8$, one dense float64 $2,800 \times 2,800$ matrix occupies 62,720,000 bytes. The three fixed targets occupy 188,160,000 bytes before auxiliary buffers.

The exact feature-space formula in Eq. (S61) avoids storing a gradient graph for all N^2 entries. A two-pass implementation caches bounded states and computes the gradient with respect to these states on the complete panel. It then replays each microbatch with identical CPU/CUDA random-number states and backpropagates the corresponding vector–Jacobian product. Dropout masks must be replayed, not redrawn. Dense full-graph evaluation on a small panel checks equality of the loss and gradients. All positive target eigenvalues are retained; the primary implementation does not substitute a rank-truncated target.

N.4. Qwen2.5-VL replication

The second student is Qwen2.5-VL-3B-Instruct [5]. The language backbone has 36 residual layers and hidden width 2,048. We supervise the final prompt-position states after layers 9, 18, 27, and 36, before any answer token enters the context. The temporal path is selected among these four ordered depth positions. The kernel therefore uses the native 2,048-dimensional states, not a separately trained EEG projection head.

The image is encoded at the same 448×448 board resolution. The prompt is the fixed four-option instruction given above. Answer loss is computed over the four A–D response tokens at the first answer position, renormalized to a four-way distribution. The tokenizer manifest verifies that these four responses use the intended distinct single-token identifiers in this prompt context. No target answer is present in the supervised prompt-position hidden state.

LoRA is applied to `q_proj`, `k_proj`, `v_proj`, and `o_proj` in all 36 language layers, with rank 16, scale 32, and dropout 0.05. The visual tower, embedding table, normalizations, and language-model head are frozen. AdamW uses learning rate 2×10^{-5} , weight decay 0.01, and gradient norm 1.0. Warm start and continuation retain the 500/1,000-update budgets. Non-reentrant activation checkpointing, PyTorch SDPA, and microbatches of one image accumulate the full-panel objective; the autoregressive key/value cache is disabled during training. The paired task-only Qwen continuation uses identical LoRA modules and optimization settings. Table S9 reports both Qwen arms and the visual-shortcut controls.

N.5. Software and hardware configuration

The training environment fixes Python 3.11.11, PyTorch 2.6.0 with CUDA 12.4, Transformers 4.51.3, PEFT 0.15.2, NumPy 2.2.5, SciPy 1.15.2, and MNE 1.9.0. Eight NVIDIA A100 80-GB GPUs process independent fold–seed jobs; each fit uses one GPU. Teacher construction uses 16 CPU threads and float64 BLAS. A worker has 256 GB RAM and local NVMe storage for the token cache. The scalar probability-analysis and mathematical tests have their own recorded runtime environment and do not depend on the CUDA training image.

O. Baselines as Explicit Objectives

O.1. Three different zero-information endpoints

Task-only ERM optimizes cross-entropy with no kernel penalty. It is exactly the $\lambda = 0$ endpoint of the stated objective, including all cached features, optimizer settings, and seed identities. Its predictions are reused for the $\lambda = 0$ row rather than regenerated as another condition.

The *fixed-path anchor* uses $\pi = (1, 1, 2, 2, 3, 3, 4, 4)$ and target G_{π}^0 , with no EEG-dependent operation. The *EEG-*

selected-path anchor retains the three paths selected by RELAY but sets $\rho = 0$, so the targets are $G_{\pi_r}^0$. It isolates the additional value of the neural target after matching the EEG-selected path. Its selector computation and associated resource use are included in its end-to-end cost. The $\rho = 0$ row is exactly this selected-path anchor, not ERM.

O.2. Predictive and relational neural baselines

The raw-EEG baseline predicts the donor-mean feature vector at each selected time window from a two-layer readout of the student state. Its MSE is normalized by the number of valid feature coordinates. The ReAlnet-style adaptation uses an EEG-predictive auxiliary loss and the same visual student; the comparison reproduces this loss family rather than its original architecture and image benchmark [6]. The auxiliary-head control places this prediction head on a parallel, trainable readout, permitting neural fit without directly matching the answer state.

The static baseline time-averages the donor features before constructing a single problem kernel and matches the final decision-state Gram. The MRGS-style adaptation matches the pairwise conceptual geometry of that time-averaged teacher [7]. The time-wise baseline uses

$$\mathcal{L}_{\text{diag}} = \frac{1}{T} \sum_{t=1}^T \frac{\|G_{tt} - \tilde{K}_{tt}\|_F^2}{n^2}. \quad (\text{S80})$$

It applies no loss or gradient to off-diagonal blocks. The separate zero-offdiagonal-target control instead computes a full Frobenius loss against a block-diagonal teacher. These are different objectives and have separate condition identifiers.

The regularization weights for the raw-EEG, ReAlnet-style, and MRGS-style adaptations are selected from $\{0.03, 0.3, 3\}$ using seven inner-family folds and the two development seeds 11 and 29. Time-wise, static, and joint geometric controls use the same $\lambda = 0.3$ primary weight after their stated per-coordinate normalization; their λ sweeps accompany the main sensitivity study. This keeps the teacher-construction comparisons separate from arbitrary rescaling of the loss.

