A New Strategic Paradigm for Academic Medical Centers

From Disease-Oriented Data Science to the Science of Living Systems

J.Konstapel,Leiden,20-7-2026.

Executive Summary

Academic medical centers worldwide — the Mayo Clinic, Johns Hopkins Medicine, Mass General Brigham, Charité Berlin, Karolinska University Hospital, and Leiden University Medical Center (LUMC) among them — have converged on a strikingly similar strategic vision: Artificial Intelligence, precision medicine, digital health, genomics, population health, and predictive analytics.

This document argues that the current paradigm is not incomplete because of insufficient data or computational power. It is incomplete because it lacks an explicit, formal model of the living system whose health medicine exists to preserve. Parts III–IV turn this observation into a concrete, sequenced, falsifiable strategy with named deliverables and kill criteria. Part V proposes a mathematical and computational formalization — grounded in Michael Levin's bioelectric morphogenesis research, a quaternion state representation, and Kuramoto synchronization — with every non-falsifiable or esoteric extension explicitly identified and removed. Part VI grounds the whole strategy in the actual financial, regulatory, and workforce pressures facing academic medical centers today. Part VII extends the model's logic into a speculative long-term device trajectory, explicitly and permanently separated from the falsifiable core.

Document Structure (Reading Map)

- **Part I** — How LUMC and comparable institutions currently define their strategy, and what is structurally missing from it (no explicit model of the organism).
- **Part II** — Why this is a scientific gap, not only a communications gap.
- **Part III** — Six sharpened strategic principles, each with a mechanism, deliverable, and comparison to the status quo.
- **Part IV** — A four-stage implementation roadmap with explicit kill criteria.
- **Part V** — The formal model itself, organized around the principle that the organism is a native Physiological Closed-Loop Control (PCLC) system independent of any device (5.0.5): quaternion state representation (5.1), Kuramoto coupling (5.2), closure via Levin's target-morphology attractor restated as native PCLC reference-tracking (5.3), the combined coherence criterion with disease reframed as a classifiable PCLC failure mode (5.4), falsifiable predictions P1–P9 (5.5), a full status table of what is established vs. assumed vs. dropped (5.6), relation to the rest of the document (5.7), explicit limits (5.8), the Systemic Coherence Layer positioning multi-device coordination as supervisory PCLC (5.9), and the "personal blueprint" formalization of DNA-as-library / field-as-reader (5.10).
- **Part VI** — Global market and technology context (financing, AI regulation, workforce, care shifting outside hospitals, relevant scientific breakthroughs, geopolitical funding volatility, global health trajectory), and how it shaped the design choices in Parts III and V.

- **Part VII** — A speculative, explicitly quarantined device trajectory (Phases A–D) extrapolating Section 5.9, plus the six breakthroughs it would require.
- **Conclusion** — Summary and how the parts connect.

Part I — The Current Situation

1.1 A remarkable convergence

Leading academic medical centers have adopted closely related strategic priorities:

- **Mayo Clinic** — Artificial Intelligence, Precision Medicine, the Mayo Clinic Platform, global data partnerships. The institution increasingly presents itself as a healthcare platform rather than a hospital.
- **Johns Hopkins Medicine** — clinical excellence, translational research, AI-supported diagnostics, Population Health, Precision Medicine, Learning Health Systems.
- **Mass General Brigham** — integrated care, community health, AI, predictive analytics, value-based healthcare.
- **Karolinska University Hospital** — Precision Medicine, Digital Health, Systems Biology, Translational Medicine.
- **Charité Berlin** — AI, digital pathology, personalized medicine, data integration, international collaboration.
- **LUMC** — "Driven by Health" (2024–2028): three scientific priorities (Population Health, Regenerative Medicine, Data-driven Healthcare & AI), plus six transformation themes (Future@Work, Digitally Close, Education of the Future, Value-Based Care, Sustainable Operations, Connection with Society and Region). Notably, the LUMC strategy contains relatively few measurable targets for scientific output or clinical performance; the emphasis is organizational, digital, and collaborative.

1.2 The shared architecture

Despite differences in language, every strategy assumes the same causal chain:

More data → better AI → better prediction → better healthcare.

This chain is rarely questioned, and it rests on an unexamined premise: that the human organism can be adequately represented as an aggregation of measurable variables (DNA, proteins, metabolites, images, lab values, diagnoses, records), and that sufficiently powerful AI applied to this aggregation will yield understanding.

1.3 The missing ontology

No major strategic document defines, in formal terms, what the living system under study actually *is*. Medicine possesses highly detailed models of parts (cardiovascular physiology, immunology, endocrinology, neuroscience) without an integrated theory of the whole. Health is discussed extensively; the entity capable of being healthy is not defined.

1.4 The bias problem

Because academic hospitals treat severe, rare, and referred disease, their datasets are structurally skewed toward pathology and away from healthy adaptation across the lifespan. AI trained on this data becomes increasingly sophisticated at describing disease while remaining largely uninformed about the organization of health.

Part II — Why the Current Strategy Is Incomplete

Data are observations; models explain observations. Without an explicit model, AI discovers correlations rather than principles — a distinction well understood in physics and evolutionary biology, but not yet operationalized in medicine. Several partial correctives exist — Systems Medicine, Network Medicine, P4 Medicine, the Free Energy Principle, autopoiesis — but each describes *properties* of living systems without supplying a formal semantic account of how biological organization emerges and is maintained. This is the gap the remainder of this document addresses.

Part III — A New Strategic Framework (Sharpened)

The original six principles are retained but each is now attached to a concrete mechanism, a measurable deliverable, and an explicit comparison against the status quo, so that the strategy can be piloted and evaluated rather than merely adopted as rhetoric.

Principle 1 — Shift the primary scientific object from disease to living organization

Mechanism: Establish a standing institutional unit (a "Living Systems Theory Group") with the explicit mandate to produce and maintain a formal ontology of the organism, reviewed and revised on the same cycle as clinical guidelines.

Deliverable: A versioned, published ontology document — analogous to a disease classification system (ICD), but for organizational states of the living system rather than diagnostic categories.

Comparison to status quo: Current strategies delegate this question implicitly to whichever specialty happens to be publishing; no institution owns it explicitly.

