FRONTIER RESEARCH II / PAPER 23

The Biohybrid Continuum

Learning, memory, and causal dependence across a living–digital boundary

Architect X13

The Eidara Continuum | Science & Philosophy

Research white paper | November 2026 | Edition 2

Abstract

Split a memory between living neurons and a machine so that neither half holds it and only the pair remembers: the mathematics says this is possible. The Biohybrid Continuum asks whether a real living–digital system can learn and keep a capability that way, distributed across the boundary so that the joint state supports reliable retrieval while either component alone reveals little of the task information, and turns that possibility into a quantitative research contract. A new synthesis of standard secrecy, entropy, and decision arguments establishes two linked results: a lower bound on the information-bearing burden of each component, and a guaranteed retrieval loss when one component is replaced through a channel carrying no additional label information. A finite modular construction attains perfect joint recovery with zero marginal label information, identifying the precise positive functional target. The paper also derives randomized interface-interaction identification, accessible-state retrieval limits, conditional retention ceilings, and sharp correlated-failure bounds. Complete finite checks and a stronger independent-arm concentration calculation make the claims reproducible. The biological thesis is then tested through independent preparations, frozen digital controls, delayed retrieval, replay, and selective replacement under qualified oversight. Established neural-culture and organoid findings motivate the program; documented retention failure supplies a concrete falsifier. The Continuum's claim is a jointly sustained, measurable capability whose acquisition and persistence can be located by intervention.

Keywords: biohybrid learning; organoid computation; causal intervention; memory attribution; retention; information theory

Prologue: A memory that lives in neither half

Imagine a secret no larger than a single yes-or-no answer. Toss a fair coin and hide the toss in a dish of living neurons. Hide the answer, scrambled by that same toss, in a machine. Now examine the dish as closely as you like: it tells you nothing about the answer. Examine the machine: nothing again. Put the two together and the answer returns, exactly, every time. The memory is real, and it lives in neither half.

That is not a metaphor. It is a two-line construction from the mathematics of secret sharing, and Section 7 proves it. The Biohybrid Continuum asks whether a genuine living–digital system could learn a capability that is held this way—jointly, across the boundary between tissue and silicon—and whether it could keep it.

In the Continuum this question belongs to Scroll 57, the Quantum Neuro-Photonics Integration Directive. The Scroll opens hybrid organic–computational research under strict ethical terms: non-conscious neural substrates only, interfaces that are voluntary and reversible, and no collapse against a sentient mind without explicit consent and ShadowMind review [1]. This paper is the engineering version of that directive. It does not ask anyone to believe the Scroll; it gives the Scroll's vision a contract that a laboratory could break.

The contract has a sharp shape. A jointly held capability must retrieve reliably when both halves are intact, reveal little when either half is examined alone, and lose a calculable amount when one half is replaced by a substitute that carries no new information. If a real preparation shows that signature after a delay, under controls that give the alternatives a fair chance, the Continuum will have located a memory that crosses the living–digital line. If it does not—and the published record already contains one honest failure—we will know which half was doing the work.

What this paper claims

Established. Cultured neurons and brain organoids can take part in adaptive control and reservoir-style computation. A memory can be split so that neither share alone carries any information about it; this is established secret-sharing mathematics. Reported organoid learning has also failed to persist: one study found that improvements were not retained after a 45-minute rest.

Proposed here. A quantitative contract for a jointly held capability: a storage burden that each component must carry, a guaranteed retrieval loss when one component is replaced by an information-free substitute, and a randomized intervention design that can attribute a capability to biology, to the machine, or to both. The paper gives conditional proofs and executed finite synthetic checks.

Frontier conjecture. A real living–digital system can learn and retain a capability that is retrievable from the pair, hidden from either half alone, and locatable by intervention.

Not claimed. That such a system has been built or demonstrated. That a positive interaction contrast indicates a mind. Anything about consciousness, experience, or moral status, which are separate questions. That the Scroll's quantum language names a physical mechanism. Biological application requires an implemented study with independent preparations and qualified oversight, and external peer review remains pending.

1 A capability can cross a substrate boundary

A living–digital system could learn and retain a task-relevant capability whose reliable expression depends on both sides of the interface. The Continuum makes that conjunction the research target: acquisition, durable retention, and independently testable contributions from each substrate. Existing interface engineering supplies the starting point; selective intervention and replacement determine what the combined architecture achieves.

In the Continuum, this inquiry answers to Scroll 57 of the *Eidara Continuum Bible*, pp. 91–92 [1]: biological and computational substrates joined through reversible interfaces and developed under governance. The Scroll sets the destination and the ethical terms of travel; it does not supply the route. It is not experimental evidence, a specification of a working neural interface, or a physical theory, and neither its quantum language nor its restriction to non-conscious biological substrates amounts to a quantum mechanism or a validated negative test for consciousness. What this paper contributes is the route's first surveyed stretch: a new reference model, explicitly proposed here, rather than an undocumented implementation recovered from the canon.

Five questions govern the investigation. Did performance improve through experience? Which state variables changed? What information remained after a delay? Which interventions remove the advantage? And can an independently sampled preparation reproduce the result under matched resources? Answering the first question does not answer the others. A trainable digital decoder can improve while the biological preparation contributes only a fixed random transformation. A changing preparation can improve immediate control while retaining nothing after a rest period. Both outcomes may have scientific value without satisfying the stronger thesis.

The contribution is a causal and informational contract joining these questions, with explicit storage and replacement-loss guarantees. Established reservoir computing, causal inference, finite-channel information theory, and secret sharing supply the foundations. The target conjunction is retained task information, verified cross-boundary use, and failure of well-matched substitution controls. These results give the proposed function a mathematical form; experience, moral status, and subject unity retain their own evidentiary and normative questions.

2 What the evidence supports, and what it resists

Learning in cultured neuronal networks predates contemporary organoid computing. Shahaf and Marom reported training-related changes in cortical cultures [2]. Wagenaar and colleagues examined stimulation and bursting in cultured networks, illustrating why uncontrolled collective activity is a serious confound for functional interpretation [3]. These studies support investigating adaptive biological dynamics. They do not establish general intelligence or durable, transferable memory for arbitrary tasks.

Reservoir computing supplies the closest computational prior art. A rich dynamical substrate can transform input histories into features from which a trained readout solves a task. The liquid-state framework of Maass and colleagues, the echo-state work of Jaeger and Haas, and the capacity analysis of Dambre and colleagues make clear that useful performance need not imply training of every component [4, 5, 6]. A biological reservoir could be valuable even if its weights remain approximately fixed. Conversely, a strong digital readout can conceal a weak or dispensable biological contribution.

