KRYONIS LAB
WORKING PAPER NO. 001

The Verification Gap in Biological Computing

A foundational analysis of the structural verification deficit in technologies built from living substrates

March 2026
kryonislab.org

This paper presents the foundational thesis of KRYONIS Lab: that no dedicated, field-wide verification framework for biological computation as computing has yet emerged, and that the verification layer for biological computing will become strategically more important than the biological computing hardware itself. The analysis rests on three independent open-source intelligence assessments conducted in March 2026, covering capability assessment, institutional audit, and global verification landscape across English, Russian, and Chinese source material.

Citation: KRYONIS Lab (2026). The Verification Gap in Biological Computing. Working Paper No. 001. KRYONIS Sovereign Systems Limited, Hong Kong. Available at: kryonislab.org

Disclosure: KRYONIS Lab is a programme of KRYONIS Sovereign Systems Limited, Hong Kong. The Lab develops verification infrastructure for biological computing, including the BCCS protocol (bccs.bio). This working paper presents the problem analysis that motivates the Lab's research programme. No classified sources were accessed. All citations reference open-source material.

CONTENTS

I. Executive Summary

Biological computing systems — organoid intelligence platforms, wetware processors, biohybrid neural arrays — are built from living human neurons that self-modify, drift, degrade, and die. Unlike silicon, where the same input reliably produces the same output, a biological substrate is non-deterministic: its synaptic weights change continuously through plasticity, its outputs vary with nutrient availability and temperature, and its internal state cannot be fully read without destroying the tissue.

Multiple laboratories have demonstrated that biological substrates can compute. Cortical Labs shipped the first commercially available biological computer in 2025. Indiana University demonstrated reservoir computing with brain organoids. UC Santa Cruz applied formal reinforcement learning benchmarks to cortical tissue. DARPA launched a 42-month program to build biological processing units for drone navigation.

> *Every one of these achievements rests on the same unresolved problem: how do you prove that a living computational substrate is performing as claimed, performing reproducibly, and performing safely? No existing standards body, regulatory framework, or verification architecture currently provides an answer.*

The standards infrastructure that does exist — ISO/TC 276 for organ-on-chip, the FDA ISTAND qualification pathway, CEN/CENELEC's roadmap, the NeuroBench neuromorphic benchmark — was designed for drug testing or silicon hardware. The computational verification problem specific to living substrates remains entirely unaddressed. The gap between what biological computing can do and what it can prove is the defining constraint of the field.

This working paper argues that whoever builds the dominant verification infrastructure — the benchmarks, assays, certification protocols, auditing tools, and compliance frameworks for biological computing — will control the terms on which this technology enters the global economy. The opportunity is currently unclaimed.

VERIFICATION IS THE CHOKEPOINT. NOT BIOLOGY.

II. The Gap Is Total

The 21st century is producing the first technologies that are not merely fabricated but metabolically alive. Biological computing substrates, biohybrid industrial systems, synthetic organisms deployed as infrastructure, and living materials used in construction consume resources, age, degrade, mutate, and die. They are thermodynamically open, temporally bounded, and subject to continuous state

change.

No existing standards body, regulatory framework, or verification architecture was designed for technologies that are alive. The field of Metabolic State Verification — the discipline of verifying, governing, and building trust infrastructure for technologies with metabolic existence — does not yet exist as a recognised domain.

The standards infrastructure adjacent to biological computing verification includes:

Body / Initiative	Designed For	Computing Verification?
ISO/TC 276/SC 2	Organ-on-chip drug testing	No
FDA ISTAND pathway	Drug Development Tools	No
CEN/CENELEC OoC Roadmap	European OoC standardisation	No
NeuroBench (Nat. Comms. 2025)	Silicon neuromorphic hardware	Explicitly excludes biology
China MOST guidelines (Apr 2025)	Organoid ethics / EEG caps	Ethics only, not computing
IEEE P2731 / P2794	BCI terminology / reporting	Interface level, not substrate

The gap is structural, not incremental. Adjacent frameworks address biological systems as drug testing tools, research models, or medical devices. The computational verification problem specific to living substrates — verification of computational performance, behavioural stability, training consistency, and cross-platform reproducibility for a substrate that is alive — remains entirely unaddressed.

III. Twelve Layers of Verification That Do Not Exist

Each layer introduces specific challenges. Failure at any layer invalidates all layers above it. No layer has a dedicated standard for biological computation.

● No standard exists ● Adjacent standard (drug testing / medical) ● Partial approach (academic / proprietary)

1 ***Cell Provenance***

Confirming that iPSC starting material is genetically stable, mutation-free, and traceable. Reported rejection rates exceed 60% of primary cell lots from commercial suppliers (CN Bio, 2025).

● ISO/DIS 23494-2 in development — drug testing context

2 ***Differentiation Protocol Quality***

Confirming neural differentiation produces intended cell types in correct proportions. Cell-type composition varies even within the same protocol.

● No minimum composition standard for computing

3 ***Network Morphology***

Confirming sufficient structural complexity for computation. Proposed thresholds: TBR1+/SATB2+ >80%, Synchrony Index >0.7 (Biomolecules, 2025).