Principle 2 — Develop an explicit theoretical model of the organism

Mechanism: Fund a multi-year formalization program (see Part V below) that integrates chemistry, cellular organization, physiology, cognition, and social/environmental interaction into one systems architecture, expressed in a notation precise enough to generate falsifiable predictions.

Deliverable: A minimal formal model, published with its assumptions and limits stated explicitly, plus a roadmap of what the model does *not* yet explain.

Comparison to status quo: Systems Medicine and Network Medicine describe interacting components; this principle asks for an explicit generative model, not a description of interactions after the fact.

Principle 3 — Measure organizational integrity, not only isolated biomarkers

Mechanism: Alongside existing lab panels, introduce longitudinal composite indices — order parameters describing the degree of coherence between physiological subsystems (see Part V, Section 5.3) — validated first in cohorts where both disease outcomes and long-term health trajectories are known.

Deliverable: A pilot "coherence index," tested for whether it predicts decline earlier than existing biomarker panels in at least one well-characterized cohort (e.g., a cardiometabolic or frailty cohort).

Comparison to status quo: Precision medicine optimizes thresholds on individual variables; this principle asks whether a relational, whole-system quantity carries independent predictive value.

Principle 4 — Reorient AI from disease prediction to organization modeling

Mechanism: Require that new medical AI projects funded under this strategy report not only diagnostic accuracy but also a measure of organizational/relational structure recovered from the data (e.g., whether the model's internal representations correspond to the coherence indices of Principle 3).

Deliverable: A methodological standard ("organizational validity reporting") added to institutional AI project review, alongside existing accuracy and fairness metrics.

Comparison to status quo: Current AI evaluation is almost entirely outcome-accuracy driven; this adds a theory-alignment criterion.

Principle 5 — Redesign datasets to include healthy adaptation, not only episodic disease

Mechanism: Partner with population-health cohorts, primary care networks, and wearable/longitudinal monitoring programs to capture developmental and adaptive trajectories, not only hospital encounters.

Deliverable: A defined minimum proportion of any AI training cohort drawn from non-hospital, longitudinal, healthy-population data, stated as an explicit dataset-composition target.

Comparison to status quo: Population Health initiatives currently focus on prevention messaging and risk stratification, not on supplying training data for organizational models.

Principle 6 — Establish Living System Laboratories

Mechanism: Convert a defined fraction of existing translational research infrastructure into dedicated facilities for testing organizational-coherence hypotheses, with joint appointments spanning medicine, chemistry, complexity science, and mathematics.

Deliverable: At least one funded pilot laboratory within the first strategic cycle (3–5 years), with pre-registered hypotheses and public reporting of negative results.

Comparison to status quo: Existing translational institutes (e.g., Leiden Bio Science Park partnerships) are organized around therapeutic modality, not around testing a general theory of organization.

Part IV — Implementation Roadmap

Stage 1 — Recognition (Year 1)

Formal institutional acknowledgment, in the strategic document itself, that no explicit theory of the organism currently exists. This is a governance act, not a research act: it creates the mandate for Stages 2–4.

Stage 2 — Integration (Years 1–2)

Convene a standing cross-disciplinary working group (medicine, systems biology, chemistry, complexity science, cognitive science, mathematics, philosophy of science). Its first deliverable is not a theory but a shared formal vocabulary — without which "integration" remains a slogan.

Stage 3 — Formalization (Years 2–4)

Translate the working group's shared vocabulary into a mathematical and computational model (Part V). Require pre-registered, falsifiable predictions before any claim of validation.

Stage 4 — Transformation (Years 4+)

Once — and only once — a model has survived falsification attempts, integrate its outputs into diagnostics, prevention, treatment design, education, and AI evaluation standards (Principles 3 and 4 above).

Governance note: Each stage should have an explicit kill criterion — a pre-specified condition under which the institution abandons or revises the approach — to avoid the common failure mode of strategic documents that commit to a direction but never specify what would count as it not working.

Part V — Toward Formalization: A Falsifiable Bioelectric-Quaternion Model of Human Coherence (Revision 2)

5.0 Preface: What This Revision Removes, and Why

The March 2026 quaternion-biofield papers on constable.blog contain two layers that must be separated before this model can be used as the backbone of a strategic mensmodel:

Retained (established or directly testable):

- The historical fact that Maxwell's 1865 equations were formulated in quaternion form, with a scalar potential later discarded by Heaviside's 1884–85 vector reduction.
- Michael Levin's experimentally established findings on bioelectric morphogenesis (L1–L5 in the original paper): resting membrane potential as an actively regulated information variable; DC intercellular field gradients encoding positional information; bioelectric pattern disruption producing

morphological change independent of genetic sequence; whole-body bioelectric pattern as a regulatory scaffold; gap-junction-mediated collective bioelectric states.

- The Kuramoto coupled-oscillator formalism, an established, peer-reviewed mathematical framework.
- A subset of the original paper's predictions (P1, P2, P3, P5, P6) that are genuinely falsifiable with current technology and do not depend on the removed claims below.

Removed (unfalsifiable, esoteric, or unsupported leaps):

- The claim that human interpersonal dynamics are "isomorphic" to Maxwell's field equations, and that Sullivan's Interpersonal Theory constitutes a proven "solution class" of that system. This is a structural analogy, not a demonstrated isomorphism — no parameter of the analogy has been fit to data.
- The DSM-as-field-physics table (depression = "collapse of φ," PTSD = "magnetic remanence," etc.), presented as a definitional mapping rather than as what it currently is: an untested hypothesis.
- Human Design, the Djed-pillar protocol, Kundalini/Sekhem terminology, and the VALIS construct (borrowed from Philip K. Dick's fiction). These are not biomathematical claims and are excluded entirely from this model.
- The relativistic "emotional mass" / spacetime-curvature section, which invokes General Relativity at energy scales where the effect is not merely small but many orders of magnitude below any conceivable measurement.

What remains, after this removal, is smaller than the original papers but stands on its own as a legitimate, falsifiable extension of Levin's bioelectric research — which is precisely the standard this document now holds itself to.