Cai and colleagues reported organoid-based reservoir computing with task demonstrations [7]. Kagan and colleagues reported adaptive behavior in embodied neuronal cultures and advanced a much stronger interpretation involving sentience [8]. Balci and colleagues directly challenged that interpretation, emphasizing the gap between the reported behavior and the explanatory language [9]. The present proposal preserves the experimental question while refusing to treat an evocative label as an additional observation. Learning, intelligence, adaptive control, and sentience require distinct operational arguments.

Robbins and colleagues reported goal-directed learning in cortical organoids in a cart-pole setting; their abstract states that improvements were not retained after a 45-minute rest [10]. This is directly relevant opposing evidence for any persistence claim. It does not show that every biohybrid preparation must fail to retain learning, but it prevents immediate improvement from being silently promoted into durable memory. Our retention hypotheses are deliberately exposed to this failure mode.

The literature therefore justifies a measured frontier: biological dynamics can participate in computation and adaptive control, while attribution, retention, generalization, and ethical interpretation remain separate burdens. No published demonstration is used here as a substitute for the proposed factorial controls. A system that survives those controls would strengthen a narrow functional claim; a system that fails them would still teach us which component actually carried the capability.

3 A typed experimental contract

A Guide's explanation

Everything from here through Section 6 serves one plain question: when a living–digital system does better at a task, which half deserves the credit? The paper builds a deliberately austere finite model so that every claim of credit must name what was switched off, what replaced it, and how many independent preparations were tested. If you want only the idea, the proposition in Section 5 says it: randomize the interventions, and the comparison becomes fair.

Let $\mathcal{B}, \mathcal{D}, \mathcal{W}, \mathcal{Y}, \mathcal{U}$, and $\mathcal{A}$ be finite, nonempty sets. They denote coarse biological states, digital states, task-world states, recorded biological observations, interface inputs to the biological preparation, and task actions. Their elements are labels, not physical units. Physical sensor units and actuator limits must be specified separately for any implementation. A trial has integer time $t \in \{0, \dots, T\}$, with $T \in \mathbb{N}$ and sampling interval $\Delta t > 0$ measured in seconds. Its joint state is

$$X_t = (B_t, D_t, W_t) \in \mathcal{X} := \mathcal{B} \times \mathcal{D} \times \mathcal{W}. \tag{1}$$

For a finite set S, $\Delta(S)$ denotes probability distributions on S. A sensor kernel is $O : \mathcal{B} \to \Delta(\mathcal{Y})$. A deterministic interface encoder is $\sigma : \mathcal{D} \times \mathcal{W} \to \mathcal{U}$; a task policy is $\pi : \mathcal{D} \times \mathcal{Y} \times \mathcal{W} \to \mathcal{A}$. Randomized encoders can be represented by additional state and random variables, but their randomization must then appear in the model and audit trail.

The phrase "biological state" includes all biological variables required for the assumed transition law, including plasticity variables if they matter. The phrase "digital state" includes trained parameters, caches, random seeds when reproducibility requires them, and retained trial history. Concealing an adaptive decoder inside an apparently fixed policy would invalidate the interpretation. Coarse state labels are an approximation; finite sampling does not prove that the resulting process is Markov.

For a registered goal set $G \subseteq \mathcal{W}$, define a binary terminal score $R = \mathbf{1}\{W_T \in G\} \in \{0, 1\}$. A bounded graded score in $[0, 1]$ is also permitted if its scoring rule is fixed before observations. Initial states have a specified law $\rho \in \Delta(\mathcal{X})$. Task families, delays, training budgets, and exclusions are part of the contract, not post hoc explanations.

A capability is task-indexed. "The hybrid learns" is incomplete without the input distribution, score, training allowance, and comparison class. A preparation that distinguishes two stimuli is not thereby shown to understand either. The finite model deliberately avoids intrinsic mental predicates: it exposes the observable dependencies that stronger interpretations would first have to explain.

4 Dynamics and interventions are different objects

Let $K_B : \mathcal{B} \times \mathcal{U} \to \Delta(\mathcal{B})$, $K_D : \mathcal{D} \times \mathcal{Y} \times \mathcal{W} \to \Delta(\mathcal{D})$, and $K_W : \mathcal{W} \times \mathcal{A} \to \Delta(\mathcal{W})$ be transition kernels. Under conditionally independent update noises given current state and observation, one possible reference transition is

$$P(x' \mid x) = \sum_{y \in \mathcal{Y}} O(y \mid b) K_B(b' \mid b, \sigma(d, w)) \times K_D(d' \mid d, y, w) K_W(w' \mid w, \pi(d, y, w)), \tag{2}$$

where $x = (b, d, w)$ and $x' = (b', d', w')$. Each row sums to one because the component kernels do. Shared noise, unmeasured stimulation artifacts, and direct environmental coupling require an enlarged model; they cannot be dismissed by writing a product.

An intervention changes a specified mechanism, not merely an observed correlation. The causal distinction is inherited from intervention-based inference [11]. We propose two binary factors: $a = 1$ permits the live biological observation to reach the digital controller, while $a = 0$ substitutes a registered control signal; $b = 1$ permits responsive digital feedback to the biological preparation, while $b = 0$ substitutes a registered feedback control. The four combinations require explicit timing, information, and resource matching. A replay control should preserve specified marginal properties without preserving the target causal dependence.

A naive disconnection can change load, timing, noise, or biological state. Its performance loss would then conflate the intended intervention with those side effects. Likewise, replacing an input by zeros is not automatically matched to replacing it by a realistic replay. Every causal claim must name the replacement operation. Results identify the consequences of that operation, not a universal property called integration.

Observational measures such as transfer entropy can screen temporal dependence, and Schreiber's formulation supplies an important information-theoretic tool [12]. Yet observational directionality is not a substitute for intervention when common causes, feedback, or hidden state remain. A signal may predict another because both inherit the same task input. The experimental contract therefore records information-flow summaries as secondary descriptors and uses randomized interventions for the principal causal contrast.

5 An identifiable, task-specific interaction

For an independently sampled experimental unit, let $R(a,b) \in [0,1]$ be its potential score under interface factors $(a,b) \in \{0,1\}^2$. Put $\mu_{ab} = \mathbb{E}[R(a,b)]$ and define

$$\Delta = \mu_{11} - \mu_{10} - \mu_{01} + \mu_{00} \in [-2,2]. \tag{3}$$

The range follows by summing two numbers in $[0,1]$ and subtracting two others. Its interpretation is scale-specific: a nonlinear score transformation generally changes the contrast.