O.3. Component controls

The joint-unshrunk control sets $\rho = 1$ after the same path-compatible projection. Joint-fixedpath fixes the uniform path but retains independent calibration and teacher construction. Pooled-calibration reuses the selector donors for moment calibration while the third cohort still constructs the teacher. Constant-rho selects one weight from $\{0, 0.25, 0.5, 0.75, 1\}$ by the same seven-family, two-seed inner validation, then uses that scalar across all three rotations in the outer fit.

Full-space is the previous calibration rule: estimate mismatch and variance against K before imposing temporal compatibility. Its target is PSD projected in the centered

space and matched by the same repeated-state student. This contrasts calibration against the complete measured target with calibration against K_π . The main component table compares all seven alternatives directly with RELAY, including intervals for their incremental effects rather than only their individual accuracies.

O.4. Neural specificity and temporal coordinates

Matched-shuffle permutes problem–EEG correspondence within training-family and displayed-correct-option strata. One permutation is shared across donors, channels, and time windows, preserving a coherent trajectory for each reasigned problem. The entire selector, calibrator, and teacher computation is repeated with the altered pairing.

For the time-rotation control, independently sampled orthogonal matrices $O_t \in \mathbb{R}^{p \times p}$ act as $Z'_{s,t} = Z_{s,t} O_t$. The same O_t acts on every donor in that window. Therefore every within-window cross-donor kernel is unchanged, whereas cross-time blocks change. Orthogonal matrices are generated by QR factorization with a fixed diagonal-sign convention. This control targets cross-time coordinates more specifically than reversing the order of windows.

Label-kernel, RT-only, mask-only, gaze-only, EOG-only, and random-PSD teachers receive their named side information only. Random-PSD targets match the training teacher’s trace and retained rank but use an independent random orthonormal basis in the centered path-compatible range. Naive-self-Gram replaces the distinct-donor estimator with the Gram of the donor mean. The human-correct and human-incorrect comparisons match the original problems and within-family response-time quintiles, equalize trial counts by subsampling, and repeat the complete teacher construction. Table S8 gives the complete 15-test exploratory family; Table S10 contains the additional matched-outcome and modality comparisons.

P. Statistical Estimands, Intervals, and Multiplicity

P.1. Accuracy and proper-score summaries

For method m , seed s , and parent problem i , let $p_{msi} \in [0, 1]^4$ be the saved probability vector, y_i the displayed correct option, and $c_{msi} = 1\{\arg \max_k p_{msik} = y_i\}$. Ties use the first displayed option. The primary statistic is

$$A_m = \frac{100}{8} \sum_{f=1}^8 \frac{1}{50} \sum_{i \in f} \frac{1}{10} \sum_s c_{msi}. \quad (\text{S81})$$

The equal family sizes make this identical to the average of all 4,000 decisions. Seed SD is the sample standard deviation of ten per-seed accuracies, with denominator nine. Seed IDs are never replaced by rank order.

NLL is $-\log p_{msi, y_i}$, with a numerical floor of 10^{-12} . The multiclass Brier score is $\sum_{k=1}^4 (p_{msik} - 1\{k = y_i\})^2$, without an additional division by four. ECE uses 15 equal-width confidence bins over $[0, 1]$ and the maximum predicted probability. Empty bins contribute zero. Probabilities and labels, not rounded accuracy summaries, define all three metrics. Calibration is reported separately from discrimination: a gain in accuracy or NLL does not imply that every ECE comparison improves.

P.2. Paired problem-level inference

For two methods, the parent-level effect is $d_i = 100 \cdot 10^{-1} \sum_s (c_{1si} - c_{0si})$. The primary interval uses 10,000 stratified bootstrap draws, sampling 50 original problems with replacement in each of the eight fixed families. The two methods share each draw. The interval is conditional on the fitted initializations and the eight family identities; donors and kernel entries are not substituted for additional behavioral test cases.

Paired sign randomization uses the same d_i and 10,000 independent sign vectors. With observed $t = |\bar{d}|$, the two-sided Monte Carlo value is

$$p = \frac{1 + \#\{b : |n^{-1} \sum_i \epsilon_{bi} d_i| \geq t\}}{10,001}. \quad (\text{S82})$$

The null is sign exchangeability of the paired original-problem effects. The procedure does not permute method labels independently for each of the ten correlated seed repetitions of a problem.

A second interval resamples both the ten seed identities and the original problems within each fixed family. Every replicate uses the same seed and item draws in the paired methods. Table S3 reports these crossed intervals. In particular, the RELAY–ERM interval is [6.25, 11.22] pp and the Qwen pair interval is [5.97, 11.88] pp. This makes initialization uncertainty visible alongside the primary conditional interval.