5.0.5 The Organism as a Native Physiological Closed-Loop Control System

Before introducing the quaternion representation, it is necessary to state the organizing principle this entire Part formalizes: **the living organism is already a Physiological Closed-Loop Control (PCLC) system, prior to and independent of any device.** PCLC — sensor, controller, effector, operating on a reference/target state — is not introduced in Section 5.9 as an engineering add-on; it is the native architecture of biological regulation itself, and everything in Sections 5.1–5.4 is a formalization of one instance of it.

Concretely: a bioelectric regulatory loop such as the ones Levin has characterized already has all three PCLC components, without any external hardware. The **sensor** is the tissue's own voltage-sensitive channels and gap-junction network, which continuously register local and neighboring membrane potential. The **controller** is the intracellular and intercellular signaling network that computes a correction relative to a target pattern. The **effector** is the set of ion channels and pumps that adjust membrane potential in response. Homeostasis, the autonomic nervous system, and hormonal feedback are further, larger-scale instances of the same triad. This is not an analogy to engineered PCLC systems — it is the same formal category, realized biologically before it was ever realized in silicon. Sections 5.1–5.4 formalize the state this native PCLC loop regulates (the quaternion $q(t)$), the coupling between multiple such loops (Kuramoto synchronization), and the loop's own reference-tracking behavior (the target-morphology attractor). Section 5.9 then extends the same PCLC framework outward to engineered devices, and Part VII extends it further into speculative whole-body engineered PCLC. The three levels — native biological PCLC, single-organ engineered PCLC, and the proposed supervisory Systemic Coherence Layer — are treated throughout this document as the same architecture operating at different scales, not as three separate ideas connected by analogy.

5.1 Why a Quaternion State Representation, Not a Maxwell Isomorphism

Levin's own published work already establishes that a living tissue's regulatory state has (at minimum) two distinct components with different physical character:

1. A **scalar, non-propagating component** — the DC bioelectric pattern (resting membrane potential distribution) that encodes positional/organizational information and persists as a slowly-changing structural blueprint (L1, L2, L4).
2. A **vector, oscillatory component** — the faster-changing rhythmic activity (neural firing, cardiac rhythm, calcium waves) that Levin and colleagues have separately documented as functionally coupled to the DC pattern (see e.g. the gap-junction/collective-state literature, L5).

A quaternion $q = a + bi + cj + dk$ is, mathematically, nothing more than the minimal algebraic object that jointly represents one scalar magnitude and one three-component vector with a well-defined multiplication rule for composing rotations. This is precisely why quaternions are the standard tool in robotics and orientation-tracking engineering for representing "a magnitude plus a 3D directional state that evolves under composition" — it is not, in that engineering context, taken to imply that a robot arm obeys Maxwell's equations. The same logic applies here: using a quaternion to jointly encode Levin's scalar DC pattern and vector oscillatory state is a modeling convenience justified by the algebra's fit to the data's structure (one scalar + one 3-vector, evolving under rotation-like transformations), not a physical claim that biological tissue is an electromagnetic field obeying Maxwell's equations in the literal sense. This is the correction to the original paper's central overreach: the quaternion is a **state-space representation**, not a **derivation from electrodynamics**.

Formal definition. For a tissue/organ subsystem i, define:

$q_i(t) = \varphi_i(t) + **V_i**(t)$

where $\varphi_i(t)$ is the local scalar bioelectric potential (Levin's DC pattern, directly measurable via voltage-sensitive dyes or electrode mapping) and $\mathbf{V_i}(t) = (v_1, v_2, v_3)$ is a three-component vector summarizing the dominant oscillatory modes at that site (e.g., derived from the first three components of a principal-component decomposition of local time-resolved bioelectric activity — a standard, off-the-shelf method, not a new physical postulate).

5.2 Coupling Between Subsystems: Kuramoto on the Vector Component

Levin's L5 finding — that cells communicate membrane-potential states to neighbors through gap junctions, producing collective bioelectric states — is the direct biological justification for modeling inter-node coupling as a synchronization process. We apply the (established, level-1) Kuramoto formalism to the oscillatory (vector) component of the quaternion state:

$$d\theta_i/dt = \omega_i + K \cdot \Sigma_j \, w_{ij} \cdot \sin(\theta_j - \theta_i)$$

where θ_i is the phase of node i's dominant oscillatory mode (extracted from $\mathbf{V_i}(t)$ via Hilbert transform, as in the earlier RAF-based draft), w_{ij} is the gap-junction/field-proximity-derived coupling strength between nodes i and j, and the order parameter

$$r(t) \cdot e^{(i\psi(t))} = (1/N) \cdot \Sigma_i \, e^{(i\theta_i(t))}$$

again gives a scalar $r(t) \in [0, 1]$ — the degree of macroscopic phase coherence. This step is unchanged from the earlier draft, but it is now anchored to Levin's specific, published gap-junction mechanism rather than to an unspecified "chemical-semantic" coupling.

5.3 Closure, Reformulated via Levin's Target-Morphology Attractor — Not Kauffman RAF Sets

The earlier draft used Kauffman's RAF (reflexively autocatalytic, food-generated) formalism to define "organizational closure." This revision replaces that with a formalism specific to the bioelectric system, because Levin's own experimental work already provides a directly biological, already-validated analogue: the **target morphology attractor**.

Definition. A subsystem's bioelectric pattern $\varphi_i(\cdot)$ is a **closed attractor** over a region R if, following an experimentally induced perturbation $\varphi_i \rightarrow \varphi_i + \delta$, the system's own regulatory dynamics drive φ_i back toward the pre-perturbation pattern (or toward a new stable target pattern that the organism then treats as its new reference state) without continued external forcing. This is not a redefinition on paper — it is an operational description of what Levin's planarian regeneration and Xenopus experiments already demonstrate directly (L3, L4): induced bioelectric states settle back into or are actively steered toward stable target configurations that the tissue itself maintains.

This gives "closure" a concrete, already-measured biological meaning (return to or active pursuit of a target bioelectric pattern) instead of an abstract chemical-network property that had never been connected to Levin's data in the first place.

This is a native PCLC loop, stated in PCLC terms. Per Section 5.0.5, the target-morphology attractor is what a PCLC controller does when it is functioning correctly: it drives a sensed state back toward a reference value using its own effectors, without external intervention. φ_i is the sensed and controlled variable, the reference value is the target pattern, and "closure" as defined above is simply the statement that the native controller is successfully tracking its reference. This reformulation matters because it means closure failure can be decomposed using standard PCLC failure-mode categories (Section 5.4), rather than remaining a single undifferentiated "loss of closure."