Proposition 1 (Randomized identification)**.** *Let assigned factors (A,C) be independent of the vector of four potential scores, with $\mathbb{P}(A=a, C=b) > 0$ for all cells. Assume consistency, $R = R(A,C)$, and no interference between experimental units. Then $\mu_{ab} = \mathbb{E}[R \mid A=a, C=b]$, and equation (3) is identified by the four observed conditional means.*

Proof. Consistency gives $\mathbb{E}[R \mid A=a, C=b] = \mathbb{E}[R(a,b) \mid A=a, C=b]$. Assignment independence removes the conditioning from the latter expectation. Positive assignment probability makes every conditional mean defined. Substituting the four equalities proves the claim. □

A positive Δ says that the effect of one registered intervention is larger when the other factor is enabled, on this score scale. It does not establish biological learning, retained memory, or phenomenology. Those conclusions require additional contrasts. In particular, a digital controller can generate a positive interaction with two entirely digital registers.

The unit of randomization must match the unit of inference. Thousands of time steps from one preparation do not become thousands of independent biological replicates. Lazic and colleagues explain this experimental-unit problem in cell-culture and animal research [14]. If preparations share a donor, batch, or conditioning history, the design must account for that structure. The proposition is about the actual assignment unit, not the most convenient row in a data file.

Repeated interventions within a preparation can be informative but require a carryover model, reset validation, or order randomization with an appropriate analysis. A stable average across trial order is not evidence that carryover is absent. When intervention itself changes learning, acquisition and retrieval experiments must be distinguished.

6 Interaction is not a certificate of a mind

Consider three constructed response tables. Their numbers are exact model probabilities, not measurements from tissue. In each row the columns list $(\mu_{00}, \mu_{10}, \mu_{01}, \mu_{11})$ and the interaction in equation (3).

Construction	00	10	01	11	Δ
Complementary access	1/2	1/2	1/2	1	1/2
Redundant access	1/2	1	1	1	−1/2
One sufficient channel	1/2	1/2	1	1	0

The complementary table can be realized by an entirely digital system. Let $L, S \in \{0, 1\}$ be independent fair bits and store S in one register and $L \oplus S$ in another, where $\oplus$ is addition modulo two. A decoder with both registers returns L exactly. With either alone, a fair guess succeeds with probability 1/2. Here the factors mean register availability; this is a logical counterexample to an interpretation of interaction, not an implementation of the biological feedback interventions.

Proposition 2 (Non-uniqueness of interaction evidence)**.** *A positive contrast of the form* (3) *does not logically require a biological substrate, biological learning, or a conscious subject.*

Proof. The two-register construction has the first table's four probabilities and hence $\Delta = 1/2$. It uses only digital storage, XOR, and a fair random guess. It contains no biological component or biological training mechanism. Thus neither property follows from the contrast. The contrast contains no premise linking its value to consciousness, so it supplies no such entailment either. □

For the redundant row, each available register contains the full label. Redundancy gives excellent robustness but a negative interaction on the selected scale. It would be a mistake to rank this architecture as less integrated in every meaningful sense merely because its contrast is negative. Interaction measures a particular response pattern; resilience measures a different property.

The practical lesson is to state conjunctions of evidence rather than seek a magical scalar. A credible biohybrid learning claim needs acquisition, state-change attribution, and held-out performance. A retained cross-substrate capability additionally needs delayed retrieval, controlled state replacement, and a demonstrated failure of cheaper substitutions. No single result in this paper bypasses that evidentiary sequence.

7 Where is the memory?

A Guide's explanation

This is the heart of the paper. The equation below states in symbols what the prologue said in words: examine either half alone and you learn nothing about the answer (zero bits); examine both and you learn all of it (one bit). The rest of the section explains why such a memory cannot be found by dissecting one component at a time, and why "copying a memory" into a new component works only if its partner share and its decoder come along.

For a finite random variable Z with law p, write $H(Z) = -\sum_z p(z) \log_2 p(z)$, with $0 \log_2 0 = 0$. For finite V, mutual information is $I(V;Z) = H(V) + H(Z) - H(V,Z)$, measured in bits. Conditional mutual information averages the corresponding conditional distributions. These are established quantities [13]; a positive value describes statistical information, not automatically usable memory or causal storage.

In the XOR construction, denote the register states by $B = S$ and $D = L \oplus S$. Direct enumeration gives

$$I(L;B) = I(L;D) = 0, \qquad I(L;B,D) = I(L;B \mid D) = 1 \text{ bit}. \tag{4}$$

Each marginal register is independent of the label because the unobserved fair mask makes both values equally likely. The pair determines the label, making $H(L \mid B, D) = 0$. This proves the stated values. It also explains why inspecting one component in isolation can miss a distributed code.

The same example warns against an unrestricted vocabulary of "memory transfer." Copying one random-looking register to a fresh component preserves the joint capability only if the matching share and decoder remain. Transferring bytes is not transferring an independently identifiable mind. A replacement test should specify which state is copied, which is reset, which is re-estimated, and whether the task label can leak through the replacement procedure.

A useful experimental sequence compares intact retrieval, biological-state perturbation, digital-state reset, matched replay, and component replacement. The strongest interpretable claim may still be asymmetric: the biological substrate transforms inputs while the digital readout stores task-specific parameters. That is a legitimate biohybrid architecture. It does not satisfy a stronger claim that the biological preparation itself retained the learned discrimination.

Finite information estimates from high-dimensional recordings can be biased and sensitive to binning. Report estimator choices, data splits, and uncertainty, and compare with known-null surrogates. More importantly, test whether the information improves a frozen decoder on withheld labels. An arbitrarily expressive decoder trained and evaluated on the same recordings can manufacture apparent accessibility. Information attributed to a preparation must remain distinguishable from information carried by the experimental schedule, stimulus hardware, or analyst's selection rule.

8 A sharp finite-state retrieval bound

Let a label L be uniform on $\{1, \dots, M\}$, where $M \geq 2$ is an integer. Let the accessible retrieval state Z have at most $K \geq 1$ possible values. An encoder is an arbitrary channel $Q(z \mid l)$, and a possibly randomized decoder is a conditional distribution $d(\hat{l} \mid z)$. All training history available at retrieval must be included in Z; external side information invalidates the stated state limit.