P.3. Three explicitly enumerated test families

The four confirmatory comparisons are RELAY versus ERM, fixed-path anchor, selected-path anchor, and matched-shuffle. Holm correction is applied once to those four raw values. The component family contains exactly seven comparisons: time-wise, joint-unshrunk, joint-fixedpath, pooled-calibration, constant-rho, full-space, and time-rotation. Holm correction is applied separately to those seven values. The 15 exploratory comparisons in Table S8 use Benjamini–Hochberg adjustment. No unseen comparison is included in any denominator.

The multiplicity ledger records the full-precision raw and adjusted values and the named correction family. Confidence intervals shown next to adjusted tests are ordinary paired 95% intervals, not simultaneous intervals. The family of

hypotheses, direction of each contrast, and inferential unit therefore remain identifiable from each table. The Qwen replication, gaze increment, and intervention interaction are named secondary endpoints with their own paired intervals.

P.4. A matching full-pipeline correspondence test

The correspondence test uses the two prespecified development seeds 11 and 29 in all eight outer folds, for both the observed statistic and every null replicate. Its observed gain is 5.75 pp. There are 199 correspondence permutations, hence $199 \times 8 \times 2 = 3,184$ null student fits. The plus-one value is $(1+e)/200$, with $e = 0$ exceedances and $p = 0.005$.

Each null permutation repeats path selection, path projection, calibration, teacher construction, and student continuation. The permutation is shared across donors and windows, within training-family/correct-option strata. RELAY’s primary $\lambda = 0.3$, visual configuration, and task-only warm-start procedure are fixed before this test, so no EEG-dependent hyperparameter is silently frozen from the observed pairing. Task-only warm starts, original DINO tokens, and donor-local preprocessing can be reused when unaffected by the correspondence permutation. Neither selected paths nor neural targets are cached across permutations.

The 10-seed confirmatory statistic is not substituted for this two-seed observed statistic. The run index names all 3,184 jobs and their fold, seed, and required stages. Table S21 accounts for this workload separately from ordinary comparisons.

P.5. Independent-risk uncertainty

Risk differences against the two-reference kernel use the unprojected cross-donor estimate, with the same reference for the compared targets. No unknown $\|K\|_N^2$ term is assigned a numerical value. The 12+8 absolute audit uses the distinct-donor fourth-order moment from all eight references. Negative finite-sample audit estimates, when they occur, are retained rather than clipped into apparently smaller uncertainty.

For absolute-audit uncertainty, first average the target-specific estimates within each named teacher condition, keeping shared reference identities aligned across all folds and seeds. Then delete one reference donor at a time and recompute both the raw reference kernel and its distinct-index moment. The jackknife standard error uses eight deleted-donor values, not 240 nominally independent target records. Intervals use a t_7 critical value. The 240-record quantiles in Table S12 describe calibration variation and are not confidence intervals for independent replications.

Q. Sensor Budgets and Decision-State Interventions

Q.1. Nested physical montages

The physical channel budgets are 4, 8, 16, 32, and 64. Twenty deterministic montage seeds generate nested sets from standard BioSemi coordinates. Each set begins with a balanced anterior/posterior and left/right choice, then adds the channel farthest from the retained set in scalp-coordinate distance, breaking ties by channel name. The full montage is unique; its repeated identity in montage-indexed arrays is not treated as 20 independent physical layouts.

Channels are removed before referencing, ocular regression, coefficient scaling, and feature extraction. Each budget refits all teacher stages with the same original training items. Montage means and confidence intervals in Table S18 average seeds first, then use 3,000 paired problem/montage resamples. The direct 64-minus-32 contrast is -0.02 pp with interval $[-0.72, 0.64]$. The interval quantifies uncertainty in the observed change rather than declaring an exactly zero benefit.

The duplication control maps a bounded 16-channel feature z to $(z, z)/\sqrt{2}$. Its Gram is exactly unchanged, making identical predictions an algebraic endpoint of the same run. The independent-noise extension adds a separately scaled nuisance block and refits the teacher. The two controls distinguish an increase in array width from additional useful measurements.

Q.2. Predicting a withheld sensor budget

Only the 4-, 8-, 16-, and 32-channel results fit $a - b \exp(-C/\tau)$, constrained to $25 \leq a \leq 100$, $0 \leq b \leq 60$, and $1 \leq \tau \leq 150$. The 64-channel value is revealed after this fit. The central prediction is 73.12%, with a bootstrap refit interval of $[69.28, 77.36]\%$. Two hundred refits use shared problem/montage draws; no arbitrary independent Gaussian noise is added to the fitted curve.

The forecast pertains to increasing the retained EEG montage under this acquisition and feature protocol. The recorded EEG+gaze comparison measures a different, actually available modality combination. The linear-Gaussian innovation analysis in Section H supplies a mathematical vocabulary for complementary sensors; the empirical channel curve is not used to assign an fMRI or ECG performance value.

Q.3. Aligned, random, and readout subspaces

The aligned intervention subspace is fitted only from training-problem student states and the corresponding neural coordinates. A rank-eight SVD of their cross-covariance gives the aligned student basis. At the final decision state, deletion subtracts its orthogonal projection. Random comparison subspaces have the same rank and match the mean removed

training-state energy within 2%; the accepted subspaces are frozen before testing and averaged over 20 random subspace seeds. This comparison measures the extra importance of neural-aligned directions at equal dimensional and energy cost.