5.4 The Combined Coherence Criterion

A region of the organism is in a state of **organizational coherence** at time t if, simultaneously:

(a) its bioelectric pattern φ is within a defined distance of a stable target-morphology attractor (Section 3), and

(b) its oscillatory order parameter $r(t)$ exceeds a threshold $r\backslash$ ***for longer than a null (randomly-rewired coupling) model of the same network would predict (Section 2).***

Disease, under this model, is a sustained departure from (a), from (b), or both — measurable directly from bioelectric mapping and time-series data, not inferred from a downstream biomarker.

This is structurally the same two-condition criterion as the earlier draft (closure + coherence), but every term in it now has a direct, citable empirical referent in Levin's published experimental work, rather than resting on an unconnected abstract network formalism.

Correction: departure from coherence is bidirectional, not only a deficit. The formulation above, as first stated, implicitly treated disease only as $r(t)$ falling too low or φ drifting too far from its target. This is incomplete. Two clinically major categories of disease are, structurally, the opposite failure: **pathologically excessive or misdirected coherence** — $r(t)$ abnormally high and locked around the wrong attractor, rather than failing to reach one. Parkinson's disease is the clearest documented case: pathological beta-band (13–30 Hz) hypersynchrony between the basal ganglia and motor cortex is a well-established feature of untreated disease,

and effective adaptive deep-brain stimulation works specifically by suppressing this excess synchrony — meaning existing closed-loop neuromodulation devices already perform a form of *r(t)*-regulation in clinical practice, without being described this way. Epileptic seizures are the same category at a different timescale: a seizure is not an absence of neural coordination but a pathological, self-reinforcing hypersynchronization that has captured a network away from its normal operating range. The coherence criterion above is corrected accordingly: **disease corresponds to *r(t)* or *φ* departing from the range that supports adaptive function, in either direction — a floor (Type I, coherence deficit) and a ceiling or wrong-attractor lock-in (Type II, pathological hyper-coherence) are both failure states.** Section 5.11 (a case-study companion to this document, referenced but not reproduced in full here) develops this into a five-part disease typology.

Disease as native PCLC failure, decomposed by failure mode. Because Section 5.3 established that closure is native PCLC reference-tracking, a departure from condition (a) or (b) above can be further classified using the standard PCLC failure categories, rather than treated as a single undifferentiated loss of coherence:

- **Sensor failure** — the tissue's own sensing apparatus (e.g., voltage-sensitive ion channels) misreports the local state. Example: a channelopathy that alters channel gating so that the sensed potential no longer reflects the true membrane state.
- **Controller failure** — the signaling network that computes the correction is disrupted, even though sensing and effectors are intact. Example: a disrupted intracellular signaling cascade that fails to translate a correctly sensed deviation into an appropriate corrective instruction.
- **Effector failure** — the correction is correctly computed but cannot be executed. Example: impaired ion pump function that leaves the cell unable to restore the target potential even though the deviation has been correctly detected and a correction correctly computed.
- **Reference failure** — sensing, computation, and execution all function correctly, but the target/ reference value itself has shifted to an abnormal state, and the loop faithfully maintains this abnormal target. This is arguably the most significant category for cancer, per Levin's own framing: a tumor is not necessarily a loop that has stopped working, but a loop that is working correctly toward the wrong target.
- **Node destruction** — the loop's physical substrate is eliminated outright (e.g., autoimmune destruction of a cell population), which is not a failure mode of a functioning loop but the removal of the loop itself, and must be classified separately from the four functional failure modes above.

This decomposition is not decorative: it generates a specific, falsifiable expectation — namely, that interventions matched to the correct failure category (e.g., correcting a reference-value error by resetting the target pattern, rather than forcing the sensed variable directly, or suppressing excess coupling in a Type II case rather than attempting to further stimulate it) should outperform interventions that treat all coherence departures identically. This sharpens the prior closure-restoring-intervention prediction into five distinct, separately testable sub-predictions rather than one.

5.5 Retained, Falsifiable Predictions from the Quaternion-Biofield Papers

Of the six predictions (P1–P6) in the original paper, the following survive scrutiny as genuinely falsifiable and independent of the removed claims:

(P1) Bioelectric signature of stable individual differences. Different individuals should show measurably different, temporally stable whole-body DC bioelectric patterns (*φ*), extractable via multi-electrode body-surface mapping. *Falsification:* if inter-individual variance in mapped *φ* is statistically indistinguishable from intra-individual measurement noise, this is falsified.

(P2) Non-contact induction. A deliberate, biofeedback-induced change in one person's bioelectric state should produce a measurable (if small) change in a nearby second person's bioelectric state, without physical contact, extending Levin's tissue-level gap-junction induction to an interpersonal near-field claim. *Falsification:* absence of any measurable effect above sensor noise at 1–3 m, across repeated trials, falsifies this specific extension (note: this is the single most speculative retained prediction, since it extrapolates a tissue-scale, contact-mediated mechanism to a non-contact interpersonal scale; it should be treated as the lowest-priority, highest-risk item in the validation program).

(P3) Coupling strength predicts rapport. The estimated coupling parameter w_{ij} between two people's oscillatory states (derived from simultaneous HRV/EEG) should correlate with independently rated interaction quality. *Falsification:* no correlation across a adequately powered sample falsifies this.

(P5) Kuramoto transition and group cohesion. Onset of cardiac/respiratory phase synchrony in a group (the Kuramoto transition, i.e., $r(t)$ crossing a critical threshold) should coincide with independently measured onset of group cohesion. *Falsification:* no temporal association falsifies this.

(P6) Scalar potential and depression severity. Depression severity on standard clinical scales should correlate inversely with the amplitude of the DC (scalar) component of the biofield, independent of oscillatory components. *Falsification:* no such correlation, or a correlation fully explained by oscillatory components alone, falsifies this specific claim.