● Academic thresholds proposed — not adopted

4 ***Signal Read/Write Fidelity***

Confirming MEA interface accuracy. Different platforms use different electrode counts, sampling rates, and spike detection algorithms.

● No standardised MEA qualification protocol

5 ***Training Response Consistency***

Confirming learning curves are reproducible and trained behaviour stable. Each laboratory uses different stimulation parameters and success criteria.

● No training reproducibility index exists

6 ***Temporal Stability***

Commercially available neurons survive on the order of six months (Cortical Labs CL1). Forgetting documented after 45-minute rest periods (UC Santa Cruz, 2026).

● No concept of verification-with-expiry in computing

7 ***Environmental Sensitivity***

Temperature, CO2, pH, humidity, medium composition, and contamination can alter computational behaviour without any change to the “code.”

● FinalSpark monitors 6+ variables — no formal standard

8 ***Cross-Lab Reproducibility***

Can two laboratories produce equivalent results from the same protocol? No inter-lab reproducibility study for organoid computation has been published.

- The most critical gap in the stack

9 ***Benchmark Comparability***

Can results from different platforms be compared on a common scale? Ad hoc benchmarks exist but no standardised suite has been adopted.

- NeuroBench exists for silicon — excludes biology

10 ***Explainability***

Synaptic weights cannot be read without destroying the tissue. Computations cannot be replayed. Biological computing systems are black boxes deeper than any silicon AI.

- Nothing exists

11 ***Ethical Status Thresholds***

Determining whether a substrate has morally relevant properties. China's MOST guidelines (April 2025) include EEG complexity caps. No other nation has adopted sentience thresholds.

- China MOST guidelines — ethics context only

12 ***Operator Control & Fail-Safe***

A biological system adversarially trained may be permanently compromised. The "tracing condition" for meaningful human control collapses when neurons self-organise.

- No fail-safe standard exists

IV. Who Is Closest — and What They Do Not Cover

Six actors have produced work adjacent to verification of biological computation. None addresses the computational verification problem directly.

DARPA — O-CIRCUIT Program

42-month program requiring biological processing units to achieve near-human Ms. Pac-Man proficiency and drone chemotaxis navigation. Creates de facto performance benchmarks through operational specification.

Gap: Verification by specification — not a transferable framework.

FDA — ISTAND Qualification

Three-stage pathway qualifying organ-on-chip as Drug Development Tools. Emulate's Liver-Chip achieved 87% sensitivity, 100% specificity. The closest regulatory model — designed for drug testing, not computation.

Gap: Drug testing context only — no computational performance criteria.

NeuroBench — Neuromorphic Benchmark

Community-driven benchmark suite for neuromorphic hardware. Dual-track architecture modelled on MLPerf. Published Nature Communications 2025. Explicitly excludes biological substrates.

Gap: Silicon only — biological extension acknowledged as future need.

Cortical Labs — CL1 & DishBrain

First commercially available biological computer (launched 2025, reported at $35K, ~800,000 neurons). Open-source CL API. Founder Brett Kagan has acknowledged: "Reproducibility is hard enough in science."

Gap: Proprietary QC — not an independent verification framework.

FinalSpark — Neuroplatform

World's first remotely accessible biocomputing platform. 1,000+ organoids, 94% stimulation reliability, MAP2 verification, continuous environmental monitoring, 30+ TB recorded data.

Gap: Most mature QC framework — platform-specific, not transferable.

ISO/TC 276/SC 2 — MPS Standardisation

International standards committee for organ-on-chip. NEN (Netherlands) secretariat. The closest institutional home for future computing standards.

Gap: Drug testing context — no computing verification workstream.

V. Why the Verification Layer Will Become More Important Than the Hardware

Six structural reasons why whoever controls verification controls the field.

1. Formal verification is impossible

Classical verification proves a system meets its specification by exhaustive state analysis. Biological substrates have effectively infinite state spaces that change continuously. Standard methods — model checking, theorem proving, temporal logic — require deterministic systems. Living substrates are not deterministic. Verification must be invented from first principles.

2. Attribution collapses

In silicon, a decision can be traced through weights and activations to training data. In biological computing, neurons self-organise, synaptic weights change non-destructively unreadable, and the silicon decoder adapts to the biology rather than the reverse. Without attribution, there is no accountability, no liability, and no compliance.

3. The substrate dies

Commercially available neurons survive on the order of six months. Every biological computing system requires periodic substrate replacement, and each replacement produces a new system that must be re-verified. Verification is the continuously consumed service; the hardware is the periodically replaced commodity.

4. The hardware bottleneck relaxes; the verification bottleneck does not

CMOS MEA hardware advances along predictable semiconductor scaling trajectories. The verification bottleneck — how to interpret, validate, and certify signals from living tissue — requires conceptual invention, not engineering optimisation. Structural bottlenecks are more valuable than engineering bottlenecks.

5. Market access requires verification, not hardware

The global pharmaceutical market exceeds $1.5 trillion annually. FDA acceptance requires qualification evidence. Whatever FDA requires becomes the global de facto standard. The verification evidence that enables market access is worth more than the hardware that generates it.