For the linear readout V , the probability-sensitive matrix is $J = (I_4 - \mathbf{1}\mathbf{1}^\top/4)V$, with rank at most three. Full deletion of its row space makes the bias-free output probabilities uniform, yielding the four-choice chance endpoint; deletion of its nullspace leaves all probabilities exactly unchanged. These are algebraic controls, not evidence that every aligned feature influences the answer. Table S20 separates these endpoints from the matched aligned/random comparison.

The paired interaction is formed at the parent-problem level using the four aligned/random intervention conditions, as derived in Section K. It is 6.05 pp with interval [4.75, 7.33]. State-patch and skipped-update controls retain the same output head and perform no test-time optimization. Case selection and trace export are specified in Section S.

R. Extended Results and Calibration Diagnostics

The main comparison aggregates the same 400 original problems and ten initializations throughout. Table S2 gives the complete four-test confirmatory family; Table S3 adds seed uncertainty to the paired problem resampling. Table S4 decomposes the average by the held pattern, and Table S5 retains the unsorted initialization identifiers. Their family and seed means reproduce the main accuracy exactly before rounding.

Table S2. The complete confirmatory multiplicity family. The same unrounded probability records define every test.

Comparator	Gain	95% CI	Raw p	Holm p
ERM	8.72	[7.38, 10.12]	0.0001	0.0004
Anchor-fixed	8.47	[7.17, 9.82]	0.0001	0.0004
Anchor-selected	8.10	[6.72, 9.53]	0.0001	0.0004
Matched-shuffle	7.50	[6.20, 8.78]	0.0001	0.0004

Table S3. Crossed seed–problem bootstrap. Each replicate resamples ten seed identities and 50 parent problems within each fixed family, jointly across compared methods.

Paired comparison	Gain (pp)	95% interval
RELAY – ERM	8.72	[6.25, 11.22]
RELAY – Anchor-fixed	8.47	[6.53, 10.47]
RELAY – Anchor-selected	8.10	[5.80, 10.65]
RELAY – Matched-shuffle	7.50	[5.10, 10.05]
RELAY – Constant-rho	4.92	[2.67, 7.35]
RELAY – Full-space	0.62	[-0.80, 2.03]
Qwen-RELAY – Qwen-ERM	8.85	[5.97, 11.88]

Table S4. All eight held-out pattern families, with 50 parent problems per family.

Pattern	ERM	RELAY	Gain	95% CI
AAABAAAB	63.00	76.00	13.00	[8.60, 17.80]
ABABCD CD	71.60	77.40	5.80	[2.80, 9.40]
ABBAABBA	65.80	75.80	10.00	[6.00, 14.60]
ABBACDDC	70.00	82.00	12.00	[7.80, 16.60]
ABBCABBC	59.20	67.60	8.40	[4.60, 12.40]
ABCAABCA	74.60	79.00	4.40	[1.80, 7.40]
ABCDDCBA	52.00	62.00	10.00	[6.00, 14.60]
ABCDEEDC	45.80	52.00	6.20	[3.20, 9.60]

Table S5. Unsorted random seed identifiers and paired results. Rows are indexed by initialization, not by accuracy.

Seed	ERM	RELAY	Gain
11	65.00	67.75	2.75
29	59.25	68.00	8.75
47	63.75	76.50	12.75
71	60.75	66.75	6.00
101	58.50	72.25	13.75
137	65.00	72.75	7.75
173	61.25	72.50	11.25
211	66.75	75.25	8.50
257	63.50	71.75	8.25
307	63.75	71.25	7.50

Table S6. All 4,000 paired item–initialization decisions. The net correction count divided by 4,000 equals the primary accuracy gain.

ERM outcome	RELAY incorrect	RELAY correct
Incorrect	1114	376
Correct	27	2483

The paired decisions contain 376 corrections and 27 regressions relative to ERM (Table S6). Their difference divided by 4,000 reproduces the 8.72-pp mean gain. Seed-to-seed gains have SD 3.23 pp, and the ERM–RELAY seed correlation is 0.41. The family and initialization views therefore expose different sources of variability instead of presenting a single nearly constant improvement on every run.

R.1. Component and alternative-information controls

Table S7 repeats the component contrasts with their raw test values and correction family, so the effect of each teacher-design choice is directly visible. Table S8 gives the full 15-comparison alternative-information set. ReAlnet-style and MRGS-style refer to matched objective adaptations on this task. Their scores do not overwrite or restate the original papers’ reported performance.

The full-space comparison is the closest numerical alternative to the revised method. Its smaller paired difference places the more distinctive contribution in the corrected estimand and independently checkable risk, while the stronger temporal, unshrunk, and constant-weight comparisons establish the practical value of the larger teacher-construction design. These conclusions use the paired component intervals, not an ordering of independently rounded averages.