(P8) Coherence topology predicts epigenetic and twin-phenotypic divergence. In monozygotic twins, (a) epigenetic marking patterns should correlate with coherence-profile measurements (HRV, EEG-coherence mapping) independent of documented environmental history, and (b) the phenotypic dimensions on which twins diverge most should correlate with the domains in which their coherence topologies diverge most. This follows from treating the genome as a shared, largely fixed information store ("library") and the individual bioelectric field as the variable "reader" that determines which parts of it are expressed (consistent with Levin's DNA-independent morphological control, L1–L4, and with the established finding that monozygotic twins diverge epigenetically over the lifespan). *Falsification:* if epigenetic divergence and coherence-topology divergence show no association beyond what is explained by known environmental exposures in a twin cohort, this is falsified. *Note:* this prediction is testable with existing, low-cost technology (standard HRV/EEG plus routine epigenome assays on existing or new twin cohorts) and does not require any of the closed-loop device infrastructure discussed in Section 5.9 — it is the nearest-term, lowest-cost validation opportunity in this document.

(P9) Failure-mode-matched intervention outperforms undifferentiated correction. Following the PCLC failure-mode decomposition in Section 5.4, an intervention matched to the correct failure category (sensor, controller, effector, or reference) should produce more durable recovery than an intervention that corrects the sensed variable directly without diagnosing which component failed — holding severity and condition constant. A specific sub-case, directly derived from Levin's own published framing of cancer as a reference-value error: normalizing the bioelectric target/reference of the surrounding tissue should outperform

interventions that act only on the tumor's sensed or effector-level state. *Falsification:* if outcomes are statistically indistinguishable between failure-mode-matched and undifferentiated interventions in matched cohorts, this is falsified.

Explicitly dropped: P4 (hysteresis in trauma) is not retained as stated, because "magnetic remanence" is a literal claim about ferromagnetic-style hysteresis in biological tissue that has no established physical mechanism connecting it to psychological trauma; a testable descendant of this idea would require first establishing *any* measurable hysteresis-like property in bioelectric tissue at all, which is not yet done. This is flagged as a possible future research question, not a current prediction.

A note on a related but separate claim (codon algebra). A recurring, independently documented structural analogy exists between the 64-unit algebra of Rowlands' nilpotent fermion formalism and the 64-codon structure of the genetic code (matching groupings of 4, resolution into 20, chirality, five-fold symmetry), described across multiple works by Rowlands and Hill (*Zero to Infinity*, 2007; "Nature's Code"; "A Mathematical Representation of the Genetic Code"). This is a genuine, repeatedly demonstrated mathematical pattern, not an unsupported assertion — but it is worth noting precisely that Rowlands and Hill's own papers consistently describe it as an "analogy" or "correspondence" ("analogous to," "corresponds closely to," "an almost perfect pattern suggests"), not as a proven formal identity. This document adopts the same standard: the pattern is retained as a documented structural analogy worth further investigation, not as an established identity between genetics and particle physics.

5.6 Explicit Status Table

| Component | Status | Basis |

|---|---|---|

| Organism as native PCLC system (Section 5.0.5) | Established engineering category applied to established biology | PCLC literature (biomedical engineering); homeostasis/ANS physiology |

| Quaternion state representation (φ + V) | Modeling convention, internally consistent | Justified by data structure (scalar + 3-vector), not by physical isomorphism |

| Kuramoto coupling on oscillatory component | Established mathematics | Kuramoto (1984), Strogatz-Mirollo (1990) |

| Gap-junction-mediated coupling mechanism | Established biology | Levin lab, L5 |

| Target-morphology attractor as "closure" | Established biology, reframed as formal criterion | Levin lab, L3–L4 |

| Predictions P1, P3, P5, P6 | Falsifiable, untested | This document |

| Prediction P2 (non-contact induction) | Falsifiable, untested, high-risk | This document — treat as lowest priority |

| Prediction P4 (trauma hysteresis) | Not currently falsifiable as stated | Removed; needs prior mechanism study |

| Prediction P8 (coherence topology vs. epigenetic/twin divergence) | Falsifiable, untested, lowest-cost near-term validation | This document, built on Fraga et al. 2005 (twin epigenetic divergence) |

| Prediction P9 (failure-mode-matched intervention) | Falsifiable, untested | This document, built on PCLC failure taxonomy + Levin's cancer-as-reference-error framing |

| DNA-as-library / field-as-reader (personal blueprint) | Reframing of established biology as formal criterion | Levin lab, L1–L4; Fraga et al. 2005 |

| Rowlands-Hill codon/nilpotent-algebra structural analogy | Documented, repeated structural analogy — not a proven identity | Rowlands & Hill, *Zero to Infinity* (2007) and related papers; their own language is "analogous to," not "identical to" |

| Maxwell-Sullivan isomorphism | Not retained | Unsupported structural analogy |

| DSM-as-field-physics table | Not retained | Unsupported structural analogy |

| Human Design / Djed / Kundalini / VALIS | Not retained | Outside biomathematics; esoteric/fictional sourcing |

| Relativistic "emotional mass" | Not retained | Effect size below any measurable threshold at biological energy scales |

5.7 Relation to the Broader Strategic Document

This model replaces the RAF-based "chemical-semantic" formalism in Part V of the LUMC strategy document with a model built on: (a) Levin's bioelectric research — independently identified as the most advanced existing biomathematical framework — and (b) a quaternion state representation used strictly as a data structure, not as a claim about electrodynamic isomorphism. The coherence index $r(t)$ and the closure criterion from the original Part V (Section 5.2–5.3) carry over unchanged in form; only their empirical grounding changes, from an untested chemical-network formalism to Levin's directly measured bioelectric attractor dynamics. The computational implementation path (Part V, Section 5.5) is unchanged in structure but should now cite Levin-lab bioelectric mapping protocols directly rather than generic RAF-detection software.

5.8 Explicit Limits (Carried Forward and Extended)

This revision does not yet address: (a) how φ and $\mathbf{V}$ should be weighted relative to each other when they conflict (e.g., a stable target morphology with degraded oscillatory coherence — which dominates?); (b) whether the quaternion multiplication structure (non-commutativity,

specifically) does any explanatory work here, or whether a simpler 4-tuple (φ, v_1, v_2, v_3) without quaternion algebra would perform identically — this should be tested directly rather than assumed; (c) validation is proposed only at tissue/individual/small-group scales, not at the organizational or societal scales the original papers gestured toward; (d) prediction P2 in particular carries a real risk of non-replication and should not be treated as more established than it is.