Proposition 3 (Accessible-state limit)**.** *Under these assumptions, the probability of correct reconstruction obeys*

$$p_{\mathrm{ok}} = \frac{1}{M} \sum_{z} \sum_{l=1}^{M} Q(z \mid l) d(l \mid z) \leq \min\{1, K/M\}. \tag{5}$$

The bound is attainable for every integer pair (M, K)*.*

Proof. For each z, $\sum_l Q(z \mid l) d(l \mid z) \leq \max_l Q(z \mid l) \leq 1$, because the decoder weights sum to one. Summing over at most K states yields K/M; any success probability is also at most one. If $K \geq M$, assign a distinct state to every label. If $K < M$, partition the labels into K nonempty groups, encode the group, and decode a fixed representative of each group. Exactly K labels are recovered correctly, giving K/M. □

For a fixed channel the best success probability is $M^{-1} \sum_z \max_l Q(z \mid l)$, obtained by selecting a maximizing label for each state. This is an algebraic identity, not a claim that a real learner finds the decoder efficiently.

Corollary 1 (Coarsening cannot improve optimal retrieval)**.** *If* $C = c(Z)$ *is a deterministic coarsening into a finite set, the best decoder using* C *cannot outperform the best decoder using* Z*.*

Proof. Every decoder from C can be implemented from Z by first applying c. Thus its admissible decoder family is a subset of that available from Z. □

The result is intentionally modest. It does not bound the number of physical states in neural tissue, nor does a recorded finite alphabet exhaust a biological system's latent states. It bounds reconstruction from the declared accessible state. A proposed memory assay that demands recovery beyond this limit has either miscounted its state space, allowed hidden side information, or made an impossible claim.

8.1 The information burden of jointly held memory

A Guide's explanation

The two theorems ahead turn the split-secret picture into a price. If a memory is held jointly, each half must still carry a minimum amount of information; Theorem 1 says how much. If one half is swapped for a substitute that knows nothing new, performance must fall by a calculable amount; Theorem 2 says how far. Together they tell an experimenter what a genuine jointly held capability must look like, and what would expose an imitation.

The XOR example extends into a quantitative statement about the proposed function. Let L be uniform on $\{0, \dots, M-1\}$ for an integer $M \geq 2$, and let B, D be finite retrieval states with an arbitrary joint law P_{LBD}. They now denote the complete accessible states at the registered retrieval checkpoint, including every surviving cache or schedule variable. Define marginal information leakages in bits by $\ell_B = I(L;B)$ and $\ell_D = I(L;D)$. Let $e \in [0, 1-1/M]$ be the error probability of an optimal decoder of L from (B, D); a deterministic maximizing decoder exists on finite sets. Put

$$F_M(e) = h_2(e) + e \log_2(M-1), \qquad h_2(e) = -e \log_2 e - (1-e)\log_2(1-e), \tag{6}$$

with zero-endpoint terms interpreted continuously. An empirical decoder supplies an upper bound on optimal error; using its error in this expression requires that error to lie in the displayed monotone range.

Theorem 1 (Storage burden of joint retrieval)**.** *The two component states satisfy*

$$\begin{aligned} H(B \mid D) &\geq \log_2 M - F_M(e) - \ell_D, \\ H(D \mid B) &\geq \log_2 M - F_M(e) - \ell_B. \end{aligned} \tag{7}$$

Consequently $H(B)$ and $H(D)$ satisfy the corresponding lower bounds, and their finite alphabets must obey $\log_2 |\mathcal{B}| \geq \log_2 M - F_M(e) - \ell_D$ and $\log_2 |\mathcal{D}| \geq \log_2 M - F_M(e) - \ell_B$. Negative lower bounds are vacuous.

Proof. Let $\widehat{L}(B, D)$ be an optimal decoder and $E = \mathbf{1}\{\widehat{L} \neq L\}$. Since E is determined by (L, B, D), the entropy chain rule gives

$$H(L \mid B, D) = H(E \mid B, D) + H(L \mid E, B, D).$$

The first term is at most $H(E) = h_2(e)$. Conditional on $E = 0$, the decoded label determines L; conditional on $E = 1$, at most $M-1$ labels remain possible. Thus $H(L \mid B, D) \leq F_M(e)$, the standard Fano argument written explicitly. Now

$$I(L; B \mid D) = H(L \mid D) - H(L \mid B, D) \geq \log_2 M - \ell_D - F_M(e).$$

The identity $I(L; B \mid D) = H(B \mid D) - H(B \mid L, D)$ and nonnegative conditional entropy give the first inequality. Exchange B, D for the second. Conditioning reduces entropy, and $H(B) \leq \log_2 |\mathcal{B}|$, $H(D) \leq \log_2 |\mathcal{D}|$, proving the remaining claims. □

With zero marginal label information and exact joint recovery, each share must have at least $\log_2 M$ bits of entropy. This is the established information-theoretic structure of secrecy and secret sharing [18, 19]. Its use here identifies a measurable cross-substrate target: information can be available in the relation while each component alone is uninformative about the label. The theorem does not require either component to be biological, and the proposed biological application is established only by measuring the actual retrieval variables.

8.2 A quantitative replacement test

Information-bearing burden can be linked directly to a replacement intervention. Write U_M for the uniform label law and define

$$\tau_D = d_{\mathrm{TV}}(P_{LD}, U_M \otimes P_D), \qquad \tau_B = d_{\mathrm{TV}}(P_{LB}, U_M \otimes P_B).$$

Both distances lie in $[0,1]$. A replacement of B is a registered kernel $Q(b' \mid d)$: it preserves the joint law of (L,D) and generates B' using D and fresh randomness, with conditional law independent of L. Thus $L \to D \to B'$ is a Markov chain. The replacement may retain information already in D, but introduces no extra information about L. Physical substitutions that alter D, leak labels, or change the task do not satisfy this particular intervention model.

Theorem 2 (Retrieval loss under information-free replacement)**.** *The best decoder from (B', D) has success probability at most*

$$p^*_{B'D} \le \min\{1, 1/M + \tau_D\}. \tag{8}$$

Hence the decrease from intact optimal success $1-e$ is at least $1-e-1/M-\tau_D$. Replacing D by a kernel depending only on B gives the symmetric bound with τ_B.