Table S7. Complete seven-comparison component family: RELAY minus each named control.

Control	Gain (pp)	95% CI	Raw p	Holm p
Time-wise	5.75	[4.58, 6.95]	0.0001	0.0007
Joint-unshrunk	5.42	[4.20, 6.70]	0.0001	0.0007
Joint-fixedpath	3.05	[2.02, 4.12]	0.0001	0.0007
Pooled-calibration	5.60	[4.40, 6.85]	0.0001	0.0007
Constant-rho	4.92	[3.80, 6.10]	0.0001	0.0007
Full-space	0.62	[-0.33, 1.57]	0.2284	0.2284
Time-rotation	4.92	[3.75, 6.08]	0.0001	0.0007

Table S8. Complete exploratory control family (15 tests), using Benjamini–Hochberg adjustment. Neural side information is training-only.

Method	Acc.	NLL	Brier	ECE	Gain	Raw p	BH q
Raw-EEG	64.50	1.008	0.475	8.75	1.75	0.0010	0.0015
Static	64.62	0.992	0.470	8.99	1.88	0.0007	0.0012
Label-kernel	63.65	1.037	0.486	9.47	0.90	0.1235	0.1323
RT-only	61.40	1.090	0.510	9.23	-1.35	0.0077	0.0096
Mask-only	59.65	1.138	0.529	9.75	-3.10	0.0001	0.0002
Gaze-only	64.58	1.015	0.475	8.72	1.82	0.0006	0.0011
EOG-only	63.38	1.031	0.485	9.45	0.62	0.2491	0.2491
Random-PSD	64.38	1.010	0.474	9.52	1.62	0.0026	0.0035
Auxiliary-head	66.07	0.969	0.458	8.66	3.33	0.0001	0.0002
Naive-self-Gram	67.65	0.950	0.443	8.15	4.90	0.0001	0.0002
Correct-trials	72.28	0.815	0.386	8.96	9.53	0.0001	0.0002
Incorrect-trials	65.42	0.981	0.462	9.29	2.67	0.0001	0.0002
Uncorrected-EEG	67.12	0.948	0.446	9.36	4.38	0.0001	0.0002
ReAlnet-style	63.80	1.038	0.486	8.56	1.05	0.0630	0.0727
MRGS-style	66.77	0.943	0.445	9.39	4.03	0.0001	0.0002

Table S9. Architecture and visual-shortcut controls. All inputs are images; candidate-only masks the sequence.

Configuration	Acc.	SD	NLL	ECE
DINO-linear	57.75	4.59	1.178	10.49
Candidate-only	24.62	1.82	2.301	27.07
ERM	62.75	2.69	1.049	9.80
RELAY	71.47	3.18	0.827	9.46
Qwen-ERM	57.83	3.51	1.175	9.84
Qwen-RELAY	66.67	4.52	0.961	8.46

Table S10. Redundancy, recorded modality, and human-outcome controls. The duplicated representation is scaled to preserve exactly the original Gram.

Control	Acc.	Seed SD	NLL
channels-16	66.70	3.94	0.959
Duplicated-16	66.70	3.94	0.959
Noise-added-16	65.60	4.09	0.983
Gaze-only	64.58	3.16	1.015
EEG+gaze	72.50	3.86	0.808
Correct-trials	72.28	3.61	0.815
Incorrect-trials	65.42	3.05	0.981

R.2. Joint family–donor transfer

Every row of Table S11 is a joint holdout: the evaluated family and reference donors are absent from the associated model’s training and teacher construction. The selected path is frozen from the remaining donor and family data. CKA is computed from centered student states and the mean of the two held-donor feature panels on the same 50 held problems. Unlike the risk audit, this descriptive CKA statistic uses an ordinary PSD reference Gram; it is not called an unbiased risk estimate. Across the same joint units, task accuracy changes from 61.85% to 72.72%, while CKA changes from 0.103 to 0.163.

Table S11. Joint family–donor holdout. Each row aggregates the same eight held-family tests and ten initializations; the two displayed reference donors are absent from every teacher stage.

Held donors	ERM acc.	RELAY acc.	Gain	ERM CKA	RELAY CKA
D01, D02	62.75	71.47	8.72	0.103	0.166
D03, D04	63.00	73.55	10.55	0.103	0.164
D05, D06	61.92	73.58	11.65	0.102	0.161
D07, D08	61.83	74.55	12.72	0.105	0.163
D09, D10	60.77	72.12	11.35	0.103	0.162
D11, D12	62.73	70.62	7.90	0.103	0.161
D13, D14	61.38	72.10	10.72	0.105	0.163
D15, D16	61.38	73.22	11.85	0.102	0.160
D17, D18	60.73	72.75	12.02	0.103	0.165
D19, D20	62.05	73.25	11.20	0.105	0.165

R.3. Independent risk and moment variation

Table S12 reports the actual row-wise quantities used by the teacher construction. Moments are estimated before the PSD projection. Variance-component clipping and anchor-error clipping are recorded as separate implementation events, rather than treated as evidence of successful transfer. The mean plug-in weight is 0.665; it is an average of ratios. Substituting the grand means 0.736 and 0.364 into a ratio computes a different summary.