5.9 The Systemic Coherence Layer: Coupling to Multi-Device Closed-Loop Trajectories

This section connects the model to an independent, market-driven technology trajectory rather than deriving its motivation from theory alone.

The trajectory. Closed-loop medical devices are developing along an identifiable four-stage path. Stage 1 (current, scaling now) consists of single-organ, single-biomarker closed-loop devices: deep-brain stimulation reacting to one pain or tremor signal, insulin pumps, theranostic oncology agents that both locate and treat one tumor, microneedle patches that sense and dose one biomarker. Stage 2 (current research frontier, 2026) has already moved beyond single devices toward coordination within one organ system: published work on multi-device collaborative theranostics operating across tissue depth, and organ-system digital twins — an ECMO digital twin coordinating gas exchange and circulatory support together, a cardiovascular digital twin (TwinCardio) synchronizing prediction and monitoring for one organ system. Stage 3, not yet published as such but a direct extrapolation of Stage 2, is the point at which a single patient carries multiple organ-system digital twins simultaneously (e.g., a cardiac twin and a neuromodulation twin), creating a coordination problem no single twin can solve on its own: an intervention optimized for one organ's local setpoint can destabilize another, absent any shared cross-organ state variable. Stage 4, further out and more speculative, is the shift from reactive correction of individual biomarkers to preventive restoration of a system-wide coherence measure before any single biomarker crosses a clinical threshold.

Where the model fits. Stage 3's coordination problem is not a new problem invented for this document — it is the direct consequence of proliferating single-purpose closed-loop devices, and it requires exactly the kind of system-wide state variable this document has already formalized. Each individual organ-system digital twin already produces, without calling it this, a local state of the same form used in Section 5.1: one slowly-changing scalar regulatory setpoint (φ_i, e.g. a target blood pressure or glucose level) plus a faster oscillatory component ($\mathbf{V}_i$, e.g. heart rhythm, stimulation frequency, secretion rhythm). This means every clinically deployed closed-loop twin is, structurally, already a node $q_i(t) = \varphi_i(t) + \mathbf{V}_i(t)$ in the sense of Section 5.1 — the representation was not designed for this purpose, but the data it operates on already has this shape.

The coordination mechanism. The Kuramoto coupling of Section 5.2 extends naturally to this setting: instead of gap-junction-derived coupling strengths between cells, w_{ij} is derived from known physiological interdependence between organ systems (e.g., the established coupling between oxygen saturation and heart rate). The resulting order parameter $r(t)$, computed jointly across a patient's active closed-loop devices, becomes a shared coordination signal that today does not exist between independently operating digital twins.

The safety criterion. The closure/target-morphology criterion of Section 5.3 translates directly into a cross-device safety condition: a combination of closed-loop devices is coherent only if the coupled system as a whole returns to a stable joint state after any single device's intervention — not merely to that device's own local setpoint. This gives a concrete, computable safeguard against the specific failure mode Stage 3 introduces: one device's local optimization destabilizing another's.

A new falsifiable prediction (P7), specific to this coupling. In patients carrying two or more simultaneous closed-loop devices, a jointly computed *r(t)* across both devices should detect impending complications (e.g., hemodynamic instability under combined ECMO and neuromodulation therapy) earlier and with fewer false alarms than the two devices monitoring their own independent thresholds. *Falsification:* if joint *r(t)* monitoring shows no predictive advantage over independent per-device thresholds in a retrospective multi-device cohort, this specific application of the model is falsified.

Architectural placement. This does not require a new device. It is proposed as a **Systemic Coherence Layer**: a data-integration layer that reads the state streams of existing, independently regulated closed-loop devices, computes the joint *r(t)* and closure estimate, and feeds this back to each device's own controller as an additional input — without replacing any device's existing certified control logic. This has a direct strategic consequence for Part I of this document: because it is a monitoring and coordination layer over already-approved devices rather than a new high-risk AI medical device in its own right, it faces a materially lower regulatory barrier under the EU AI Act's high-risk classification than a novel standalone closed-loop system would.

Relation to Physiological Closed-Loop Control (PCLC). The sensor-controller-effector architecture assumed throughout this section is not a new engineering category invented for this document — it is the established formal framework already used to classify implanted and wearable closed-loop medical devices in the biomedical engineering literature, where such systems are reviewed as a distinct class (deep-brain stimulators, closed-loop insulin delivery, and related devices are treated as instances of a common PCLC architecture). Positioning the Systemic Coherence Layer explicitly as a PCLC-class system — specifically, a supervisory layer coordinating multiple lower-level PCLC controllers rather than a first-order controller itself — connects Section 5.9 to an existing body of safety, verification, and regulatory literature on closed-loop medical automation, rather than requiring a new engineering category to be defined and validated from scratch.

Explicit limit. This section describes a plausible integration point between the coherence model and an independently observed technology trend; it does not itself constitute evidence that the coherence model is correct. The trajectory from Stage 2 to Stage 3 is an extrapolation, not a documented industry roadmap, and should be labeled as such in any external use of this document.

5.10 The Personal Blueprint: DNA as Shared Library, Field as Individual Reader

This section formalizes a related concept — that individual and family-level variation arises not from genetic sequence differences alone, but from differences in which parts of a largely

shared genetic "library" are actively expressed — using the same apparatus as Sections 5.1–5.4, and ties it to prediction P8 above.

Formal statement. Let L denote the genetic information store (DNA) of an individual, treated here as a largely fixed resource shared within a family or species. Let $q_i(t) = \varphi_i(t) + V_i(t)$ be the bioelectric field state of subsystem i, as defined in Section 5.1. Define the **expression-selection function** $E(L, q(t))$ as the (currently unspecified, empirically determinable) mapping from library plus field state to actual gene expression at time t. The central claim — consistent with Levin's findings (L1–L4) that morphology can be altered by field manipulation without altering L — is that phenotypic variation between individuals sharing similar or identical L (e.g., monozygotic twins) is substantially explained by variation in $q(t)$ at the point of developmental individuation, not by variation in L itself.