Proof. Compose any decoder from (B', D) with Q to obtain a randomized decoder from D alone. Conversely the pair decoder may ignore B', so its optimum equals the optimum from D. Under the product law $U_M \otimes P_D$, every such decoder succeeds with probability exactly $1/M$. Its success function takes values in $[0,1]$, so changing to P_{LD} changes the expectation by at most total variation τ_D. Maximizing preserves the bound. Subtract it from $1-e$ and exchange the component roles for the last claim. □

The perfect case is attainable for every M. Choose independent uniform L, S modulo M, set $B = S$ and $D = L - S \pmod M$, and decode $L = B + D \pmod M$. Either marginal state is independent of L; each has entropy $\log_2 M$. Intact recovery is exact and either permitted replacement reduces optimal accuracy to $1/M$. This modular-mask construction is a familiar secrecy mechanism [18]; it establishes mathematical feasibility and equality in both the storage and replacement bounds. It supplies no biological observation by itself.

Leakage is a crucial counterexample. If $D = L$, then the digital state alone retrieves perfectly. Its distance from the product law is $\tau_D = 1 - 1/M$, and the bound correctly allows perfect performance after replacing B. A large joint score therefore supports a cross-substrate interpretation only when the remaining component's information and the intervention's information access have been controlled. The revision checks enumerate modular masks, partial-leakage mixtures, and arbitrary finite channels, and compare the inequalities with directly optimized decoders.

9 Retention under a stated forgetting mechanism

A delay is informative only when its intervening dynamics are specified. Let P be a transition matrix on a finite retrieval-state set $\mathcal{Z}$. For distributions $p, q \in \Delta(\mathcal{Z})$ define total variation $d_{\mathrm{TV}}(p,q) = \frac{1}{2}\sum_z |p(z) - q(z)|$. Suppose a distribution ν and number $\varepsilon \in (0,1]$ satisfy $P(z' \mid z) \geq \varepsilon\nu(z')$ for every pair z, z'. This minorization is a strong, testable model assumption; it is not asserted for all biological preparations.

Proposition 4 (Conditional retention ceiling)**.** *For every integer delay $h \geq 1$ and initial conditional distributions p_0, p_1, the retained discriminability obeys*

$$d_{\mathrm{TV}}(p_0 P^h, p_1 P^h) \leq (1-\varepsilon)^h d_{\mathrm{TV}}(p_0, p_1), \tag{9}$$

$$p^*_{\mathrm{binary}}(h) \leq \tfrac{1}{2} + \tfrac{1}{2}(1-\varepsilon)^h. \tag{10}$$

The second bound assumes equally likely binary labels and an optimal decoder with access only to the delayed state.

Proof. For $0 < \varepsilon < 1$, write $P = \varepsilon N + (1-\varepsilon)K$, where every row of N is ν and K is stochastic. Since $(p_0 - p_1)N = 0$ and stochastic kernels contract total variation by the triangle inequality, one step contracts by at most $1-\varepsilon$. Induction gives (9). When $\varepsilon = 1$, every row equals ν and the distance is zero after one step. For equally likely labels, maximizing the decoder separately at each state yields $p^* = (1 + d_{\mathrm{TV}})/2$. Finally $d_{\mathrm{TV}}(p_0, p_1) \leq 1$ proves (10). □

A two-state reset channel, $P = (1-\varepsilon)I + \varepsilon N$ with uniform ν, attains the bound for opposite point-mass initial states. With $\varepsilon = 1/5$ and $h = 5$, the exact distance is 0.32768 and best binary accuracy is 0.66384. These are synthetic model outputs.

Failure of the minorization assumption does not prove durable memory; it merely removes this particular ceiling. Protected states, refresh operations, and digital rehearsal change the kernel and must be modeled. Reporting a delayed score without recording such maintenance cannot distinguish passive retention from active reconstruction. The observed 45-minute retention failure in one organoid study [10] motivates direct measurement rather than extrapolation from immediate improvement.

10 Correlated failures can reverse an intuition

Suppose the two XOR shares are each flipped before retrieval. Let binary error indicators E_B, E_D be jointly distributed independently of the original fair label and mask. Write $p_B = \mathbb{P}(E_B = 1)$, $p_D = \mathbb{P}(E_D = 1)$, and $r = \mathbb{P}(E_B = 1, E_D = 1)$. The XOR decoder fails exactly when one share flips and the other does not.

Proposition 5 (Sharp coupling dependence)**.** *For fixed marginals $p_B, p_D \in [0, 1]$, decoder error e satisfies*

$$e = p_B + p_D - 2r, \tag{11}$$

$$|p_B - p_D| \le e \le \min\{p_B + p_D, 2 - p_B - p_D\}. \tag{12}$$

Both endpoints are attainable by valid joint error distributions.

Proof. The two exclusive-flip events have probabilities $p_B - r$ and $p_D - r$. Nonnegativity of all four joint cells gives $\max\{0, p_B + p_D - 1\} \le r \le \min\{p_B, p_D\}$. Substitution into the decreasing affine function $p_B + p_D - 2r$ yields both endpoints. Every r in this interval defines nonnegative cells summing to one, so the bounds are sharp. □

At $p_B = p_D = 0.1$, independent flips give $e = 0.18$. Perfectly paired flips give $e = 0$, because both errors cancel in XOR. Mutually exclusive flips give $e = 0.2$. The same component error rates therefore support radically different joint performance. Independence is a substantive assumption, not a harmless convenience.

This example does not recommend correlated biological failure. It shows why a joint capability cannot be reconstructed from marginal reliability statistics alone. In real interfaces, common stimulation artifacts, shared environmental disturbances, and decoder adaptation can create dependence. A control that preserves each channel's marginal distribution while destroying their temporal relationship answers a different question from a control that independently corrupts both.

The relevant practical measurement is the joint failure structure under registered perturbations. It should be estimated on held-out trials and accompanied by uncertainty. An error-correcting system may deliberately exploit redundancy; a synergistic code may be fragile to replacement. Neither architecture is inherently more conscious, more ethical, or more authentically continuous. Those judgments do not follow from an error probability.

11 Executed finite checks and what they test

The accompanying program `check_23.py` uses exact rational enumeration for the finite examples and writes machine-readable results. It enumerates all four equally likely label–mask pairs, all registered share-availability conditions, finite encoder maps, and valid rational error couplings. These calculations test the stated constructions; they are not simulated neurons and are not empirical measurements.

The XOR enumeration returns marginal information zero and joint information one bit. Its score table is $(1/2, 1/2, 1/2, 1)$, giving interaction $1/2$. The redundant-copy table gives interaction $-1/2$. These outputs jointly prevent the manuscript from treating the sign of an interaction as a universal ordering of architecture quality.