The primary two-reference risk differences favor RELAY by -0.664 and -0.034 units of 10^{-6} relative to the selected-path anchor and projected unshrunk teacher. The independent 12+8 audit in Table S13 supplies absolute compatible-space estimates and reference-donor uncertainty. Figure S1 evaluates frozen interpolation weights using independent reference features, rather than drawing a parabola from the same fitted a and v values. The oracle identity predicts the expectation before projection; finite-reference fluctuations and the PSD step affect the observed curve.

Table S12. All 240 calibration records. The final weight is the mean of record-level ratios, not a ratio of grand means.

Quantity	Mean	10th pct.	90th pct.
a	0.736	0.572	0.880
v	0.364	0.340	0.389
rho	0.665	0.611	0.715
delta_anchor	-0.664	-1.063	-0.346
delta_unshrunk	-0.034	-0.186	0.162
path_floor	0.237	0.217	0.258
rank8_tail	0.557	0.422	0.736

Table S13. Independent-reference risk audit. Values are 10^6 times squared normalized kernel error. Relative audit: 18 training plus two reference donors. Absolute audit: 12 training plus eight reference donors. Intervals use a reference-donor jackknife and a t_7 critical value, after aggregating targets.

Absolute audit	Risk	Donor SE	95% interval
Anchor	0.753	0.017	[0.712, 0.794]
RELAY	0.300	0.023	[0.246, 0.355]
Unshrunk	0.491	0.028	[0.426, 0.557]




Figure S1. **Independent-reference risk along the weight path.** The vertical axis is risk difference from the anchor, estimated from raw held-reference features at each fixed weight. The line is the record mean and the band spans the 10th–90th record percentiles, not a confidence interval. Their extrema need not equal a plug-in weight derived from calibration donors.

R.4. Loss, capacity, and information sensitivity

The $\lambda = 0$ row in Table S14 is exactly ERM; $\rho = 0$ is the selected-path anchor. The two controls have different objectives and are not relabeled versions of one another. Width and computation depth in Table S15 are sensitivities of the complete model. They are not claimed as matched estimates of the neural effect at each new architecture.

Table S16 varies the number of teacher donors and the number of EEG-paired training problems, keeping the full task-label set available. For fewer paired items, a fixed stratified subset supplies the neural Gram, while cross-entropy still uses all 350 training problems. Table S17 varies timing, feature width, and the nuisance definition. Table S18 reports the retained physical montage with its paired resampling intervals.

Table S14. Loss-weight sensitivity; $\lambda = 0$ reuses the task-only run exactly.

Configuration	Acc.	SD	NLL	ECE
lambda-0	62.75	2.69	1.049	9.80
lambda-.01	66.92	2.70	0.940	8.96
lambda-.03	66.80	3.18	0.938	9.23
lambda-.1	69.83	4.31	0.868	8.75
lambda-.3	71.47	3.18	0.827	9.46
lambda-1	70.33	4.89	0.863	8.80
lambda-3	65.65	3.07	0.981	8.93

Table S15. Workspace width and computation depth. These are sensitivity analyses of the complete method, not paired component-effect estimates.

Configuration	Acc.	SD	NLL	ECE
width-64	68.12	3.79	0.918	8.93
width-128	69.85	2.66	0.872	9.25
width-256	71.47	3.18	0.827	9.46
width-512	70.83	3.27	0.854	9.14
steps-2	66.65	2.20	0.956	8.88
steps-4	71.47	3.18	0.827	9.46
steps-6	72.25	5.19	0.833	7.54
steps-8	74.25	3.74	0.761	9.70

Table S16. Training information budgets; item counts are EEG-paired parent problems, while all task labels remain available.

Configuration	Acc.	SD	NLL	ECE
donors-12	71.73	3.74	0.842	7.65
donors-15	69.92	4.78	0.866	8.94
donors-18	71.47	3.18	0.827	9.46
items-35	66.17	4.20	0.974	8.80
items-70	69.12	4.49	0.884	8.78
items-140	69.30	3.68	0.895	8.45
items-245	70.83	3.78	0.844	9.20
items-350	71.47	3.18	0.827	9.46

Table S17. Timing, feature, and nuisance sensitivity. The primary representation uses eight windows and four DCT coefficients.

Configuration	Acc.	SD	NLL	ECE
bins-4	69.45	2.39	0.873	9.28
bins-6	69.15	4.24	0.891	8.94
bins-8	71.47	3.18	0.827	9.46
bins-12	71.38	4.35	0.826	9.18
DCT-2	71.00	2.66	0.851	8.50
DCT-4	71.47	3.18	0.827	9.46
DCT-8	70.12	3.91	0.869	8.58
guard-.15	70.60	3.73	0.849	8.97
guard-.30	71.47	3.18	0.827	9.46
guard-.45	69.62	4.64	0.873	9.15
stimulus-clock	70.60	3.82	0.866	7.90
No-EOG-regression	68.70	2.50	0.895	9.23
Nuisance-visual	70.65	3.26	0.832	9.10
Nuisance-label	69.75	4.37	0.891	8.37

Table S18. Retained-channel results. Intervals resample parent problems and nested montage identities; the 64-channel montage is unique.