Why this is not a new claim, only a new framing. This does not add a new causal mechanism to the model — it restates the closure/coherence apparatus of Sections 5.1–5.4 applied to gene expression selection instead of morphological attractor states. It is included as a separate section because it generates the specific, near-term-testable prediction P8, which is otherwise easy to overlook among the tissue-scale predictions (P1–P6).

Relation to reduced genetic diversity in closed populations. A secondary, weaker implication is that populations with historically reduced genetic diversity (founder effects in isolated communities) have a narrower library L, which — independent of any claim about the field — is already an established genetic mechanism (reduced redundancy lowers the chance that a second, unaffected gene copy can compensate for a harmful variant). This document does not extend the coherence-field claim to this mechanism; it notes only that the two are compatible and should not be conflated: reduced library diversity is standard population genetics, while the field-as-reader claim is the separate, less-established part of this section that P8 is designed to test.

Explicit limit. The expression-selection function $E(L, q(t))$ is currently a placeholder for an unspecified relationship, not a derived formula. This section identifies what would need to be measured (coherence topology plus epigenetic pattern, per P8) to begin constraining its form; it does not yet propose that form.

Part VI — Global Market and Technology Context

This part consolidates the market and technology research gathered separately during this project. It exists to answer a question Parts I–V do not: not "what should the model be," but "what environment will any such model have to survive in."

6.1 Financial and Structural Pressure

Global healthcare is shifting from "innovation at any cost" to "survival through efficiency": health systems are consolidating and adopting AI specifically to offset record medical inflation and labor shortages, not primarily to pursue scientific advancement. OECD countries spent an average of 9.3% of GDP on health in 2024, above pre-pandemic levels,

even as economies slowed. The financial logic of the sector has shifted from volume-based growth to "tech-enabled margin expansion." Any strategy — including the one in this document — must be positioned within, not against, this efficiency pressure: Part III's deliverables were deliberately written to require no new hospital infrastructure and to attach to existing budgets rather than compete with them.

6.2 AI Regulation and Infrastructure

Hospitals are moving from simple AI chatbots toward autonomous "AI agents" handling scheduling, prior authorization, and documentation, while consolidating around unified platforms (e.g., Epic, Oracle Health) rather than niche point solutions. Regulation is simultaneously tightening and being renegotiated: the EU AI Act classifies AI used for diagnosis, clinical decision support, treatment recommendation, triage, and patient monitoring as high-risk, with a political agreement (May 2026) postponing enforcement of high-risk obligations to December 2027 (standalone systems) and August 2028 (systems embedded in regulated medical products). Approximately 75% of currently deployed commercial AI medical devices are concentrated in radiology — meaning most existing regulatory and deployment experience does not transfer directly to a systemic, whole-body model of the kind proposed in Part V. The Systemic Coherence Layer proposed in Section 5.9 was deliberately designed as a data-integration layer over already-certified devices specifically to reduce exposure to this regulatory bottleneck.

6.3 Workforce Shortage

WHO projects a shortfall of roughly 11 million health workers by 2030, concentrated in low- and middle-income countries, but felt everywhere: 40% of UK general practitioners expect to retire or leave within five years. This is structural, not cyclical, pressure toward automation and toward moving diagnostic/monitoring burden onto devices and AI — reinforcing, independently of any scientific argument, the strategic case for Section 5.9's device-coordination layer and Part III's Principle 4 (AI evaluated on organizational validity, not accuracy alone).

6.4 Care Moving Outside the Hospital

Surgery and chronic care are shifting to ambulatory centers and retail clinics, and home healthcare is projected to exceed USD 700 billion, driven by over 13,000 active startups. This directly challenges the traditional academic hospital's role as the primary site of care and data collection — reinforcing Part III's Principle 5 (redesigning datasets to include non-hospital, longitudinal, healthy-population data) as a matter of practical necessity, not only scientific completeness.

6.5 Relevant Scientific Breakthroughs

Personalized gene therapies are entering early clinical use, accelerated by AI tools (e.g., Stanford's CRISPR-GPT, mid-2025) that compress experiment design from years to months; regenerative medicine and 3D bioprinting are advancing tissue repair and engineering. This is the applied context into which Levin's bioelectric findings (the empirical foundation of Part

V) are increasingly being received — the field is actively searching for mechanisms beyond genetic sequence alone, which is the same gap this document's model addresses.

6.6 Geopolitical and Funding Volatility

U.S. science funding has become structurally less predictable, creating downstream uncertainty for university-industry partnerships and international clinical trial networks; development finance institutions are taking on a larger role globally as traditional aid budgets contract. This affects the funding environment for the multi-year formalization program proposed in Part IV, Stage 3, and argues for the staged, kill-criterion-based funding structure already specified there rather than a single long-horizon commitment.

6.7 Global Health Trajectory

The WHO's 2026 World Health Statistics report finds that global progress toward the health-related Sustainable Development Goals is uneven, slowing, and in places reversing: universal health coverage service-index growth has nearly stalled (68 to 71 between 2015–2023), and 1.6 billion people were pushed into or further into poverty by out-of-pocket health costs in 2022 — alongside real gains, such as a 40% fall in new HIV infections since 2010. This unevenness is a caution against a strategy that assumes uniformly rising resources: Part IV's staged approach and Part III's emphasis on deliverables usable without new infrastructure are consistent with, rather than contrary to, this constrained global picture.

6.8 Synthesis: What the Market Context Implies for This Document

None of Sections 6.1–6.7 argue for or against the scientific content of Part V. What they establish is a timing and framing constraint: the sector is currently allocating resources toward efficiency, compliance, and workforce substitution, not toward foundational theory. A strategy that requires new funding for pure theory-building will compete poorly for attention in this environment. The strategy in Part III was written to avoid this trap by tying every principle to an existing budget line or existing regulatory process; Section 5.9's Systemic Coherence Layer and Section 5.5's Prediction P8 were both selected, in part, because they are inexpensive to test using infrastructure or cohorts that already exist. This is a deliberate response to the market context in this Part, not a coincidence.

Part VII — Speculative Device Trajectory (Not Part of the Falsifiable Core)

This entire Part is explicitly quarantined from the falsifiable model in Part V, in the same way Section 5.0 quarantined the Human Design/Djed/ Kundalini/VALIS material. Nothing here is a prediction of the model; it is a forward-looking engineering narrative built on top of the model, included because it was directly requested, and it should be read, cited, and used accordingly — as speculation about where the field could go, not as a claim about what the model shows.