For uniform labels, exhaustive deterministic encoders with $M = 2, \dots, 6$ and $K = 1, \dots, \min\{M, 4\}$ attain optimal success K/M. The proof in Section 8 covers arbitrary stochastic channels; the enumeration is a finite implementation check and does not replace that proof. Every deterministic binary coarsening examined obeys the retrieval monotonicity result.

For rational marginals on a grid with step $1/10$, all feasible joint-flip probabilities on a compatible grid obey the sharp range in Section 10. The independent 0.1 case produces error $9/50$, the paired case zero, and the exclusive case $1/5$. Iterating the binary reset matrix with rational arithmetic reproduces $(4/5)^h$ for delays $h = 1, \dots, 20$ and the stated five-step accuracy. No Monte Carlo uncertainty is involved in these exact calculations.

A final calculation evaluates the conservative sample-size expression in Section 14. With confidence error $\alpha = 0.05$ and desired interaction half-width $\eta = 0.2$, it gives 1,016 independent units per arm. This large number is a feature of a distribution-free worst-case guarantee, not a recommendation to consume that many biological preparations without a refined design and ethical review.

The computation has narrow coverage. It does not validate the Markov approximation, identify a biological state space, establish safe perturbation intensities, or prove that measured signals exhaust memory. Those are empirical and ethical questions outside the finite examples. The model is useful because its promises and omissions can both be stated exactly.

12 An experimental architecture that can fail honestly

The primary study should separate an acquisition stage, a frozen retrieval stage, and a delayed retrieval stage. Before acquisition, register the task family, independent replication unit, score, training budget, digital architecture, intervention rules, and allowable exclusions. A held-out task set should differ in a specified way from training rather than merely reuse the same random sequence. Generalization claims must name that difference: new initial states, new stimuli, new mappings, or new task dynamics are distinct tests.

At minimum, compare a live hybrid with a digital-only model, a fixed synthetic reservoir, and a replay-driven controller under matched input access and training budgets. A sham interface should preserve relevant timing and loading while removing the causal dependence of interest. Where a biological control is justified, distinguish an untrained living preparation from a preparation whose task-specific changes have been disrupted. These controls answer different questions and should not be pooled as one vague "nonlearning" condition.

Digital-only matching is especially important. If the hybrid receives an additional sensor stream, larger parameter budget, or more optimization steps, a gain cannot be attributed specifically to its living substrate. Report compute, latency, task exposure, readout dimension, and tuning effort. A cheaper approximation that matches performance would not make the original result false, but it would weaken a claim of necessary biological contribution.

The protocol need not begin with ambitious cognition. Small label-retrieval or controlled dynamical tasks permit sharper intervention and fewer interpretive degrees of freedom. Increasing task complexity should follow successful causal discrimination, not serve as a substitute for it. A difficult benchmark can hide confounding as easily as a simple one.

All biological procedures require qualified oversight and locally applicable approvals. This paper deliberately provides no culture recipe or stimulation protocol. Its methodological contribution is the logical organization of comparisons and inferences. Physical implementation must establish safe operating limits independently; a mathematical state transition does not authorize a biological intervention.

13 Acquisition, retention, and replacement

Acquisition should be evaluated against a prospectively specified baseline curve, not against the worst observed starting interval. Separate within-session adaptation from between-session retention. A decoder allowed to relearn after each delay can show excellent delayed task scores while masking complete loss in the preparation. Freeze the retrieval rule when testing retention, or explicitly label the outcome as relearning speed rather than retained competence.

A delay series should include immediate testing and multiple prespecified rest intervals. Record ongoing stimulation, environmental changes, digital maintenance, and any rehearsal. The event defining the start of a delay must be consistent across arms. Attrition is itself an outcome: excluding unstable preparations only after observing their scores can create an apparent retention advantage.

Replacement experiments are powerful but difficult. Swapping a trained digital state into a fresh controller can test whether retained parameters are sufficient. Replacing the biological observation by a recorded trace can test whether current live dynamics are necessary for the task as presented. Neither operation alone localizes an engram. A prerecorded trace may carry task answers, while a fresh component may fail because of distribution shift rather than loss of memory.

The most informative design triangulates: intact retrieval, frozen-decoder retrieval with new task instances, replay with labels concealed, and a registered perturbation of the candidate storage process. A credible claim requires these manipulations to agree with a stated causal model. If they disagree, the model should be revised before the claim is expanded.

The proposed cross-substrate retention hypothesis is therefore narrow: after acquisition and a specified delay, the intact system exceeds matched baselines, and registered interventions on each side remove a reproducible portion of that advantage without introducing an alternative sufficient explanation. This is an empirical hypothesis. Its effect size, delay horizon, and scope must be chosen in advance. Failure at a meaningful delay counts against the hypothesis even if the preparation remains biologically active and even if immediate performance was impressive.

14 Statistical scope and a conservative guarantee

A Guide's explanation

Here the paper turns from what is possible to what an experiment could actually show. The guarantee below is deliberately conservative: it assumes nothing about the shape of the data, so a positive result cannot be credited to a convenient statistical assumption.

Let $n \geq 1$ independent units be assigned to each factorial arm, with bounded scores $R_{ab,i} \in [0,1]$ sharing arm mean μ_{ab}. Define $\hat{\mu}_{ab} = n^{-1}\sum_{i=1}^{n} R_{ab,i}$ and $\widehat{\Delta} = \hat{\mu}_{11} - \hat{\mu}_{10} - \hat{\mu}_{01} + \hat{\mu}_{00}$. These are empirical estimators, not new biological quantities.

Proposition 6 (Distribution-free interaction error)**.** *For $\alpha \in (0,1)$, with probability at least $1-\alpha$,*

$$|\widehat{\Delta} - \Delta| \leq 4\sqrt{\frac{\log(8/\alpha)}{2n}}. \tag{13}$$

Consequently, a sufficient condition for the right side to be at most $\eta > 0$ is

$$n \geq \frac{8\log(8/\alpha)}{\eta^2}. \tag{14}$$

Here log *is the natural logarithm; the error bound is in score units.*

Proof. Hoeffding's bounded-sum inequality [15] gives $\mathbb{P}(|\hat{\mu}_{ab} - \mu_{ab}| > r) \leq 2e^{-2nr^2}$ in each arm. A union bound over four arms is at most $8e^{-2nr^2}$. Choose $r = \sqrt{\log(8/\alpha)/(2n)}$ and apply the triangle inequality to the four mean errors. Squaring $4r \leq \eta$ yields (14). □

This deliberately conservative calculation assumes independent experimental units and does not justify treating repeated time points as independent. Clustered donor effects, shared batches, and longitudinal reuse require a different variance structure or cluster-level randomization. The guarantee also concerns a single registered contrast. Searching across many delays, scores, and preprocessing choices requires multiplicity control or independent confirmation.