Channels	Accuracy	95% interval
4	62.41	[58.46, 66.35]
8	65.48	[61.66, 69.23]
16	67.80	[64.09, 71.53]
32	71.49	[67.92, 75.14]
64	71.47	[67.87, 75.12]

Table S19. Conditional paired-variance design sensitivity. The normal approximation uses the parent-level paired SD with seeds averaged; it is a design calculation, not post-hoc evidence for the observed alternative.

Problems	SE (pp)	80% MDE (pp)	Power at 2 pp	4 pp	6 pp
100	1.44	4.04	0.283	0.792	0.986
200	1.02	2.86	0.500	0.975	1.000
400	0.72	2.02	0.792	1.000	1.000

S. Visual Atlas: Kernels, Decisions, and Interventions

S.1. A joint kernel with its required repeated blocks

Figure S2 uses one selected path and the same displayed problems in all panels. The anchor and student matrices are computed from $F_\pi F_\pi^\top$, not generated as independent heatmaps. If $\pi(t) = \pi(u)$, their complete block rows and columns are equal. The path-compatible teacher also lies in the range of W_π , so its repeated groups follow the same temporal assignment. This property is checked numerically before exporting the images.

The saved display object includes the four underlying state arrays, the path, the centered stacked matrices, and the resulting Grams. The color scale is shared across all three images. Comparing colors without this shared scale would conflate signal strength with plot normalization. The displayed compact panel is a diagnostic of the defined matrices; quantitative teacher-risk summaries use their prescribed normalized inner products.




Figure S2. **Joint time-problem kernels.** Block equality is inherited from the saved path and state arrays. The matrices share one display scale, problem ordering, and temporal assignment.

S.2. Input, answer distribution, and a difficult case

Figure S3 shows the complete input format, including candidate answers. The family specification belongs to the analyst's label file and is not rendered in the model input. Figure S4 shows two stored cases at seed 11: the first correctly answered ABCAABCA item and the first incorrectly answered ABCDEEDC item in stable parent-ID order. The stored case record includes the original problem ID, seed, correct option, and all four probability vectors. Each colored line is an option probability at one of the four state-readout locations. A rising correct-option probability can be interrupted or reversed; no monotonicity is imposed on the intermediate states.




Figure S3. **The complete board and four choices.** Simplified icons illustrate the released ABCAABCA pattern. Option D completes this example. The pattern string is absent from the board itself.




Figure S4. **Decision-state traces.** Final vectors equal the corresponding saved prediction records exactly. Intermediate probabilities are read from the same four-way output map at each state.

S.3. Interventions with shared paired units

Table S20 supplies unmodified-reference drops and paired intervals for each intervention. The aligned and matched-random deletions target equally sized state subspaces. Their difference-in-differences is formed jointly across both methods and all four conditions, giving 6.05 pp with interval [4.75, 7.33]. Readout row-space deletion provides the uniform-probability endpoint; readout nullspace deletion preserves the output exactly. These checks distinguish a decision-relevant empirical comparison from an architecture-wide claim that every supervised coordinate must be used.

Table S20. Decision-state interventions. Drops are paired against each unmodified model; nullspace deletion is applied at the final linear readout.

Intervention	ERM acc.	Drop	95% CI	RELAY acc.	Drop	95% CI
aligned	52.40	10.35	[9.00, 11.78]	55.10	16.38	[14.47, 18.30]
random	57.00	5.75	[4.85, 6.70]	65.75	5.72	[4.83, 6.65]
row-space	25.00	37.75	[32.23, 43.38]	25.00	46.48	[41.15, 51.90]
nullspace	62.75	0.00	[0.00, 0.00]	71.47	0.00	[0.00, 0.00]
patch	54.83	7.92	[6.80, 9.07]	60.12	11.35	[9.93, 12.85]
skip-step	54.30	8.45	[7.30, 9.65]	61.25	10.22	[8.85, 11.62]




Figure S5. **Aligned versus matched-random deletion.** Paired accuracy drops are computed from the same item–seed records; the interaction compares the two drop differences jointly.

T. Computation, Records, and Reproduction

T.1. Stage accounting and the correspondence workload

Table S21 enumerates the principal run allocations. Primary and exploratory suites retain their own run identities; named controls repeated in a separate suite consume that suite’s allocation. The montage allocation counts four smaller budgets, 20 montage seeds, eight folds, and ten initializations; the unique 64-channel condition reuses RELAY. Additional joint holdouts count the nine donor pairs beyond the primary reference pair. The ledger separates optimization GPU hours from serial stage time, so CPU teacher construction is not silently charged as GPU optimization.