7.1 Why Include This

Parts III–VI describe what is defensible now. This Part answers a different, legitimate question: if the trajectory described in Section 5.9 (Stages 1–4 of closed-loop device coordination, itself framed as an instance of the established Physiological Closed-Loop Control (PCLC) engineering category) continues, what is the most advanced device concept consistent with the model's own logic, and what would have to be true for it to exist? Labeling this speculative does not mean it is arbitrary — each step below is anchored either to an existing technology or an identified, named research gap — but it means the document does not present it as validated or as a deliverable with a kill criterion, unlike Parts III–V. Phase D, in particular, should be read as a hypothetical whole-body extension of PCLC architecture, not as a departure from it.

7.2 Development Path

Phase A (now–~2029): Local q(t) from combined existing sensors. Existing point sensors (ECG, EEG, bio-impedance, galvanic skin response) are combined in software into a single local quaternion estimate per body region. This is the Systemic Coherence Layer of Section 5.9, requiring no new hardware — only integration of what already exists.

Phase B (~2029–2034): From point measurement to field reconstruction. A dense sensor mesh — extending existing bioadhesive ultrasound patches and microneedle arrays — continuously reconstructs the bioelectric field over a body surface as a spatial map $\varphi(x,t)$, rather than a handful of discrete electrode values. Technically, this is an extension of existing EEG source-localization algorithms from the head to the whole body.

Phase C (~2034–2040): From measuring to remote, deep correction. Temporal interference stimulation — using two high-frequency external fields that interfere so that only their intersection, deep inside tissue, produces a functional low-frequency stimulation signal — already exists for non-invasive deep-brain targeting. The speculative extension: the same interference approach, applied from multiple external sources simultaneously and guided by the Phase B field map, to "write" a target φ at an arbitrary depth in an arbitrary organ, not only the brain.

Phase D — the speculative "superdevice" (~2040+): A closed, whole-body bioelectric control system. A non-invasive or minimally invasive mesh (skin-patch level, potentially supplemented by a small number of ingested or injected microscale nodes) that simultaneously (1) continuously reconstructs the full quaternion field $q(x,t) = \varphi(x,t) + V(x,t)$ over the body, (2) computes $r(t)$ and closure-distance to the target attractor in real time (Sections 5.3–5.4), and (3) delivers targeted correction via multi-point interference fields to the specific regions where coherence is declining — before any single biomarker crosses a clinical threshold. In effect, this reverses the single-organ closed-loop paradigm of Section 5.9's Stage 1: one system, all organs, driven by coherence rather than by a single biomarker.

7.3 Breakthroughs Required, in Order of Urgency

3. **High-resolution, non-invasive, wearable field imaging at depth.** Optically-pumped magnetometers already enable wearable MEG at head level; the breakthrough needed is extending this to whole-body, chronic-use cost and form factor.
4. **Fast inverse source reconstruction suitable for closed-loop use, not offline analysis.** EEG source localization is currently a typically offline computation; needed is the same class of dedicated low-

power AI chip already used for closed-loop seizure detection, generalized to whole-body field reconstruction.

5. **A validated inter-organ coupling matrix w_{ij}.** This does not currently exist in any systematic form — no one has measured how the bioelectric states of different organ systems influence each other at the resolution this would require. This is itself a foundational research program, not an engineering task, and is likely the single largest bottleneck of the six listed here.
6. **Deep, targeted non-invasive actuation outside the brain.** Temporal interference has been demonstrated almost exclusively for brain tissue; extension to other organs (heart, gut, liver) is unproven and would require independent validation for each.
7. **A validated, measurable closure criterion at the tissue level beyond Levin's regeneration experiments.** This is the gap already flagged in Section 5.8 — without it, "correcting toward the attractor" remains an assumption for organs beyond those Levin has directly studied.
8. **Biocompatible, long-term-stable sensor/actuator materials** for daily chronic use without rejection or signal degradation.

7.4 Explicit Status

This Part is a narrative extrapolation, not a research plan with deliverables or kill criteria. It should not be cited as evidence that the model in Part V is correct, nor used in funding or regulatory contexts as a roadmap without independent technical review. Its function in this document is to show what the falsifiable core in Part V would imply at the limit, if every open question in Sections 5.8 and 7.3 were resolved favorably — which is not assumed.

Conclusion

Medicine has entered the age of Artificial Intelligence without first resolving one of its oldest questions: what is the living system it exists to preserve? The sharpened strategy above turns that observation into a sequenced, falsifiable program with named deliverables and kill criteria. Part V's central organizing principle, stated explicitly in Section 5.0.5, is that the organism is already a native Physiological Closed-Loop Control (PCLC) system — sensor, controller, effector, tracking a reference state — independent of any device; Sections 5.1–5.4 formalize this native loop's state (the quaternion representation), its coupling across subsystems (Kuramoto synchronization), and its reference-tracking behavior (the target-morphology attractor), and reframe disease itself as a specific, classifiable PCLC failure mode (sensor, controller, effector, or reference failure) rather than an undifferentiated loss of coherence, yielding Prediction P9. Section 5.9 extends the same PCLC architecture outward to engineered closed-loop devices, positioning the proposed Systemic Coherence Layer as a supervisory PCLC coordinating both native and artificial PCLC subsystems within a single, already-established engineering category — giving the model a market-driven rationale in Part VI in addition to its scientific one. Section 5.10 extends the same apparatus to individual and familial variation via the "personal blueprint" concept, yielding Prediction P8 — the lowest-cost, nearest-term test in this document. Part VI situates all of this within the actual financial, regulatory, and workforce pressures facing academic medical centers worldwide, and shows that the strategy's design — deliverables attached to

existing budgets, a coordination layer over already-certified devices, tests using existing cohorts — is a deliberate response to that environment rather than an assumption of favorable conditions. Part VII extends the model's own logic into a speculative long-term device trajectory, explicitly quarantined from the falsifiable core in the same way the esoteric material was quarantined in Section 5.0, so that ambition and evidence remain visibly separate throughout the document.