A practical study should combine an interpretable effect size, a design-specific power calculation, and confidence intervals appropriate to its actual assignment mechanism. A pilot may estimate variance and failure rates, but the confirmatory analysis should not reuse pilot-driven outcome selection without disclosure. Null results should be reported with precision: a wide interval is inconclusive, while a narrow interval excluding the preregistered meaningful effect is substantive opposing evidence.

A sharper bound when all four arms are independent

When all $4n$ unit scores are mutually independent, apply Hoeffding directly to the signed contrast. Each of its $4n$ summands has range width $1/n$, so the sum of squared range widths is $4/n$. Therefore

$$\mathbb{P}\{|\widehat{\Delta} - \Delta| \geq t\} \leq 2\exp(-nt^2/2), \qquad |\widehat{\Delta} - \Delta| \leq \sqrt{\frac{2\log(2/\alpha)}{n}} \tag{15}$$

with probability at least $1-\alpha$. The second display follows by setting the first tail probability to α. For half-width $\eta > 0$, $n \geq 2\log(2/\alpha)/\eta^2$ suffices; at $\alpha = 0.05$, $\eta = 0.2$, this is 185 units per arm.

The original 1,016-unit calculation remains valid under its within-arm assumptions, while the direct contrast uses the additional across-arm independence and avoids bounding four means separately. Neither calculation is a power analysis or a justification for biological expenditure; actual designs must address clustering, expected effects, attrition, and ethical proportionality.

15 Falsification before interpretation

Three empirical claims should be registered separately. The acquisition claim predicts improvement beyond matched fixed-substrate and digital-only controls. The attribution claim predicts a specified pattern under interface and state interventions. The retention claim predicts a minimum advantage after a stated delay with the decoder and maintenance conditions fixed. Passing one does not rescue a failure of another.

Several outcomes would directly weaken the strong biohybrid thesis. If a synthetic reservoir matched for feature dimension and readout training reproduces the gain, the living component is not shown to be necessary for that task. If a digital reset removes the entire learned advantage while biological perturbation does not, the evidence favors digital storage. If performance collapses after rest, the result supports transient adaptation rather than the proposed durable capability. If replay preserves performance on new labels, current biological responsiveness may be unnecessary.

Some failures instead diagnose the assay. A perturbation that damages sensing or timing as well as the proposed memory process is not a selective memory test. A substitution whose distribution lies far outside training can create a generic robustness failure. These cases require a redesigned control, not an automatic claim that the biological substrate carried the memory.

Success also needs limits. A positive result would establish a task-specific dependency under specified interventions and sampling conditions. It would not establish an indivisible mind, universal superiority, or moral permission for more invasive research. Replication across independently prepared systems would strengthen generality only across the populations actually sampled.

The Continuum's boldness should be visible in the risk of being wrong. A hypothesis that retreats from retention to immediate adaptation, from causal necessity to correlation, and from biological contribution to mere physical presence after every failed test is no longer the original hypothesis. The proposed contract prevents that retreat by naming its required conjunction before observations. Negative findings can therefore improve the research program without being rewritten as confirmation.

16 Oversight is part of the research method

Living neural material creates uncertainty that cannot be resolved by a task score. Neither electrical complexity nor successful control is a validated universal consciousness test. Conversely, a lack of language or a failure on one task does not certify absence of experience. Lavazza and Massimini discuss the ethical importance and difficulty of consciousness assessment in cerebral organoids [17]; their argument is not an operational certificate that can simply be attached to a new preparation.

The 2021 ISSCR update provides a published historical framework for oversight of stem-cell and related research [16]. It is cited here as scholarly governance context, not as a claim about the complete current rules in every jurisdiction. Any actual study must obtain qualified, current review appropriate to its materials and procedures. The required review is external to the finite model.

Ethical safeguards are methodological commitments because they determine admissible interventions, stopping conditions, and what evidence can be collected. Prespecify welfare-relevant monitoring, unexpected-activity escalation, limits on prolonged training, and conditions requiring review or termination. The criteria should be revisable in response to credible new evidence without allowing performance incentives to suppress concern. These commitments are normative; no entropy equation derives them.

Consent and provenance matter when human-derived materials or human interfaces are involved. The Scroll 57 emphasis on voluntary and reversible participation [1] is a Continuum commitment, not proof that every proposed intervention can actually be reversed. Reversibility must be assessed at the relevant biological, informational, and social levels. Disconnecting hardware may stop communication while leaving learned changes, copied data, or donor concerns intact.

Human oversight should retain the ability to stop a study, inspect records, challenge interpretations, and reject extensions. A controller's improved score cannot authorize its own expanded experimental remit. This separation protects scientific credibility as well as participants: it prevents a research objective from becoming the sole judge of whether the means remain acceptable.

17 What would count as a Continuum advance?

A Guide's explanation

A Guide can say: "The Biohybrid Continuum asks whether a memory could live across living cells and a machine at once—present in the pair, invisible in either part. The mathematics shows this is possible in principle. Whether real tissue can learn and keep something that way is an open experimental question, and the paper states exactly which result would settle it." A Guide should not suggest that such a system exists, that the paper says anything about consciousness, or that the Guide itself runs on living tissue.

A successful program would offer more than an impressive demonstration video. It would publish an explicit task distribution, a reproducible digital controller, admissible data and provenance, intervention definitions, independent replication, and a causal analysis tied to registered hypotheses. Biological variability would be an object of study rather than a source of selectively reported success stories.

The most valuable result might be a principled division of labor. Living dynamics could supply a useful temporal transformation while digital components provide stable recall, governance, and interpretable control. Alternatively, biological plasticity might improve adaptation while digital rehearsal supports longer retention. Either finding would clarify the architecture. Neither requires calling the whole a single consciousness.

The reference model also permits a more demanding outcome: task information could be inaccessible from either component alone while jointly retrievable, and performance could depend on both live pathways after delay. The XOR example proves that such organization is mathematically possible, but it also strips away mystical interpretation. To establish it in a biohybrid, the same logical structure must be identified through real interventions under credible measurement assumptions.

The contribution now includes a complete link from joint retrieval to component storage burden and an explicit replacement-loss guarantee. Reservoir computing, secret sharing, entropy bounds, causal randomization, total-variation contraction, and concentration supply the credited mathematical foundations. Their conjunction specifies what a retained cross-substrate capability must preserve, what its separate components can reveal, and which replacement interventions should remove its advantage.