6. The cloud and semiconductor analogies confirm the pattern

In cloud computing, value migrated from hardware to trust infrastructure: SOC 2 audits, compliance, SLA enforcement. In semiconductors, the most profitable segments are EDA tools and verification IP — not fabrication. Biological computing follows both trajectories, compressed by the additional factor that the hardware is alive.

> *Whoever builds the dominant verification infrastructure — the benchmarks, assays, certification protocols, auditing tools, and compliance frameworks for biological computing — will control the terms on which this technology enters the global economy.*

THE OPPORTUNITY IS CURRENTLY UNCLAIMED.

VI. Building the Trust Architecture

KRYONIS Lab exists to close this gap. The Lab's research programme and protocol infrastructure address the verification problem at every layer of the stack.

The Lab — Research

kryonislab.org

Five research tracks spanning living systems verification, thermodynamic asset architecture, biocapital governance, biohybrid computation infrastructure, and ontological boundaries. AI-native methodology. Open-access working papers.

BCCS — Protocol

bccs.bio

The Biological Computing Clearing System. BAIN ID for substrate identity. Proof-of-Physical-State for metabolic verification. Eight-state lifecycle model. Base L2 settlement. The infrastructure layer that makes biological computing auditable.

The Verification Gap Brief — Intelligence

product page at kryonislab.org

Flagship intelligence product. The structural verification deficit in biological computing, aligned with concerns raised in recent governance literature, including Sirbu & Floridi (2026, Science and Engineering Ethics, Springer Nature).

MSV Protocol — Standard

in development, open review Q3 2026

Metabolic State Verification. A formal protocol for verifying the state of living technological systems. The computing-specific verification framework the field needs.

VII. Research Basis and Methodology

This working paper rests on three open-source intelligence assessments conducted in March 2026, each applying distinct methodology to a different aspect of the verification problem.

Analysis I — Capability Assessment

Comprehensive mapping of Russia's biological computing, organoid intelligence, and neuromorphic hardware capabilities. Identified the Lobachevsky–Kurchatov–MIPT institutional triangle, the BioCAM4096 infrastructure, and the complete absence of verification frameworks in the Russian ecosystem.

Analysis II — Institutional Audit

Six-track second-pass investigation stress-testing the first assessment. Identified previously undetected actors, the Gordleeva–Sarov defence connection, and the Tsinghua–Lobachevsky co-authorship pathway. Bayesian probability model for undisclosed work. Confirmed that the verification gap extends to classified and dual-use programmes.

Analysis III — Global Verification Landscape

Exhaustive investigation of verification, benchmarking, standardisation, and governance across all major jurisdictions. English, Russian, and Chinese source material. Identified the 12-layer verification stack, mapped all active standards bodies, assessed country positioning, and modelled time horizons through 2075. Confirmed the total absence of a dedicated verification framework for biological computation as computing.

Source transparency. All three analyses were conducted using open-source intelligence methods. Primary and near-primary sources were preferred. Confidence labels were applied to all major claims. No classified material was accessed. Full source documentation available on request for institutional partners.

Methodological note. The three analyses were designed to be independent and mutually stress-testing. Analysis II challenged the conclusions of Analysis I. Analysis III operated without reference to the Russia-specific findings, confirming the verification gap as a global structural condition rather than a regional observation.

VIII. Selected References

This working paper draws on the three underlying research documents cited above. Key external sources referenced in those documents and in this paper include:

Cortical Labs (2025). CL1 Product Launch. corticallabs.com/cl1

DARPA (2026). O-CIRCUIT Special Notice DARPA-SN-26-48. darpa.mil/research/programs/o-circuit

FDA (2025). ISTAND Pilot Program: First Organ-Chip Technology Acceptance. fda.gov

FinalSpark (2024). Neuroplatform: Open Research Access. Frontiers in Artificial Intelligence.

Kagan, B. et al. (2022). In vitro neurons learn and exhibit sentience when embodied in a simulated game-world. Neuron.

NeuroBench Consortium (2025). NeuroBench: Advancing Neuromorphic Computing. Nature Communications.

Sirbu, A. & Floridi, L. (2026). Ethics and governance of biological computing. Science and Engineering Ethics, Springer Nature.

Stem Cell Reports Working Group (2024). Reproducibility challenges in brain organoid research. Stem Cell Reports.

Teodorescu, M. et al. (2026). Goal-directed learning in brain organoids. Cell Reports.

UNESCO (2025). Recommendation on the Ethics of Neurotechnology. unesco.org

University of Bristol (2026). Tactile classification benchmark for organoid computing. arXiv.

CN Bio (2025). Pre-Validating Cells for Organ-on-Chip Assays. cn-bio.com

Biomolecules (2025). Quantitative maturity benchmarks for brain organoids. PMC.

ISO/TC 276/SC 2 (2024). Microphysiological Systems standardisation. iso.org

China MOST (2025). Human Organoid Research Ethical Guidelines. PMC.

KRYONIS Lab · A programme of KRYONIS Sovereign Systems Limited, Hong Kong

Building the verification layer for technologies that are alive.