The Qwen allocation uses 900 minutes of student optimization per fit, including its full-panel warm start and continuation. With approximately 300 language/visual tokens per board, 525,000 board presentations, and a three-billion-parameter dense backbone, a $6ND$ operation-count approximation is 2.8×10^{18} operations before accounting for checkpointing and kernel efficiency. Full-panel gradient accumulation therefore has a substantially different budget from a short minibatch adaptation.

The correspondence allocation is exactly 3,184 null fits. The two-seed observed statistic is already present in the primary prediction record. Its run manifest names the required selector, projection, calibration, teacher, and continuation stages. Cache entries are keyed by source snapshot, outer family, donor identity, preprocessing configuration, seed, and item order. A changed EEG correspondence invalidates paths, calibration records, and targets even when the waveform-feature cache remains reusable.

Table S21. Configuration-level computation ledger. GPU hours count student optimization; serial-stage hours additionally count preprocessing/cache and teacher construction, which may run on CPUs.

Experiment	Fits	GPU h	Serial stage h
Primary comparison set	720	129.6	174.0
Exploratory controls	1200	216.0	290.0
Core extra conditions	400	72.0	96.7
Nested montage sweep	6400	1152.0	1546.7
Additional joint holdouts	1440	259.2	348.0
Qwen pair	160	2400.0	2412.8
Correspondence null	3184	573.1	652.7
Inner baseline selection	1568	282.2	378.9
Total	15072	5084.2	5899.8

T.2. Record schemas and numerical checks

The probability schema contains method, outer family, initialization seed, original problem ID, displayed-correct-option index, and four normalized probabilities. The calibration schema additionally names the selector, calibrator, teacher, and reference identities, path, raw moments, clipped components, weight, and target-spectrum diagnostics. The multiplicity ledger records all raw values at full precision before adjustment. Separate resource and job manifests record suite membership and matching observation/null protocols.

The executable reference implements the original U-statistic, exact distinct-index moments, PSD projection, exhaustive and branch-and-bound path search, the path-compatible teacher, both independent-risk estimators, the diagonal-only loss gradient, and centered-readout projectors. Algebraic tests check each object separately from the visual training system. Probability checks verify normalization, displayed answer mapping, the endpoint aliases, the paired correction count, and all reported summary identities. The kernel display includes explicit repeated-block assertions; the final trace vectors are matched to the probability file.

The core checks comprise 12 original tests and nine revision tests. The supplied PyTorch workspace separately verifies the stated trainable-parameter count and output tensor dimensions. These checks exercise the mathematical and reporting implementations. Dataset ingestion, pretrained checkpoint acquisition, and end-to-end model fitting are distinct executable stages with their own manifests.

T.3. The role of the auxiliary theory

The primary result is the path-compatible teacher-risk identity and its independent audit. The kernel-ridge stability analysis concerns its specified linear predictor, and the rank-floor result concerns best approximation of a finite trajectory kernel. The Gaussian sensor model concerns conditional information under its measurement assumptions. They are auxiliary analyses that explain different design questions; the nonlinear student's task accuracy is evaluated through its saved predictions and paired controls.

References

- [1] M. Ostrow, A. Eisen, L. Kozachkov, and I. Fiete, “Beyond geometry: Comparing the temporal structure of computation in neural circuits with dynamical similarity analysis,” *arXiv preprint arXiv:2306.10168*, 2023. [Online]. Available: <https://arxiv.org/abs/2306.10168> 8
- [2] P. Godara, P. S. Tay, and M. G. Mattar, “Beyond DSA: Conjugacy-based comparison of dynamical systems,” *arXiv preprint arXiv:2607.04493*, 2026. [Online]. Available: <https://arxiv.org/abs/2607.04493> 8
- [3] C. O. M. Pinier and M. Nunez, “Large language models show signs of alignment with human neurocognition during abstract reasoning,” University of Amsterdam, version 5, 2026, cC BY 4.0. [Online]. Available: <https://doi.org/10.21942/uva.29573534.v5> 13
- [4] C. Pinier, S. Acuña Vargas, M. Steeghs-Turchina, D. Matzke, C. E. Stevenson, and M. D. Nunez, “Large language models show signs of alignment with human neurocognition during abstract reasoning,” *arXiv preprint arXiv:2508.10057*, 2025. [Online]. Available: <https://arxiv.org/abs/2508.10057> 13
- [5] Qwen Team, “Qwen2.5-VL-3B-Instruct: Model configuration,” 2025. [Online]. Available: <https://huggingface.co/Qwen/Qwen2.5-VL-3B-Instruct/blob/main/config.json> 16
- [6] Z. Lu, Y. Wang, and J. D. Golomb, “Achieving more human brain-like vision via human EEG representational alignment,” *Communications Biology*, vol. 9, p. 463, 2026. 16
- [7] J. Chen, Y. Qi, Y. Wang, and G. Pan, “Human-like cognitive generalization for large models via mental representation-guided supervision,” *Nature Communications*, vol. 17, p. 4709, 2026, related preprint: arXiv:2505.09085, 2025. 16