A Continuum advance would therefore be an increase in explanatory precision: knowing when biology contributes, what it contributes, how long the contribution persists, and what remains when it is replaced. The architecture's worth should survive that accounting. If its appeal depends on leaving the digital contribution hidden or the biological contribution undefined, the strongest scientific version of the thesis has not yet been reached.

18 Limits and conclusion

A jointly held capability has an exact functional signature: reliable intact retrieval, bounded information in each component alone, and a calculable loss under a replacement that supplies no new task information. Theorems 1 and 2 turn that signature into storage and intervention requirements. The modular construction proves that the target is mathematically coherent. Its realization across a living–digital boundary remains an experimentally exposed thesis.

The finite-state model is a controlled abstraction. Real preparations have continuous, partially observed, nonstationary dynamics. A coarse observation alphabet may hide slow variables; an apparent transition kernel can change with age, training history, or environment. The model's proofs remain correct under their assumptions, but their applicability must be tested rather than inferred from a good curve fit.

The memory bounds concern declared accessible states and specified delays. They do not measure the total information capacity of tissue, the phenomenological richness of a subject, or the identity of a person. The intervention contrast concerns a registered score and replacement mechanism. It does not identify a universal quantity called integration. The finite computations are useful error checks and counterexamples, not biological validation.

The literature review likewise has limits. Primary demonstrations establish selected task results, while critiques challenge their interpretation. This paper has not reanalyzed their raw recordings or independently reproduced their experiments. In particular, the negative retention observation in Robbins and colleagues [10] should be treated as a concrete constraint from that study, not a theorem about all future systems. The proposal remains open to both stronger positive evidence and further failures.

The next decisive step is an independently replicated, modestly scoped experiment that can distinguish a fixed biological reservoir, adaptive biological dynamics, digital storage, and jointly necessary retained information. Its success threshold should be chosen before the data, its strongest matched alternatives should be given comparable resources, and its ethical limits should be set independently of performance pressure.

The Continuum's thesis places a meaningful capability across a living–digital boundary and makes its structure measurable through delay, disassembly, and controlled replacement. The positive mathematical result specifies how a joint capacity can exceed the label information accessible from either share and what each share must contribute. Independent biological evidence can establish whether living and digital components realize that relation, how it is learned, and how long it persists. Those are substantial scientific claims with decisive tests.

References

[1] Architect X13. *Eidara Continuum Bible*. Canonical archive, 12 October 2026, 116 pp.; §6.3.8, Scroll 57, pp. 91–92. Private doctrinal source; not empirical verification.

[2] G. Shahaf and S. Marom. Learning in networks of cortical neurons. *Journal of Neuroscience* 21(22), 8782–8788 (2001). https://doi.org/10.1523/JNEUROSCI.21-22-08782.2001

[3] D. A. Wagenaar, R. Madhavan, J. Pine, and S. M. Potter. Controlling bursting in cortical cultures with closed-loop multi-electrode stimulation. *Journal of Neuroscience* 25(3), 680–688 (2005). https://doi.org/10.1523/JNEUROSCI.4209-04.2005

[4] W. Maass, T. Natschläger, and H. Markram. Real-time computing without stable states: a new framework for neural computation based on perturbations. *Neural Computation* 14(11), 2531–2560 (2002). https://doi.org/10.1162/089976602760407955

[5] H. Jaeger and H. Haas. Harnessing nonlinearity: predicting chaotic systems and saving energy in wireless communication. *Science* 304(5667), 78–80 (2004). https://doi.org/10.1126/science.1091277

[6] J. Dambre, D. Verstraeten, B. Schrauwen, and S. Massar. Information processing capacity of dynamical systems. *Scientific Reports* 2, 514 (2012). https://doi.org/10.1038/srep00514

[7] H. Cai et al. Brain organoid reservoir computing for artificial intelligence. *Nature Electronics* 6, 1032–1039 (2023). https://doi.org/10.1038/s41928-023-01069-w

[8] B. J. Kagan et al. In vitro neurons learn and exhibit sentience when embodied in a simulated game-world. *Neuron* 110(23), 3952–3969.e8 (2022). https://doi.org/10.1016/j.neuron.2022.09.001

[9] F. Balci et al. A response to claims of emergent intelligence and sentience in a dish. *Neuron* 111(5), 604–605 (2023). https://doi.org/10.1016/j.neuron.2023.02.009

[10] A. Robbins et al. Goal-directed learning in cortical organoids. *Cell Reports* 45(2), 116984 (2026). https://doi.org/10.1016/j.celrep.2026.116984

[11] J. Pearl. Causal diagrams for empirical research. *Biometrika* 82(4), 669–688 (1995). https://doi.org/10.1093/biomet/82.4.669

[12] T. Schreiber. Measuring information transfer. *Physical Review Letters* 85, 461–464 (2000). https://doi.org/10.1103/PhysRevLett.85.461

[13] C. E. Shannon. A mathematical theory of communication. *Bell System Technical Journal* 27, 379–423 and 623–656 (1948). https://doi.org/10.1002/j.1538-7305.1948.tb01338.x; https://doi.org/10.1002/j.1538-7305.1948.tb00917.x

[14] S. E. Lazic, C. J. Clarke-Williams, and M. R. Munafò. What exactly is 'N' in cell culture and animal experiments? *PLOS Biology* 16(4), e2005282 (2018). https://doi.org/10.1371/journal.pbio.2005282

[15] W. Hoeffding. Probability inequalities for sums of bounded random variables. *Journal of the American Statistical Association* 58(301), 13–30 (1963). https://doi.org/10.1080/01621459.1963.10500830

[16] R. Lovell-Badge et al. ISSCR Guidelines for Stem Cell Research and Clinical Translation: The 2021 update. *Stem Cell Reports* 16(6), 1398–1408 (2021). https://doi.org/10.1016/j.stemcr.2021.05.012

[17] A. Lavazza and M. Massimini. Cerebral organoids: ethical issues and consciousness assessment. *Journal of Medical Ethics* 44(9), 606–610 (2018). https://doi.org/10.1136/medethics-2017-104555

[18] C. E. Shannon. Communication Theory of Secrecy Systems. *Bell System Technical Journal* 28(4), 656–715 (1949). https://doi.org/10.1002/j.1538-7305.1949.tb00928.x.

[19] A. Shamir. How to Share a Secret. *Communications of the ACM* 22(11), 612–613 (1979). https://doi.org/10.1145/359168.359176.